



Silver Peak Unity Orchestrator

Operator's Guide

Orchestrator 8.0
March 2016
PN 200095-001 Rev P

Silver Peak Unity Orchestrator Operator's Guide

Document PN 200095-001 Rev P

Date: March 2016

Copyright © 2016 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, Suite 100
Santa Clara, CA 95050

1.877.210.7325 (toll-free in USA)
+1.408.935.1850

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

Contents

Preface	ix
Who Should Read This Manual?	ix
Manual Organization	ix
Support	x
Chapter 1 Getting Started	1
Overview	2
Completing the Orchestrator's Getting Started Wizard	2
Assumptions	3
One more thing	3
What to Configure Next in a WAN Optimization Network	4
Understanding Topology and Layout	6
Alarms	6
Topology Settings & Legend	7
Other	8
Managing Orchestrator User Accounts and Authentication	9
Adding to the Subnet Table	11
Chapter 2 Unity Overlays	13
Introduction to Unity Overlays for SD-WAN	14
Discovered Appliances	16
Deployment Profiles	17
Mapping Labels to Interfaces	18
LAN-side Configuration: DHCP	18
WAN-side Configuration	18
Definitions	19
DHCP Server Definitions	19
DHCP Relay Definitions	19
A More Comprehensive Guide to Basic Deployments	19
Bridge Mode	20
Router Mode	21
Server Mode	24
How You Can Adjust the Basic Deployments	25
Configuring Gigabit Etherchannel Bonding	25
Adding Data Interfaces	26
Business Intent Overlays	27
Topology	28
Overlay Policy	28
Best Practices	28
Apply Overlays	29
Interface Labels	30
Licenses	31
Silver Peak Cloud Portal	32
DHCP Server	33
DHCP Server Definitions	34
DHCP Relay Definitions	34

IPSec Pre-shared Key Rotation	35
Configuration Wizard	36
Chapter 3 Configuration Templates	39
Using Configuration Templates	41
System Template	42
Tunnels Template	44
Shaper Template	46
Dynamic Rate Control	47
User Defined Apps Template	50
Application Groups Template	52
Access Lists Template	53
Route Policies Template	55
QoS Policies Template	57
Handling and Marking DSCP Packets	58
Applying DSCP Markings to Optimized (Tunnelized) Traffic	58
Applying DSCP Markings to Pass-through Traffic	59
Optimization Policies Template	61
TCP Acceleration Options	63
NAT Policies Template	66
Advanced Settings	67
SSL Certificates Template	69
SSL CA Certificates Template	71
SSL for SaaS Template	72
Threshold Crossing Alerts Template	74
Metrics and Defaults	75
Auth/Radius/TACACS+ Template	76
SNMP Template	78
NetFlow Template	80
DNS Template	81
Logging Template	82
Banner Messages Template	84
Cloud Portal Registration Template	85
SaaS Optimization Template	86
VRRP Template	88
Definitions (alphabetically)	88
CLI Template	89
Session Management Template	90
Default Users Template	91
Date/Time Template	93
Chapter 4 System, Network, and Policy Configuration Tabs	95
Deployment Tab	96
Interfaces Tab	98
Bridge Interfaces Tab	99
Tunnels Tab	100
Troubleshooting	100
Advanced Tunnel Options	101
Definitions (alphabetically)	101
Tunnel Groups Tab	104
Topology	104

Interfaces	105
Shaper Tab	106
Subnets Tab	108
SSL Certificates Tab	110
SSL CA Certificates Tab	112
SSL for SaaS Tab	113
VRRP Tab	115
WCCP Tab	117
Route Policies Tab	120
QoS Policies Tab	122
Optimization Policies Tab	124
Access Lists Tab	126
User Defined Applications Tab	127
Application Groups Tab	128
NAT Policies Tab	129
Advanced Settings	131
SaaS Optimization Tab	132
Threshold Crossing Alerts Tab	134
Chapter 5 Appliance Administration Tabs	137
Date/Time Tab	138
Domain Name Servers (DNS) Tab	139
SNMP Tab	140
NetFlow Tab	141
Logging Tab	142
Appliance User Accounts Tab	144
Auth/RADIUS/TACACS+ Tab	145
Banners Tab	147
Chapter 6 Alarms	149
Understanding Alarms	150
Categories of Alarms	150
Types of Appliance Alarms	151
Viewing Alarms	161
Specifying Alarm Recipients	163
Chapter 7 Monitoring Status and Performance	165
About Reports	166
Types of Reports	166
Interpreting Charts	166
Line Charts	166
Bar Charts	167
Configuring and Distributing Custom Reports	168
Data Collection & Management	169
Viewing Appliance Statistics	170
Health Dashboard	171
Appliance Data Transfer & Reduction	172
Appliance Max Bandwidth	172
Appliance Bandwidth Utilization	173
Appliance Bandwidth Trends	173
Appliance Bandwidth Cost Savings	174

Calculations	174
Appliance Flow Count	175
Appliance Packet Count	175
Viewing Application Statistics	176
Application Reduction	176
Application Pie Charts	177
Application Trends	178
DSCP Reduction	179
Traffic Class Reduction	179
Viewing Tunnel Statistics	180
Tunnel Data Transfer & Reduction	180
Tunnel Bandwidth Trends	181
Tunnel Max Bandwidth	181
Tunnel Bandwidth Utilization	182
Latency	182
Latency Trends	183
Loss	183
Loss Trends	184
Out of Order Packets	184
Out of Order Packets Trends	185
Tunnel Flow Count	186
Tunnel Packet Count	186
Tunnels Summary	187
Viewing Flows	188
How Flows Are Counted	188
How Flows are Organized	189
Customizing Which Columns Display	191
Flow Details	192
Error Reasons for TCP Acceleration Failure	196
Error Reasons for CIFS Acceleration Failure	199
Error Reasons for SSL Acceleration Failure	200
Error Reasons for Citrix Acceleration Failure	203
Resetting Flows to Improve Performance	205
Monitoring Status & Reporting	206
View Reports	206
Scheduled & Historical Jobs	207
Realtime Charts	208
Historical Charts	208
Reachability Tab	209
Chapter 8 Orchestrator Administration	211
Viewing Orchestrator Server Information	212
Restart, Reboot, or Shutdown	212
Managing the Orchestrator Server License	212
Managing Orchestrator Users	213
Guidelines for Creating Passwords	213
User Menu Access	214
Remote Authentication	216
Debug Files	217
Silver Peak Cloud Portal	218
SMTP Server Settings	219

Overlay Manager Settings	220
Schedule Timezone	221
Audit Logs	222
Getting Started Wizard	223
Proxy Configuration	224
Managing Orchestrator Software	225
Checking for Orchestrator and Appliance Software Updates	225
Upgrading Orchestrator Software	225
Switching Software Versions	226
Backing Up the Orchestrator Database	227
Backing Up on Demand	227
Scheduling Orchestrator Database Backup	227
Chapter 9 Maintenance and Support	229
Viewing System Information	230
Software Versions	231
Upgrading Appliance Software	232
Backing Up Appliance Configuration Files	233
Restoring a Backup to an Appliance	234
Viewing Configuration History	235
Disk Management	236
Synchronizing Appliance Configuration	237
Putting the Appliance in System Bypass Mode	238
Broadcasting CLI Commands	239
Testing Link Integrity	240
TCPERF Version 1.4.8	241
Erasing Network Memory	245
Rebooting or Shutting Down an Appliance	246
Scheduling an Appliance Reboot	247
Scheduling QoS Map Activation	248
Managing Tech Support Files	249
Logging in to the Support Portal	251
Appendix A TCP/IP Ports Used by the Orchestrator and Silver Peak Appliances	253
List of ports used by the Orchestrator	254
List of ports used by Silver Peak Appliances	255
Data Plane	255
Management Plane	255
Diagrams of TCP/IP Port Use	256



Preface

Silver Peak's Unity Orchestrator provides simplified appliance configuration for rapid, large-scale deployment of Silver Peak appliances in your network.

Who Should Read This Manual?

Anyone who wants to centrally manage Silver Peak appliances should read this manual. Users should have some background in Windows[®] terminology, Web browser operation, and a knowledge of where to find the TCP/IP and subnet mask information for your system.

Manual Organization

This section outlines the chapters and summarizes their content.

[Chapter 1, “Getting Started,”](#) provides an overview of the Unity Orchestrator's functions and features and a summary of the tasks for getting started.

[Chapter 2, “Unity Overlays,”](#) describes the screens used for configuring SD-WAN overlays and provides an overview of the workflow.

[Chapter 3, “Configuration Templates,”](#) describes how to use the **Configuration** templates to manage appliances.

[Chapter 4, “System, Network, and Policy Configuration Tabs,”](#) describes the reports that display appliance configuration parameters.

[Chapter 5, “Appliance Administration Tabs,”](#) describes the reports that display appliance administration parameters.

[Chapter 6, “Alarms,”](#) describes alarm categories and definitions. It also describes how to configure, view, and handle alarm notifications.

[Chapter 7, “Monitoring Status and Performance,”](#) focuses on reports related to the status and performance of appliances, applications, and tunnels.

[Chapter 8, “Orchestrator Administration,”](#) describes the administrative tasks that directly relate to managing **Orchestrator-related events and tasks only**. These activities do not relate to managing appliances.

[Chapter 9, “Maintenance and Support,”](#) describes the activities and tools related to maintaining the appliances. This includes database and software image management, as well as reboot operations.

[Appendix A, “TCP/IP Ports Used by the Orchestrator and Silver Peak Appliances,”](#) uses tables and diagrams to list the ports that the Orchestrator and Appliance Manager use for TCP/IP.

Support

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

- **1.877.210.7325 (toll-free in USA)**
- **+1.408.935.1850**
- **www.silver-peak.com/support**

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

- If you have suggestions or feedback for our documentation, please send an e-mail to **techpubs@silver-peak.com**.
- If you have comments or feedback about the GUI's ease of use, please send an e-mail to **usability@silver-peak.com**.



Getting Started

The Orchestrator allows you to manage either an *SD-WAN network* or a *WAN Optimization network*. This chapter outlines the typical tasks involved in setting up the Orchestrator and using it to monitor and manage your Silver Peak appliances.

In This Chapter

- **Overview** See page 2.
- **What to Configure Next in a WAN Optimization Network** See page 4.
If you're setting up an SD-WAN network, skip this subheading and refer to *Chapter 2, "Unity Overlays,"* for setup and configuration guidance.
- **Understanding Topology and Layout** See page 6.
- **Managing Orchestrator User Accounts and Authentication** See page 9.
- **Adding to the Subnet Table** See page 11.

Overview

The Orchestrator enables you to globally monitor performance and manage Silver Peak appliances, whether your focus is configuring primarily for WAN optimization or for .

This section discusses the following:

- **Completing the Orchestrator's Getting Started Wizard** See page 2.
- **Assumptions** See page 3.
- **What to Configure Next in a WAN Optimization Network** See page 4.
- **Understanding Topology and Layout** See page 6.

Completing the Orchestrator's Getting Started Wizard

After you first install the Orchestrator and use a web browser to go to the IP address you've assigned it, its **Getting Started Wizard** appears.

It takes you through the basics of configuring the following:

- **Orchestrator Name, management IP address, and password**
 - The default for username and password is **admin**.
- **License and Registration**
 - EdgeConnect registration is required for Cloud-based features and products, including CPX and SaaS. The associated **Account Name** and **Account Key** enable the Orchestrator to discover EdgeConnect appliances via the Silver Peak Cloud Portal, as they're added to your network.
 - If you have NX, VX, and VRX appliances, you will also have an Orchestrator License.
- **Date/Time**
 - Silver Peak strongly recommends using an NTP server so that data across the Orchestrator and appliances is synchronized.
- **Email**
 - Change the default settings to your Company's SMTP server, and then test.
 - Separate fields are provided for **Global Report** recipients and **Alarm** recipients.

- **Add Appliances**
 - [Optional] You can use this now to add NX, VX, and VRX appliances that are *already* up and running in your network. Or you can add them later.
- **Backup**
 - Specifies the database backup destination, transfer protocol, and backup schedule.

If you don't **Apply** the configuration after you complete the last screen, the Orchestrator's wizard reappears at the next login.

To access the Orchestrator wizard again after initial configuration, go to **Orchestrator Administration > Getting Started Wizard**.

Assumptions

The assumptions here are as follows:

- Any appliance that you add has already been deployed with Appliance Manager, either *in-line* (Bridge mode) or *out-of-path* (Router or Server¹ modes).
- Any necessary flow redirection is already configured on the deployed appliance and, if necessary, the appropriate router.



For detailed appliance configuration information, refer to the *Appliance Manager Operator's Guide*. Also see the *Network Deployment Guide* for specific scenarios.

One more thing ...

This is also a good time to **add users** to the Orchestrator server database. By default, the Orchestrator uses this local database for authentication. However, you can also point to a RADIUS or TACACS+ server for that function.

Related Menus

[Orchestrator Administration > User Management](#)

[Orchestrator Administration > Authentication](#)

For more information, see “Managing Orchestrator User Accounts and Authentication” on page 9.

1. **Server mode** is a subset of Router mode. It uses one interface for both management and datapath traffic.

What to Configure Next in a WAN Optimization Network

This is the general workflow to follow when working with **NX**, **VX**, and **VRX** appliances in a WAN Optimization network (as opposed to an SD-WAN network).

Initially, you'll configure the more generic items. For example:

- 1 In the navigation pane, use contextual menus to **create a group** or groups to which you'll assign each appliance. For example, you may choose to create a group for Engineering or Finance.
- 2 The Orchestrator **adds appliances by discovery**.

When you add an NX, VX, or VRX appliance to your network, you use Appliance Manager's **Monitoring > Orchestrator Reachability** page to add the Orchestrator's IP address.



Note To configure deployment for an NX, VX, or VRX appliance, you need to configure it in the appliance itself, using Appliance Manager.

As soon as the Orchestrator establishes communication, all of the appliance's existing configuration, alarm, and statistical data is available immediately.

- 3 **Create and apply configuration templates**. Create templates for non-unique variables and apply across one or more appliances. They include templates for SNMP, DNS, date and time, tunnel characteristics, SSL certificates, web-related parameters, user-defined applications, policies, logging, etc.

For more information, see Chapter 3, "Configuration Templates."

IMPORTANT: Templates will **REPLACE** all settings on the appliance with the template settings unless the template has a **MERGE** option and that option is selected.

However, in the case of templates for policies (Route, Optimization, QoS, NAT) and ACLs:

- You can create template rules with priority from **1000 – 9999**, inclusive. When you apply the template to an appliance, the Orchestrator deletes all appliance entries in that range before applying its policies.
- If you access an appliance directly (via the WebUI or the command line interface), you can create rules that have higher priority (**1 – 999**) than Orchestrator rules and rules that have lower priority (**10000 – 65534**).

Related Menus

[Configuration > Templates](#)

- 4 **Subnet sharing** is a method for automatically routing a flow into the appropriate tunnel for optimization based on destination IP alone. The appliance builds a subnet table from entries added automatically by the system or manually by a user. When two appliances are connected by a tunnel, they exchange this information ("learn" it) and use it to route traffic to each other.

Locally connected networks are automatically added to the subnet table. You will need to add any additional local subnets manually.

For more information, see "Adding to the Subnet Table" on page 11.

Related Menus

[Configuration > Subnets](#)

- 5 If **tunnels** don't already exist, then:
 - You can enable each appliance's **auto tunnel** feature. This feature automatically creates tunnels between Silver Peak appliances that have network connectivity and active flows.

Related Menus

[Configuration > Templates > System](#)

- If you prefer to retain more control and configure the tunnels yourself, you can disable the **auto tunnel** feature in the appliance's system configuration and create the configurations manually.

Related Menus

[Configuration > Templates > System](#)

[Configuration > Templates > Tunnels](#)

[Configuration > Tunnels](#)

[Configuration > Tunnel Groups](#)

- 6 Generate your first reports.

For more information, see “Configuring and Distributing Custom Reports” on page 168.

Related Menus

[Monitoring > Schedule & Run Reports](#)

Understanding Topology and Layout

Alarms

- The **Alarms Summary** shows the total number of Orchestrator and appliance alarms, and color-codes them.
- Click the summary bar to hyperlink to the **Alarms** page.

Topology Settings & Legend

- The **Legend** details the appliances' management and operational states.

The screenshot shows the 'Topology Settings & Legend' panel. It includes the following elements and callouts:

- Map - Upload**: A button to upload a topology map of your own, or select from the drop-down list.
- World-Large**: A drop-down menu for selecting a map.
- Link Display Limit**: A text input field set to 600, representing the maximum number of tunnels to draw on the map.
- Legend**: A section containing various status indicators for appliances and links.
- Appliances**: A sub-section with the following items:
 - Alarms**: A blue square with the number 2, indicating the number of most severe alarms on the appliance. A callout notes: "Number of most severe alarms on appliance. This one shows two **Warnings**."
 - Bypass**: A red square with a white 'B', referring to *hardware bypass*.
 - Orphaned Tunnel**: A blue square with a white 'T'.
 - Unsaved Changes**: A grey square with a white 'U'.
 - Unreachable**: A red circle.
 - Unknown**: A pink circle.
 - Unsupported**: A cyan circle.
 - Not synchronized**: An orange circle.
 - Synchronizing**: A purple circle.
- Links**: A sub-section with the following items:
 - Up**: A solid green line.
 - Partially Down**: A dashed orange line.
 - Down**: A solid red line.
 - Pending**: A solid cyan line.
 - One-way**: A dashed orange line.

- **Bypass** refers to *hardware bypass*. If there is a major problem with the appliance hardware, software, or power, all traffic goes through the appliance without any processing. Additionally, you can manually put the appliance into **Bypass** mode as an aid to troubleshooting or during maintenance events.
- If an appliance displays **Unsaved Changes**, you must log into the appliance directly to save the changes.
- An **Unreachable** appliance is one that the Orchestrator can't contact.
- The Orchestrator acts as configurations cache for the appliances. When the Orchestrator doesn't have a configuration cache from an appliance, it is **Not synchronized**.
- An appliance is **Unsupported** when the Orchestrator software version doesn't support the appliance's software version.

Other

- Tunnel states are color-coded, and rollover with the mouse displays the state. For example, **Up**.
- Tables are sortable by column.
- Clicking the **Edit** icon provides direct access to editing a specific appliance by opening the corresponding Appliance Manager page in a separate browser tab.

[Orchestrator]

SNMP ?

Show 25 ▾ Search

Edit	Mgmt IP	Appliance Name	Enable SNMP	Enable SNMP Traps	Enable V3 User	Trap Receivers	
						Trap Receiver 1	Trap Receiver
	10.0.236.198	Tallinn	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
	10.0.238.69	laine-vxb	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
	10.0.238.71	laine-vxa	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
	10.0.238.20	laine2-vxa	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		
	10.0.238.21	laine2-vxb	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>		

Click to open Appliance Manager for this **mgmt0** IP address

[Appliance Manager]

Silver Peak

Name **laine-vxa** IP **10.0.238.71**
 Up Time 17d 0h 23m 17s VXOA 6.2.7.0-33/89
 Time 2014/12/24 01:06:38 UTC User remote [log out]

Save Changes

Alarms 0 Critical 0 Major 0 Minor 0 Warning

Application View | Network View | Monitoring | Configuration | Administration | Maintenance

SNMP ?

Enable SNMP
 Enable SNMP Traps
 Read-Only Community *****
 Default Trap Community *****

SNMP V3

Enable Admin User

Authentication Private

Type SHA1 AES-128
 Password

Trap Receivers

Add

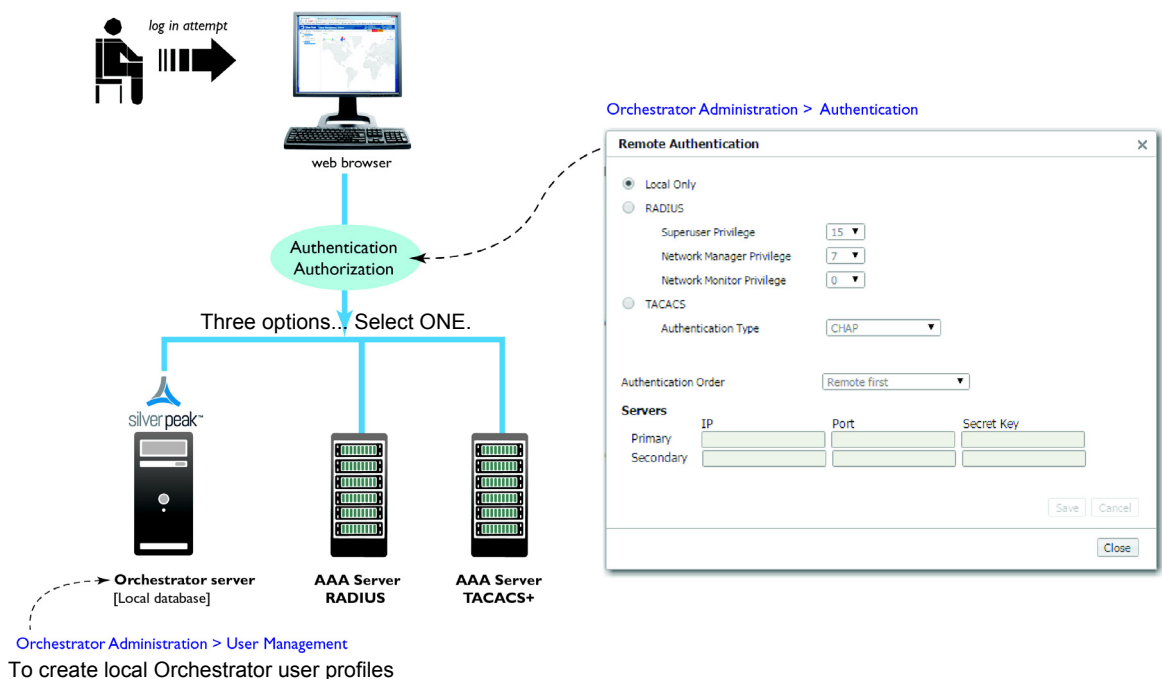
Enabled

Managing Orchestrator User Accounts and Authentication

For a user to successfully log into the Orchestrator client, the Orchestrator server must authenticate and authorize the user. Only then does the user have access to the Orchestrator server and, by extension, the appliances.

Based on its configuration, the Orchestrator authenticates the user via its own built-in local database or via a network server used for access control.

- The AAA server (Authentication Authorization Accounting server) can be either a **RADIUS** server or a **TACACS+** server.
- Add users to the Orchestrator server's local database via the Orchestrator client's **Orchestrator Administration > User Management** menu. The user profile includes the user role, which maps to a particular level of authorization and determines what the user can do.

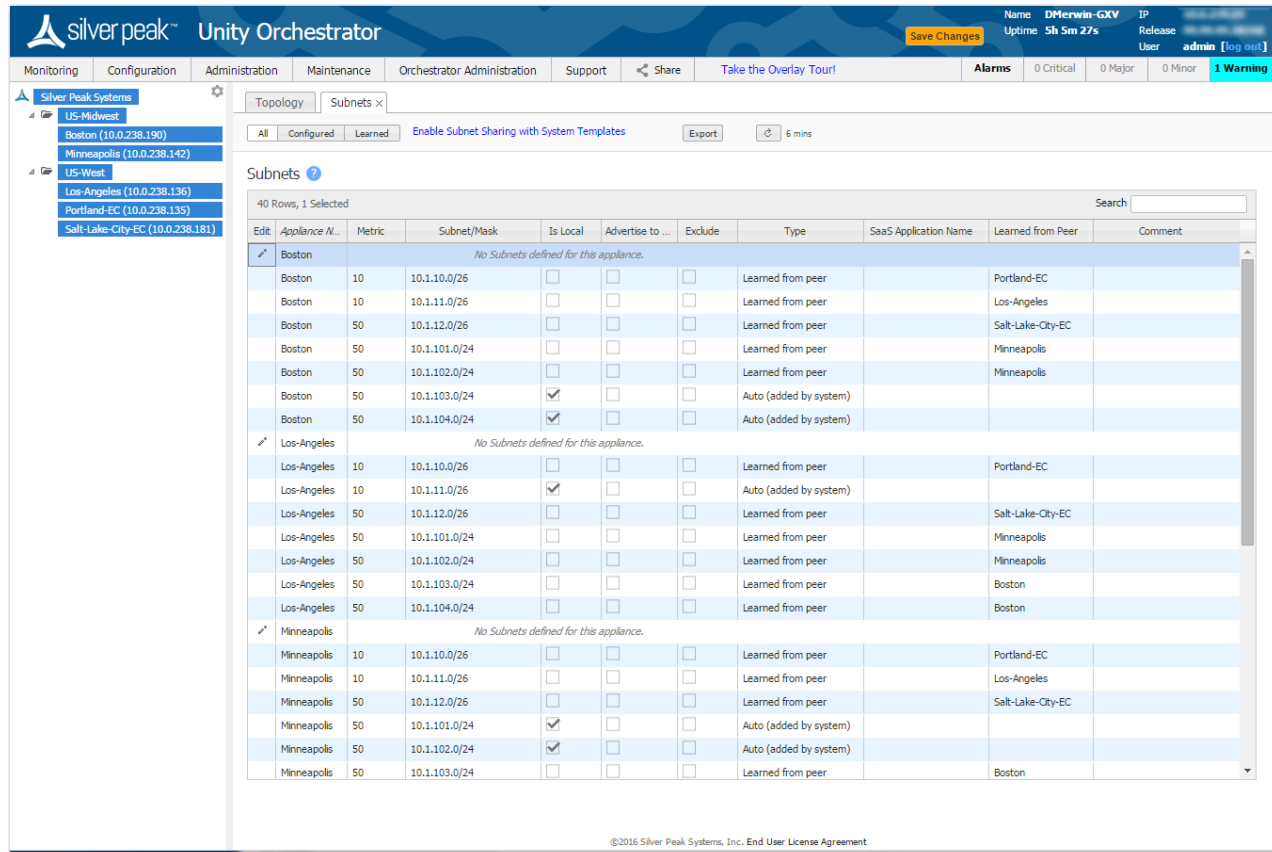


- The Orchestrator has three user roles: **Admin Manager (Superuser)**, **Network Manager**, and **Network Monitor**. Authorization always maps to one of these three levels:
 - **Admin Manager** has all privileges. It's the equivalent of **Superuser**.
 - **Network Manager** has read/write privileges. In practice, these are the same privileges that Admin Manager has.
 - **Network Monitor** has view-only privileges.
- Although there are three authentication options to choose from, you can only configure one.
 - If **Local Only** is selected, then authentication defaults to the Orchestrator server's local database.
 - If **Local Only is not** selected, then either a (remote) **RADIUS** or **TACACS+** server is also involved.
 - If **Remote first** is selected and fails, then the Orchestrator tries the **Local** database.
 - If **Local first** is selected and fails, then the Orchestrator tries the **Remote** database.
- The **Secret Key** enables the Orchestrator to talk to the access control server. The Orchestrator has hard-coded keys for TACACS+, so no user entry is required.

- You can also use Orchestrator templates to create remote authentication profiles for direct access to individual appliances via Appliance Manager or the CLI. Be aware, though, that that is different than creating a remote authentication profile for the Orchestrator.

Adding to the Subnet Table

To add, edit, or delete a subnet, you must select an individual subnet from the navigation panel and click in **Edit**. That opens a new browser tab on the specific appliance's **Subnets** page.



What is subnet sharing ?

Subnet sharing is one of the three strategies that Silver Peak uses to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. Auto-optimization strategies reduce the need to create explicit route map entries to optimize traffic. The other two strategies are **TCP-based** auto-opt and **IP-based** auto-opt.



Note Enabled by default, the global settings for all three reside on the **Templates** tab, under **System**.

How is subnet sharing implemented?

Each appliance builds a subnet table from entries added automatically by the system and manually by a user. When two appliances are connected by a tunnel, they exchange this information ("learn" it) and use it to route traffic to each other.

When would you need to use a Route Policy template?

Subnet sharing takes care of optimizing IP traffic.

Use and apply a Route Policy template for flows that are to be:

- sent pass-through (shaped or unshaped)
- dropped

- configured for a specific high-availability deployment
- routed based on application, ports, VLAN, DSCP, or ACL (Access Control List)

Subnet table columns

- **Subnet/Mask:** Actual subnet to be shared or learned
- **Metric:** Metric of the subnet. Value must be between 0 and 100. When a peer has more than one tunnel with a matching subnet (for example, in a high availability deployment), it chooses the tunnel with the greater numerical value.
- **Is Local:** Specifies if the subnet is local to this site.

The appliance sets this parameter for **automatically** for locally connected subnets of the appliance.

Also, you can select the parameter when manually adding a subnet:

- Select this option for a **manually** added subnet if all the IP addresses in the subnet are known to be local.
- Deselect this option if the subnet is so large (for example, 0.0.0.0/0) that it may include IP addresses that are not local to this appliance. If a subnet is too wide, and it's marked **local**, then the stats will count any pass-through packets with an IP address within that range as WAN-to-LAN.
- **Advertise to Peers:** Selected by default, it shares the subnet information with peers. Peers then learn it. To add a subnet to the table without divulging it to peers, yet, deselect this option.
- **Exclude:** Use this option to prevent optimization of more specific subnets from a wider advertised subnet range.
- **Type** of subnet:
 - **Auto (added by system)** = automatically added subnets of interfaces on this appliance
 - **Added by user** = manually added/configured subnets for this appliance
 - **Learned from peer** = subnets added as a result of exchanging information with peer appliances
- **SaaS Application Name:** If the subnet is associated with a SaaS service, the name displays here.
- **Learned from Peer:** Which peer appliance advertised (and shared) this subnet information



Unity Overlays

This chapter describes the screens related to creating SD-WAN overlays.

In This Chapter

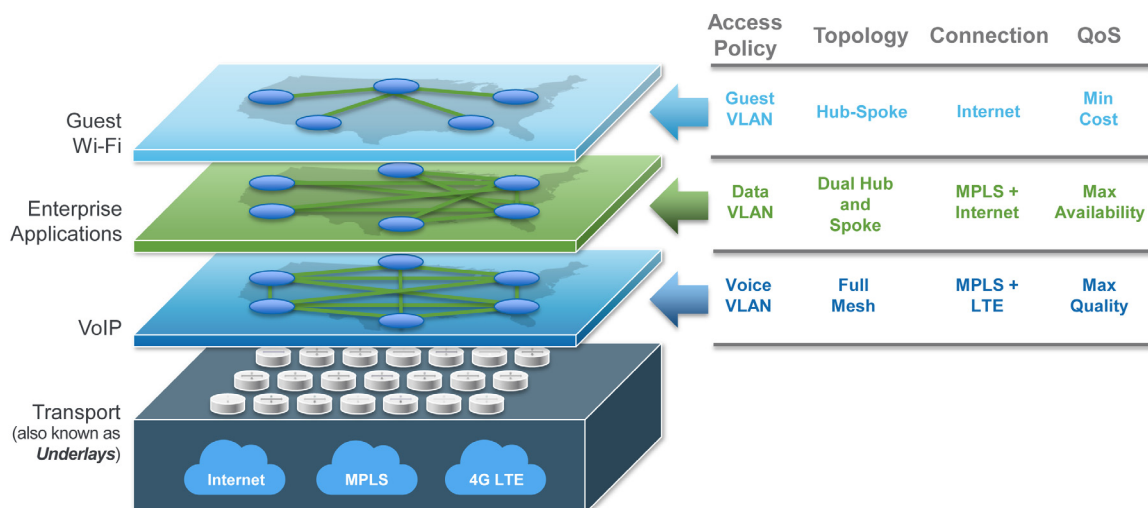
- **Introduction to Unity Overlays for SD-WAN** See page 14.
- **Discovered Appliances** See page 16.
- **Deployment Profiles** See page 17.
- **Business Intent Overlays** See page 27.
- **Apply Overlays** See page 29.
- **Interface Labels** See page 30.
- **Licenses** See page 31.
- **Silver Peak Cloud Portal** See page 32.
- **DHCP Server** See page 33.
- **IPSec Pre-shared Key Rotation** See page 35.
- **Configuration Wizard** See page 36.

Introduction to Unity Overlays for SD-WAN

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

- **Interface Labels** associate each interface with a use.
 - **LAN** labels refer to traffic type, such as **VoIP**, **data**, or **replication**.
 - **WAN** labels refer to the service or connection type, such as **MPLS**, **internet**, or **Verizon**.
- **Deployment Profiles** configure the interfaces and map the labels to them, to characterize the appliance.
- **Business Intent Overlays** use the Labels specified in Deployment Profiles to define how traffic is routed and optimized between sites. These overlays can specify preferred paths and can link bonding policies based on *application*, *VLAN*, or *subnet*, independent of the brand and physical routing attributes of the underlay.

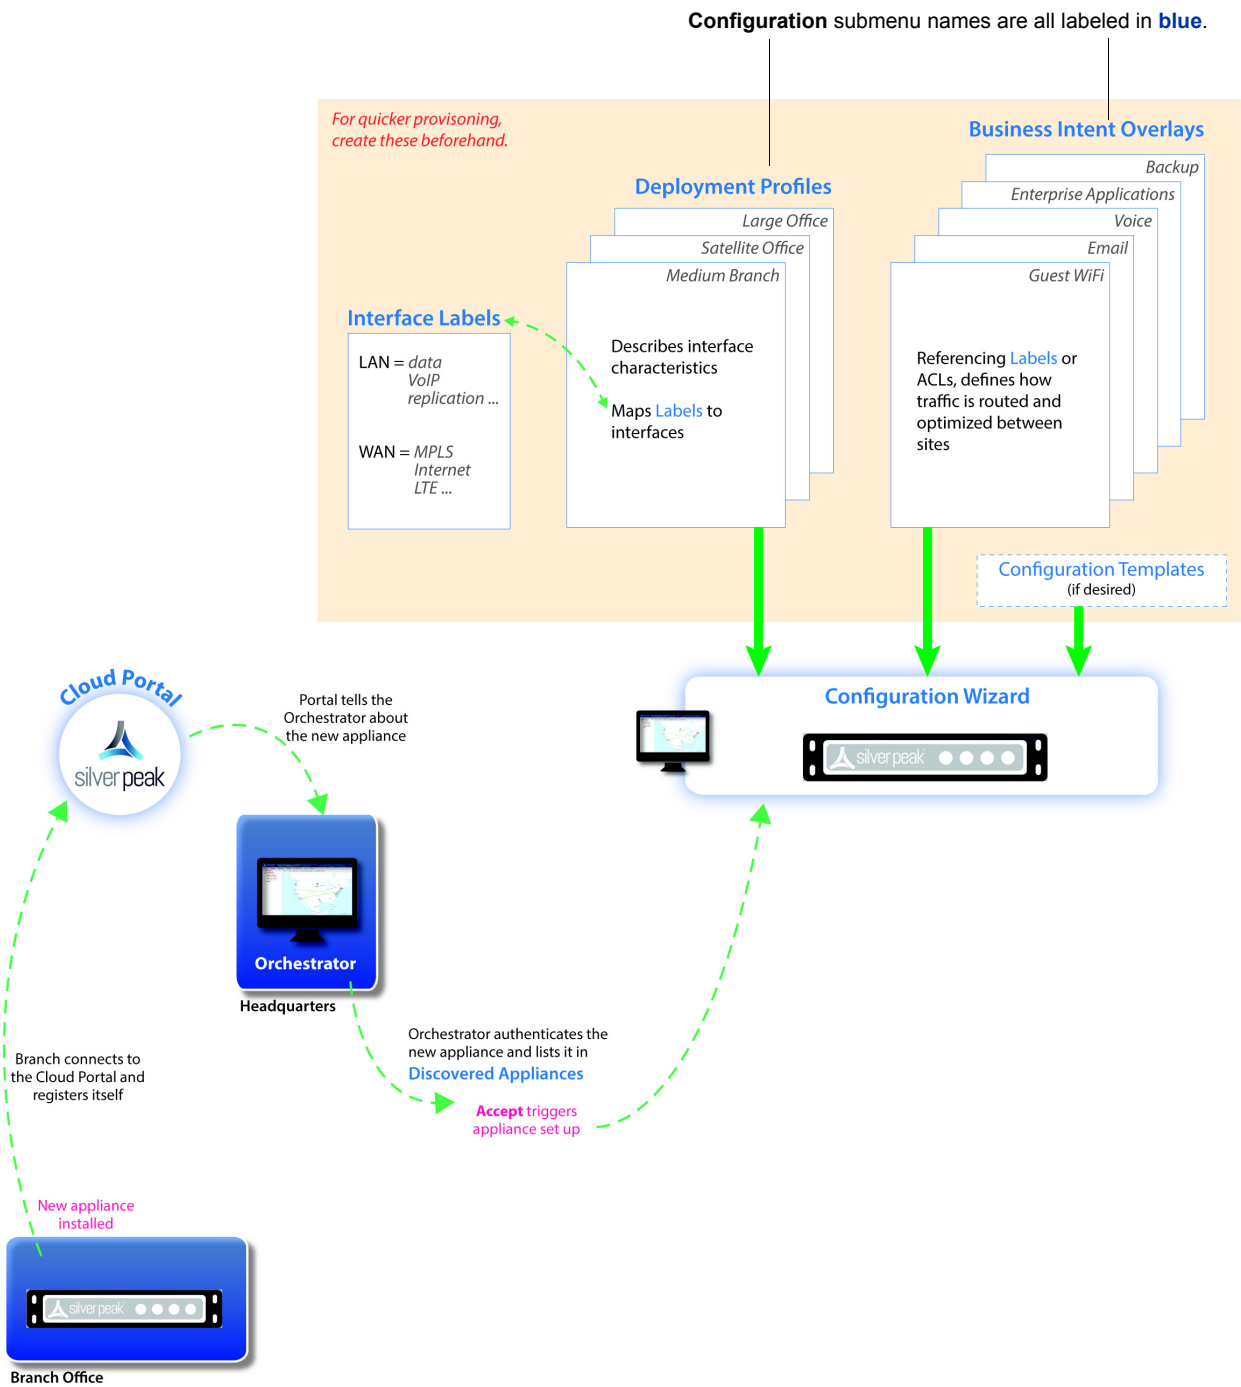
This diagram shows the basic architecture and capabilities of **Overlays**.



Including a new appliance into the Unity fabric consists of two basic steps:

- 1 **Registration and discovery.** After you **Accept** the discovered appliance, it opens the **Configuration Wizard**.
- 2 **Provisioning.** Since the wizard prompts you to select profiles, it's easiest to create these ahead of time.

Figure 2-1 The process of installing and provisioning an appliance for SD-WAN.



Discovered Appliances

Configuration > [Unity Overlays] Discovered Appliances

This page lists each appliance that the Orchestrator discovers.

The screenshot shows the 'Discovered Appliances' page in the Unity Orchestrator. At the top, there are tabs for 'Topology' and 'Discovered Appliances x'. Below the tabs, there is a search bar and a 'Discovery Email Recipients' field with the value 'testAccount@silver-peak.com'. The main area contains a table with 40 rows and 1 selected row. The table has the following columns: Serial Number, Hostname, IP Address, Public IP Address, Location, Tag, Discovered Time, Approve, Deny, and Model. Each row represents a discovered appliance, and the 'Approve' and 'Deny' columns contain buttons for each appliance.

Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Approve	Deny	Model
000C294D1E7C	JENKINS2-VXB	10.0.238.51	10.0.238.5		Unassigned	03-Feb-16 19:08	Approve	Deny	VX-1000
000C29B3044E	JENKINS2-VXA	10.0.238.50	10.0.238.5		Unassigned	03-Feb-16 19:08	Approve	Deny	VX-1000
000C29863CB6	JENKINS-VX1000B	10.0.238.45			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C295E943D	SEL-VX1000A	10.0.233.107			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29ADA2A4	SEL7-VX1000B	10.0.238.86			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29ECB055	SEL7-VX1000A	10.0.238.86			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
001BBC035678	SEL-VX1000A	10.0.233.107			Unassigned	19-Jan-16 21:15	Approve	Deny	EC-V
000C291C7F7A	AUTO2-VX1000B	10.0.239.30			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29A8280D	JENKINS-VX1000A	10.0.238.44			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29FEC411	AUTO2-VX1000A	10.0.239.28			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29AC8619	AWA-VX1000A	10.0.237.208			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29C275D7	AWA-VX1000E	10.0.237.242			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29ADA2A4	SEL-VX1000B	10.0.233.108			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C295E943D	SEL-VX1000A	10.0.233.107			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C29AA1D79	system1-vxb	10.0.237.69			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C297F690F	SEL3-VX1000A	10.0.238.161			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000
000C2946B44F	AWA-VX2000B	10.0.237.209			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-2000
000C29A47F2C	SELS-VX2000B	10.0.233.95			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-2000
000C292DB6C7	SELS-VX2000A	10.0.233.126			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-2000
000C29788A7	AUTO1-VX1000B	10.0.233.5			Unassigned	19-Jan-16 21:15	Approve	Deny	VX-1000

- To enable the Orchestrator to manage an appliance after you verify its credentials, click **Approve**.
- If the appliance doesn't belong in your network, click **Deny**. If you want to include it later, click **Show Denied Appliances**, locate it in the table, and click **Approve**.

Deployment Profiles

Configuration > [Unity Overlays] Deployment Profiles

Instead of configuring each appliance separately, you can create various **Deployment Profiles** and provision a device by applying the profile you want. For example, you can create a standard format for your branch.



Tip For smoother workflow, complete the **Configuration > DHCP Server** tab before creating Deployment Profiles.

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

The screenshot displays the 'Deployment Profiles' configuration page for a Silver Peak device. The profile name is 'MediumBranch'. The device type is 'Router'. The configuration is divided into LAN and WAN interface sections. The LAN interface 'lan0' has a 'None' label and 'No DHCP' configuration. The WAN interface 'wan0' has an 'MPLS' label and 'Internet' service. Shaping Kbps are set to 100,000 for both directions. Total Outbound and Inbound traffic are both 150,000 Kbps. There are also options for EdgeConnect Licensing and Boost.



Note IP/Mask fields are not editable because they are appliance-specific.

Information for this tab is organized as follows:

- **Mapping Labels to Interfaces** See page 18.
- **LAN-side Configuration: DHCP** See page 18.
- **WAN-side Configuration** See page 18.
- **Definitions** See page 19.
- **A More Comprehensive Guide to Basic Deployments** See page 19.
- **How You Can Adjust the Basic Deployments** See page 25.
- **Adding Data Interfaces** See page 26.

Mapping Labels to Interfaces

- On the **LAN** side, labels identify the data, such as *data*, *VoIP*, or *replication*.
- On the **WAN** side, labels identify the service, such as *MPLS* or *Internet*.
- To create a global pool of labels, either:
 - Click the **Edit** icon next to **Label**.
 - Select **Configuration > Interface Labels**.
- If you edit a label, that change propagates appropriately. For example, it renames tunnels that use that labeled interface.

LAN-side Configuration: DHCP

- By default, *each* LAN IP acts as a **DHCP Server** when the appliance is in (the default) Router mode.
- The global defaults are set in **Configuration > DHCP Server** and pre-populate this page. The other choices are **No DHCP** and having the appliance act as a **DHCP Relay**.
- To customize an individual interface in the Deployment Profile, click the Edit icon under the **IP/Mask** field, to the right of the displayed DHCP label.

WAN-side Configuration

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

- For traffic inbound from the WAN, the appliance accepts *only* IPsec tunnel packets.
- For traffic outbound to the WAN, the appliance *only* allows IPsec tunnel packets and management traffic.
- Click the *lock icon* to toggle between hardening and unhardening an interface.

NAT: If the appliance is behind a NAT-ed interface, select **NAT** (without the strikethrough). 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.

Shaping: You can limit bandwidth selectively on each WAN interface.

- **Total Outbound** bandwidth is licensed by model. It's the same as max system bandwidth.
- To enter values for shaping inbound traffic, which is optional, you must first select **Shape Inbound Traffic**.

EdgeConnect Licensing: Only visible on EC appliances

- By default, every EC has a max system bandwidth of 200 Mbps. For more bandwidth, you can purchase **Plus**, and then select it here for this profile.
- If you've purchased a reserve of **Boost** for your network, you can allocate a portion of it in a Deployment Profile. You can also direct allocations to specific types of traffic in the Business Intent Overlays.
- To view how you've distributed **Plus** and **Boost**, view the **Configuration > Licenses** tab.

Definitions

Following are the definitions for DHCP servers and DHCP relays.

DHCP Server Definitions

- **DHCP Pool Subnet/Mask** is the full range of IP addresses that you make available for your network.
- **Subnet Mask** is a mask that specifies the default number of IP addresses reserved for any subnet. For example, entering 24 reserves 256 IP addresses.
- **Start Offset** specifies how many addresses not to allocate at the beginning of the subnet's range. For example, entering 10 means that the first ten IP addresses in the subnet aren't available.
- **End Offset** specifies how many IP addresses are not available at the end of the subnet's range.
- **Default lease** and **Maximum lease** specify, in hours, how long an interface can keep a DHCP-assigned IP address.
- **Default gateway**, when selected, indicates that
- **DNS server(s)** specifies the associated Domain Name System server(s).
- **NTP server(s)** specifies the associated Network Time Protocol server(s).
- **NetBIOS name server(s)** is used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
- The **NetBIOS node type** of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types:
 - **B-node** = 0x01 Broadcast
 - **P-node** = 0x02 Peer (WINS only)
 - **M-node** = 0x04 Mixed (broadcast, then WINS)
 - **H-node** = 0x08 Hybrid (WINS, then broadcast)

DHCP Relay Definitions

- **Destination DHCP Server** is the IP address of the DHCP server assigning the IP addresses.
- **Enable Option 82**, when selected, inserts additional information into the packet header to identify the client's point of attachment.
- **Option 82 Policy** tells the relay what to do with the hex string it receives. The choices are **append**, **replace**, **forward**, or **discard**.

A More Comprehensive Guide to Basic Deployments

This section discusses the basics of three deployment modes: **Bridge**, **Router**, and **Server** modes.

It describes common scenarios, considerations when selecting a deployment, redirection concerns, and some adaptations.

For detailed deployment examples, refer to the *Silver Peak Network Deployment Guide*.


In Bridge Mode and in Router Mode, you can provide security on any WAN-side interface by **hardening the interface**. This means:

- For traffic inbound from the WAN, the appliance accepts **only** IPsec tunnel packets.
- For traffic outbound to the WAN, the appliance **only** allows IPsec tunnel packets and management traffic.
- Click the *lock icon* to toggle between hardening and unhardening an interface.

Bridge Mode

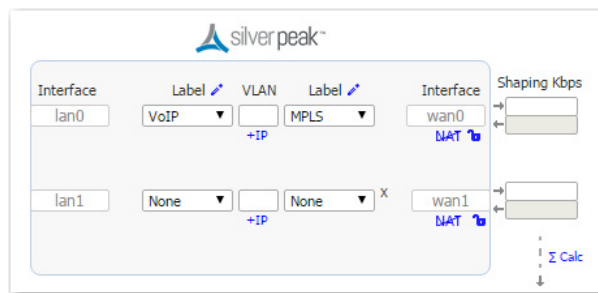
■ Single WAN-side Router

In this deployment, the appliance is in-line between a single WAN router and a single LAN-side switch.



■ Dual WAN-side Routers

This is the most common 4-port bridge configuration.



- 2 WAN egress routers / 1 or 2 subnets / 1 appliance
- 2 separate service providers or WAN services (MPLS, IPsec VPN, MetroEthernet, etc.)

■ Considerations for Bridge Mode Deployments

- Do you have a physical appliance or a virtual appliance?
- A virtual appliance has no fail-to-wire, so you would 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.

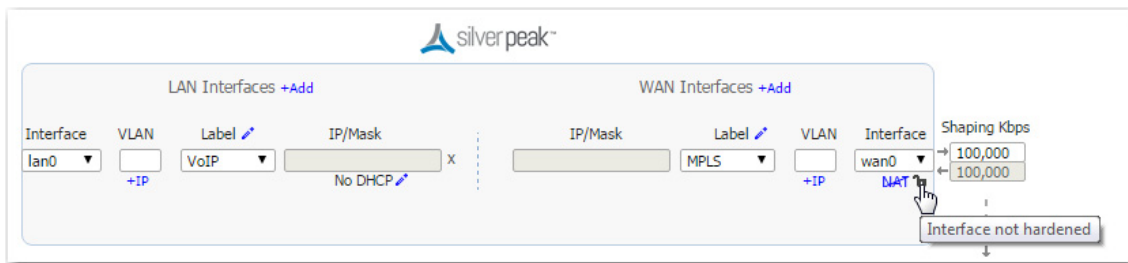
Router Mode

There are four options to consider:

- 1 Single LAN interface & single WAN interface
- 2 Dual LAN interfaces & dual WAN interfaces
- 3 Single WAN interface sharing LAN and WAN traffic
- 4 Dual WAN interfaces sharing LAN and WAN traffic

For best performance, visibility, and control, Silver Peak recommends Options #1 and #2, which use separate LAN and WAN interfaces. And when using NAT, use Options #1 or #2 to ensure that addressing works properly.

■ #1 - Single LAN Interface & Single WAN Interface

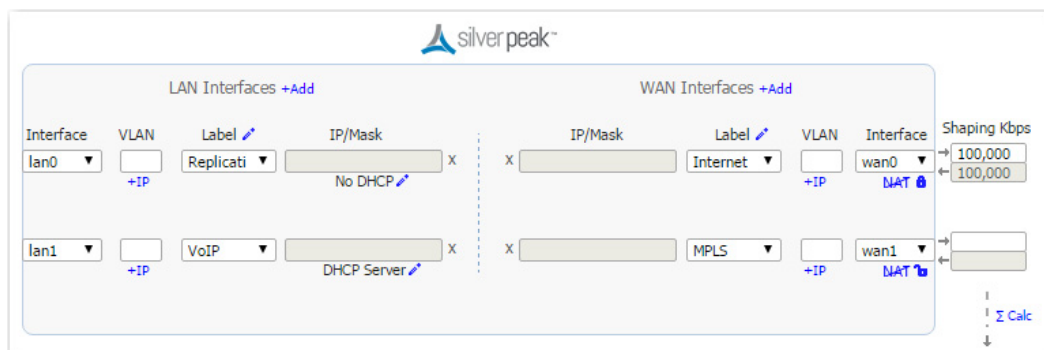


For this deployment, you have two options:

- a You can put Silver Peak *in-path*. In this case, if there is a failure, you need other redundant paths for high availability.
- b You can put Silver Peak *out-of-path*. You can redirect LAN-side traffic and WAN-side traffic from a router or L3 switch to the corresponding Silverpeak interface, using WCCP or PBR (Policy-Based Routing).

To use this deployment with a single router that has only one interface, you could use multiple VLANs.

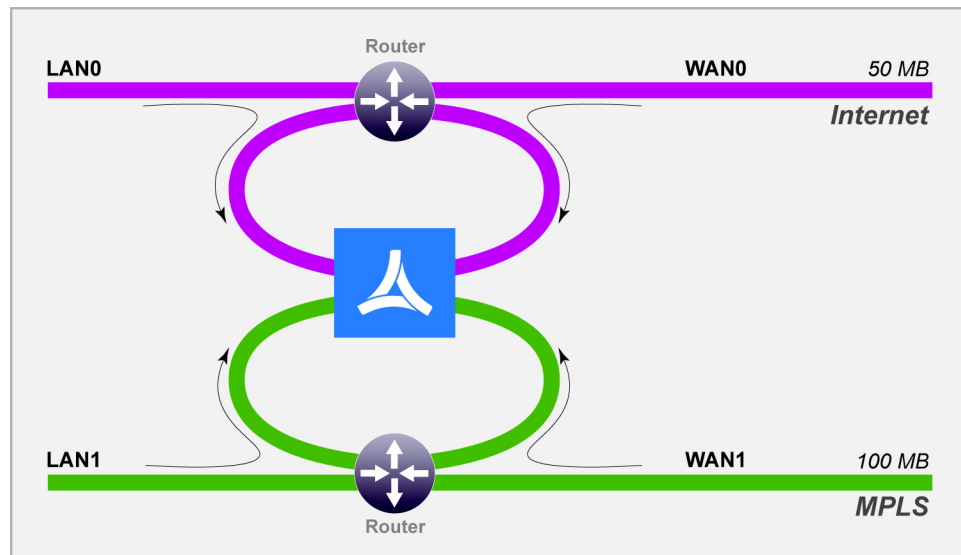
■ #2 - Dual LAN Interfaces & Dual WAN Interfaces



This deployment redirects traffic from two LAN interfaces to two WAN interfaces on a single Silver Peak appliance.

- 2 WAN next-hops / 2 subnets / 1 appliance
- 2 separate service providers or WAN services (MPLS, IPsec VPN, MetroEthernet, etc.)

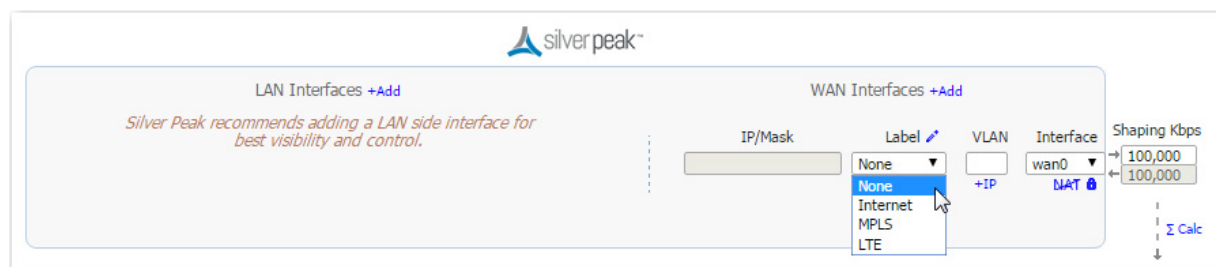
Out-of-path dual LAN and dual WAN interfaces



For this deployment, you have two options:

- You can put Silver Peak *in-path*. In this case, if there is a failure, you need other redundant paths for high availability.
- You can put Silver Peak *out-of-path*. You can redirect LAN-side traffic and WAN-side traffic from a router or L3 switch to the corresponding Silverpeak interface, using WCCP or PBR (Policy-Based Routing).

■ #3 - Single WAN Interface Sharing LAN and WAN traffic



This deployment redirects traffic from a single router (or L3 switch) to a single subnet on the Silver Peak appliance.

- This mode only supports *out-of-path*.

- When using two Silver Peaks at the same site, this is also the most common deployment for high availability (redundancy) and load balancing.
- For better performance, control, and visibility, Silver Peak recommends Router mode **Option #1** instead of this option.

■ #4 - Dual WAN Interfaces Sharing LAN and WAN traffic

The screenshot shows the Silver Peak configuration interface for WAN Interfaces. The interface is titled "WAN Interfaces +Add" and contains a table with the following columns: IP/Mask, Label, VLAN, Interface, and Shaping Kbps. There are two rows of configuration:

IP/Mask	Label	VLAN	Interface	Shaping Kbps
x []	MPLS_1	[] +IP	wan0 NAT	100,000 / 100,000
x []	MPLS_2	[] +IP	lan0 NAT	[] / []

Below the table, there are controls for "Calc" and a downward arrow.

This deployment redirects traffic from two routers to two interfaces on a single Silver Peak appliance.

This is also known as **Dual-Homed Router Mode**.

- 2 WAN next-hops / 2 subnets / 1 appliance
- 2 separate service providers or WAN services (MPLS, IPsec VPN, MetroEthernet, etc.)
- This mode only supports *out-of-path*.
- For better performance, control, and visibility, Silver Peak recommends Router mode **Option #2** instead of this option.

■ Considerations for Router Mode Deployments

- Do you want your traffic to be in-path or out-of-path? This mode supports both deployments. In-path deployment offers much simpler configuration.
- Does your router support VRRP, WCCP, or PBR? If so, you may want to consider out-of-path Router mode deployment. You can set up more complex configurations, which offer load balancing and high availability.
- Are you planning to use host routes on the server/end station?
- In the rare case when you need to send inbound WAN traffic to a router other than the WAN next-hop router, use LAN-side routes.

■ Examining the Need for Traffic Redirection

Whenever 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):

- **PBR** (Policy-Based Routing) — configured on the router. No other special configuration required on the appliance. This is also known as FBR (Filter-Based Forwarding).

If you want to deploy two Silver Peaks at the site, for redundancy or load balancing, then you also need to use VRRP (Virtual Router Redundancy Protocol).

- **WCCP** (Web Cache Communication Protocol) — configured on both the router and the Silver Peak appliance. You can 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 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).

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.

How you plan to optimize traffic also affects whether or not you also need inbound redirection from the WAN router (known as WAN-side redirection):


- If you use **subnet sharing** (which relies on advertising local subnets between Silver Peak appliances) or route policies (which specify destination IP addresses), then you only need LAN-side redirection.
- If, instead, you rely on **TCP-based** or **IP-based** auto-optimization (which relies on initial handshaking outside a tunnel), then you must also set up inbound and outbound redirection on the WAN router.
- For TCP flows to be optimized, both directions must travel through the same client and server appliances. If the TCP flows are asymmetric, you need to configure flow redirection among local appliances.

A tunnel must exist before auto-optimization can proceed. There are three options for tunnel creation:

- If you enable **auto-tunnel**, then the initial **TCP-based** or **IP-based** handshaking creates the tunnel. That means that the appropriate LAN-side and WAN-side redirection must be in place.
- You can let the **Initial Configuration Wizard** create the tunnel to the remote appliance.
- You can create a tunnel manually on the **Configuration - Tunnels** page.

Server Mode

This mode uses the **mgmt0** interface for management and datapath traffic.



How You Can Adjust the Basic Deployments

When you choose a deployment, only the appropriate options are accessible.



Bonding

- When using an NX or EC appliance with four 1Gbps Ethernet ports, you can bond like pairs into a single 2Gbps port with one IP address. For example, **wan0** plus **wan1** bond to form **bwan0**. This increases throughput on a very high-end appliance and/or provides interface-level redundancy.
- For bonding on a virtual appliance, you would need configure the host instead of the appliance. For example, on a VMware ESXi host, you would configure NIC teaming to get the equivalent of etherchannel bonding.
- Whether you use a physical or a virtual appliance, etherchannel must also be configured on the directly connected switch/router. See *“Configuring Gigabit Etherchannel Bonding”*, below.

Use Fiber Ports

Choose this when you want to enable 10Gbps ports on a physical appliance.

Propagate Link Down

Forces the WAN interface to go down when the corresponding LAN interface goes down, or vice versa.

4-port single bridge

This is a corner case. Here, four ports form a single bridge with a single WAN next-hop. This is in contrast to having dual WAN routers with two separate bridges.

Configuring Gigabit Etherchannel Bonding

When using a four-port Silver Peak appliance, you can bond pairs of Ethernet ports into a single port with one IP address. This feature provides the capability to carry 2 Gbps in and out of an appliance when both ports are in service.

When you configure bonding, the following is true:

- **lan0** plus **lan1** bond to form **blan0**, which uses the **lan0** IP address.
- **wan0** plus **wan1** bond to form **bwan0**, which uses the **wan0** IP address.
- The appliances use flow-based load balancing across the links.
- This configuration provides failover in case one link goes down.
- You can view the statistics on the **Monitoring - Interfaces** page. If you’re using bonding, you’ll see statistics for **blan0** and **bwan0**, as well as for the interfaces that comprise them (**lan0**, **lan1**, **wan0**, and **wan1**).
- If a WCCP or VRRP deployment already exists, then you must reconfigure the deployment on the bonding interface. In other words, if you previously configured on **wan0**, then after bonding you must reconfigure on **bwan0**.
- Rollback to non-bonding mode returns the intact, non-bonded configuration.
- Enabling/disabling bonding requires an appliance reboot.

◆ To configure etherchannel bonding

To enable bonding, you need to configure both the appliance and the router for bonding.

- 1 Access the **Configuration - Deployment** page. The three available bonding modes are:
 - a Out-of-path (Router/Server mode) with a single WAN-side router
 - b Out-of-path (Router/Server mode) with dual WAN-side routers
 - c In-path (Bridge mode) with dual WAN-side routers
- 2 Complete the various fields and click **Apply**.
- 3 When prompted, reboot the appliance.
- 4 Now, configure the Cisco router. Following is an example of the commands, where angle brackets indicate variables:

```
config t
interface range <g1/0/6-7>
channel-group <1> mode on

show etherchannel
show interface port-channel <1>
```

Adding Data Interfaces

- You can create additional data-plane Layer 3 interfaces, to use as tunnel endpoints.
- To add a new logical interface, click **+IP**.

Business Intent Overlays

Configuration > [Unity Overlays] Business Intent Overlays

Use **Business Intent Overlays** to create separate, logical networks that are individually tailored to your applications and requirements.

Essentially, a Business Intent Overlay describes where and how to build tunnels.

The screenshot shows the configuration page for Business Intent Overlays. On the left, there is a list of overlays: Voice, Backup, Guest_Wifi, and Email. The 'Voice' overlay is selected. The main configuration area is divided into several sections:

- Topology:** Mesh is selected. Hubs include Los-Angeles and Denver-EC.
- Overlay Policy:**
 - Traffic Access Policy:** VoIP is selected.
 - Link Brownout Thresholds:** Loss is 0.5%, Latency is 40ms, and Jitter is 50ms.
 - Route Matched Traffic to these WAN Ports:** Internet, MPLS, and LTE are selected.
- Link Bonding Policy:** High Quality is selected.
- Overlay Down Action:** Drop is selected.
- Shaping Traffic Class:** 1 (default) is selected.
- Boost License:** Boost this Traffic is checked.

- Each logical network is independent from the others in terms of topology, traffic type, security management and QoS queues, SaaS optimization, and WAN optimization. For example, you could have different overlays for voice, backup, Guest WiFi, email, and Salesforce.
- Overlays are independent of the brand and physical routing attributes of the underlay (physical network of switches and routers).
- Questions to ask when configuring overlays
 - What are the access policies to this logical network?
 - Which application, at the branch, is mapped into each one of these slices?
 - What service can a virtual network use?
 - If you've purchased **Boost**, then for which apps do you want to allocate some of the total WAN optimization bandwidth you've provisioned?

Topology

- You can choose either a **Mesh** or a **Hub & Spoke** topology.
- If choosing **Hub & Spoke**, choose the hubs you need from the **Select Hubs** area. If one you need isn't displayed, click **+Add**, as needed.
- Orchestrator builds the topology when you apply a Business Intent Overlay to appliances that have already been assigned a Deployment Profile.

Overlay Policy

Traffic Access Policy - Select the traffic you want to manage by choosing a labeled LAN port or an Access Control List (ACL).

Link Brownout Thresholds specify the triggers for switching from a **Primary** service to a **Backup** service. Exceeding any one of the three thresholds is sufficient.

Route Matched Traffic to these WAN Ports

- Traffic is routed to the Primary service unless a threshold for Loss, Latency, or Jitter has been exceeded.
- When **Blackout** is selected, then the Backup service is used only if the Primary service goes down completely.
- When **Brownout** is selected, then the Backup service is used until each aspect of the Primary service is again within normal limits.
- Choosing to **Cross Connect Providers** enables load balancing among Primary services by creating bonded tunnels. Also, failover of one doesn't force failover for the other.

Link Bonding Policy

When there are multiple tunnels between two appliances, you need to specify the criteria for selecting the best route. This is managed by *packet-based* Dynamic Path Control (DPC).

- **High Availability** – for critical real-time services that cannot accept any interruption at all. For example, call center voice or critical VDI traffic.
- **High Quality** – for typical real-time services, such as VoIP or video conferencing. For example, WebEx or business-quality Skype, VDI traffic.
- **High Throughput** – for anything where maximum speed is more important than quality. For example, data replication, NFS, file transfers, etc.
- **High Efficiency** – for everything else, including most TCP applications. This option sends load balance info on multiple links, with no FEC, no overhead, and just raw packets.

Path conditioning consists of proprietary algorithms to mitigate performance degradation from loss, jitter, or latency.

Silver Peak uses additional bandwidth to send parity packets, enabling data reconstruction at the remote site. As the number of parity packets is decreased, *BW Efficiency* increases.

Overlay Down Action - The options are **pass-through**, **pass-through unshaped**, or **drop**.

Shaping Traffic Class - Select from the 10 traffic classes.

Boost License - Select if you've purchased **Boost** and want to apply it to this overlay.

Best Practices

Create ACLs beforehand, using the configuration Template, and push them to the appliances. You'll then be able to select from them in the **Traffic Access Policy**, to match traffic.

Apply Overlays

Configuration > [Unity Overlays] Apply Overlays

Use this page to **add or remove overlays** from appliances.

Displays appliances selected in the tree view

The screenshot displays the 'Apply Overlays' configuration page. The page title is 'Apply Overlays' with a search icon and a link to 'Edit Overlays'. Below the title, there is a 'View Status' section with a green 'Add' button and an orange 'Remove' button. The main content is a table with 5 rows, showing a list of appliances. The table has columns for 'Appliance' (Hostname and IP Address) and 'Overlays' (Present and Changes). The 'Present' column lists the overlays applied to each appliance: Voice, Backup, Guest_Wifi, and Email. On the left side of the table, there are checkboxes for each overlay type: Voice, Backup, Guest_Wifi, and Email. Below the checkboxes are 'Apply' and 'Cancel' buttons.

Appliance		Overlays	
Hostname	IP Address	Present	Changes
Chicago	10.0.238.190	Voice, Backup, Guest_Wifi, Email	
Dallas	10.0.238.189	Voice, Backup, Guest_Wifi, Email	
Denver-EC	10.0.238.181	Voice, Backup, Guest_Wifi, Email	
Los-Angeles	10.0.238.136	Voice, Backup, Guest_Wifi, Email	
Seattle-EC	10.0.238.135	Voice, Backup, Guest_Wifi, Email	

Interface Labels

Configuration > [Unity Overlays] Interface Labels

Use this dialog box to create labels for the WAN and LAN interfaces.

Type	Label	
lan	Data	X
lan	Voice	X
wan	LTE	X
wan	Internet	X
wan	MPLS	X

Licenses

Configuration > [Unity Overlays] Licenses

This page lists each appliance's make, model, license terms, and registered services.

Topology
Licenses ×

Licenses ?
🔄 2 mins

Appliances	5
NX	0
VX	5

EC Base	0/10
EC Plus	0/10
EC Boost	0.0 Kbps/1.0 Gbps

SaaS	Valid Until 03-Sep-17 0:00
EC	Valid Until 03-Sep-17 0:00
Orchestrator	License Not Required.

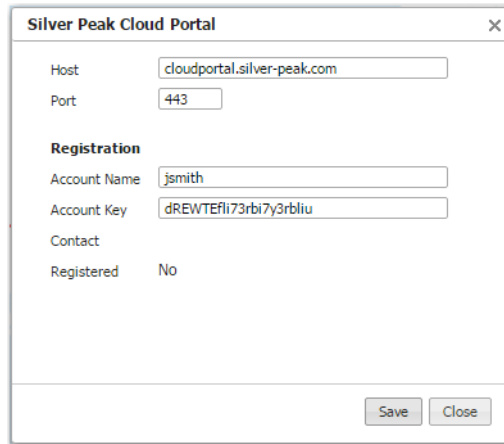
[Configure EC Licenses](#)

5 Rows	Search	Edit	Host Name	Model	Serial No	License Start	License End	SaaS
		✎	Los-Angeles	VX3000	00-1B-BC-03-00-00	30-Jul-14 17:00:00	Perpetual	Yes
		✎	Seattle-EC	VX3000	00-0C-29-47-5A-FE	28-Jan-13 16:00:00	Perpetual	Yes
		✎	Denver-EC	VX5000	00-1B-BC-03-00-01	27-Jul-15 17:00:00	26-Oct-17 16:59:59	Yes
		✎	Dallas	VX1000	00-0C-29-0A-25-73	04-Nov-15 16:00:00	Perpetual	No
		✎	Chicago	VX1000	00-0C-29-9C-71-3E	04-Nov-15 16:00:00	Perpetual	No

Silver Peak Cloud Portal

Configuration > [Unity Overlays] Silver Peak Cloud Portal
Orchestrator Administration > [General] Silver Peak Cloud Portal

The **Silver Peak Cloud Portal** is used to register cloud-based features and services, such as *SaaS optimization*, *EdgeConnect*, and *CPX*.



The screenshot shows a dialog box titled "Silver Peak Cloud Portal" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Host:** A text input field containing "cloudportal.silver-peak.com".
- Port:** A text input field containing "443".
- Registration:** A section header.
- Account Name:** A text input field containing "jsmith".
- Account Key:** A text input field containing "dREWTEfll73rbi7y3rbliu".
- Contact:** A text input field that is currently empty.
- Registered:** A checkbox labeled "Registered" with the value "No".
- Buttons:** "Save" and "Close" buttons are located at the bottom right of the dialog.

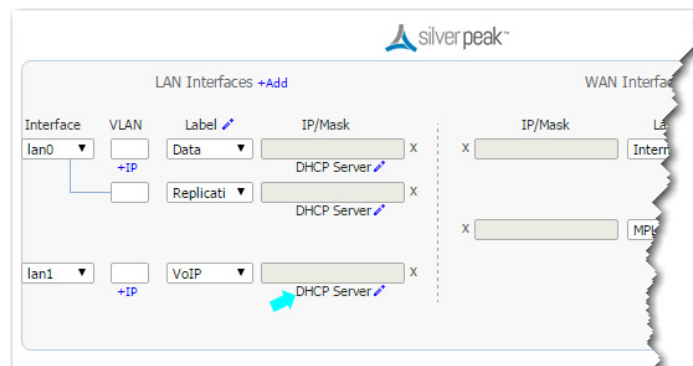
- When you purchase one of these services, Silver Peak sends you an **Account Name** and instructions to obtain your **Account Key**. You will use these to register your appliance(s).
- The cloud portal populates the **Contact** field from information included in your purchase order.
- Use of these services requires that your appliance(s) can access the cloud portal via the Internet.

DHCP Server

Configuration > [Unity Overlays] DHCP Server

You can reduce your workload by using this page to configure global defaults for Dynamic Host Configuration Protocol (DHCP).

- These defaults apply to the LAN interfaces in **Deployment Profiles** that specify Router mode.
- There are three choices:
 - **No DHCP**
 - Each LAN interface acts as a **DHCP Server**.
 - The Silver Peak appliance acts as a **DHCP Relay** between a DHCP server at a data center and clients needing an IP address.
- On the **Configuration > Deployment Profiles** tab, the selected default displays consistently under each LAN-side **IP/Mask** field.



For any LAN-side interface, you can override the global default by clicking the Edit icon to the right of the label and changing the values or selection.

- Changes you save to the global default only apply to new configurations.
- To view or revise the list of reserved subnets, click **Monitor**.

DHCP Server Definitions

- **DHCP Pool Subnet/Mask** is the full range of IP addresses that you make available for your network.
- **Subnet Mask** is a mask that specifies the default number of IP addresses reserved for any subnet. For example, entering **24** reserves 256 IP addresses.
- **Start Offset** specifies how many addresses not to allocate at the beginning of the subnet's range. For example, entering **10** means that the first ten IP addresses in the subnet aren't available.
- **End Offset** specifies how many IP addresses are not available at the end of the subnet's range.
- **Default lease** and **Maximum lease** specify, in hours, how long an interface can keep a DHCP-assigned IP address.
- **Default gateway**, when selected, indicates that
- **DNS server(s)** specifies the associated Domain Name System server(s).
- **NTP server(s)** specifies the associated Network Time Protocol server(s).
- **NetBIOS name server(s)** is used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
- The **NetBIOS node type** of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types:
 - **B-node** = 0x01 Broadcast
 - **P-node** = 0x02 Peer (WINS only)
 - **M-node** = 0x04 Mixed (broadcast, then WINS)
 - **H-node** = 0x08 Hybrid (WINS, then broadcast)

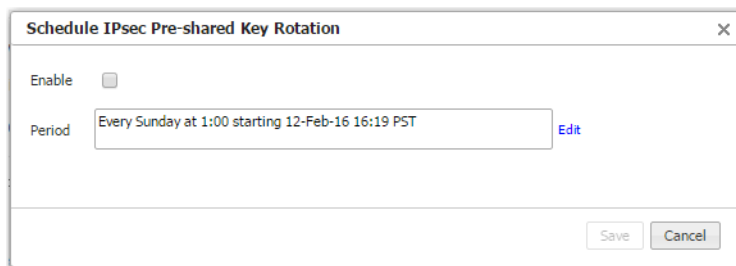
DHCP Relay Definitions

- **Destination DHCP Server** is the IP address of the DHCP server assigning the IP addresses.
- **Enable Option 82**, when selected, inserts additional information into the packet header to identify the client's point of attachment.
- **Option 82 Policy** tells the relay what to do with the hex string it receives. The choices are **append**, **replace**, **forward**, or **discard**.

IPSec Pre-shared Key Rotation

Configuration > [Unity Overlays] IPSec Pre-shared Key Rotation

Use this dialog box to scheduled the rotatation of auto-generated IPsec pre-shared keys.



The screenshot shows a dialog box titled "Schedule IPsec Pre-shared Key Rotation" with a close button (X) in the top right corner. Inside the dialog, there is an "Enable" checkbox which is currently unchecked. Below it, there is a "Period" label followed by a text input field containing the text "Every Sunday at 1:00 starting 12-Feb-16 16:19 PST". To the right of the text field is a blue "Edit" link. At the bottom right of the dialog, there are two buttons: "Save" and "Cancel".

Configuration Wizard

Configuration > [Unity Overlays] Configuration Wizard

Use this wizard to set up a newly added appliance or to reconfigure an appliance that's already in your network.

The screenshot shows the 'Appliance Wizard' window titled 'Appliance Setup'. At the top, there are four numbered steps: 1 (highlighted in blue), 2, 3, and 4. The form contains the following fields:

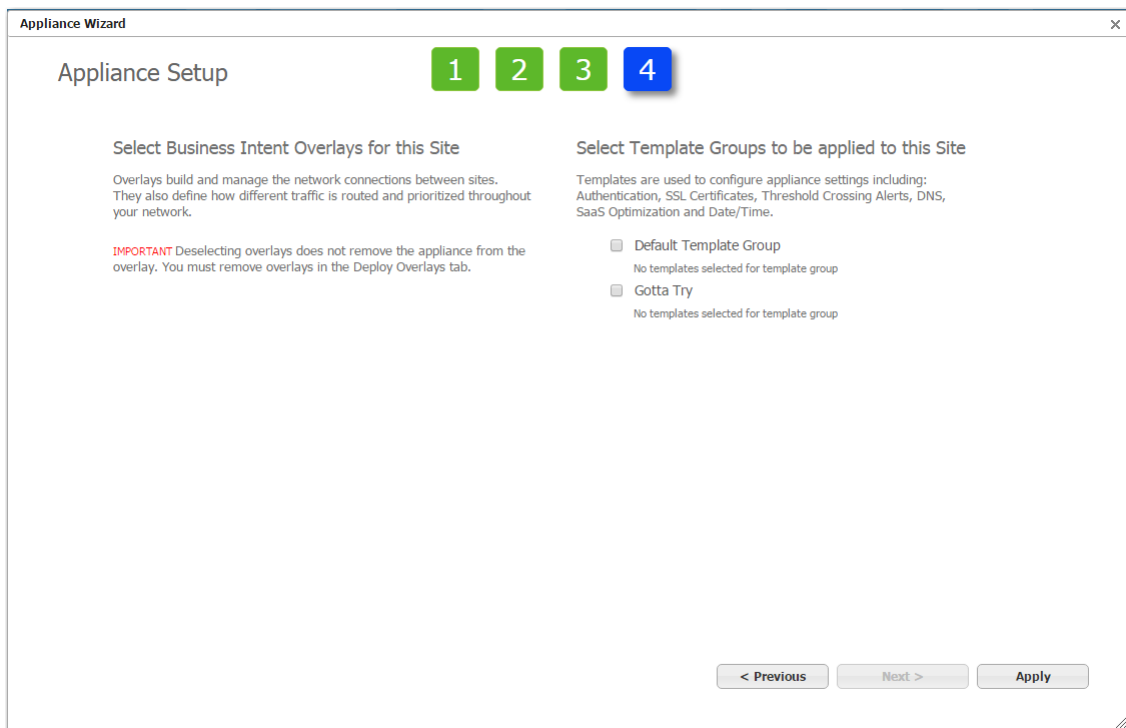
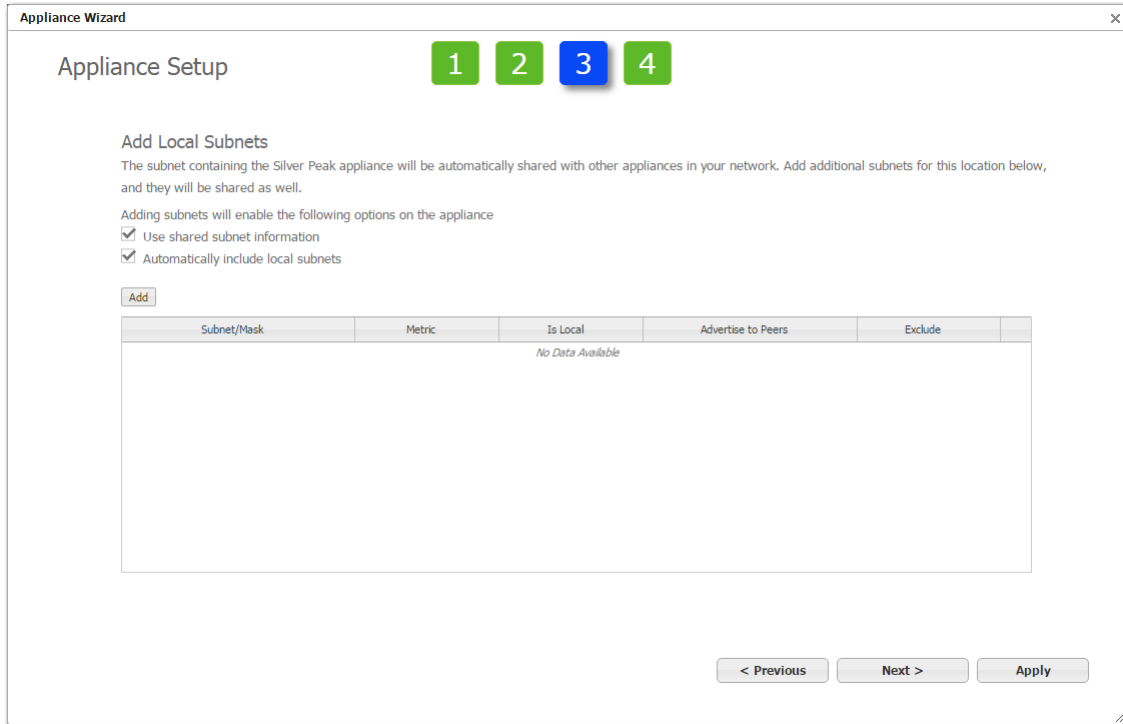
- Hostname*:
- Group*:
- Hub Site?*: Not a hub Hub
- Contact Name:
- Contact Email:
- Serial Number*:
- Site Name:
- Site Priority:
- Location:
 - Address 1:
 - Address 2:
 - City:
 - State:
 - Zip Code:
 - Country:

At the bottom right, there are three buttons: '< Previous', 'Next >', and 'Apply'.

The screenshot shows the 'Appliance Wizard' window titled 'Appliance Setup'. At the top, there are four numbered steps: 1, 2 (highlighted in blue), 3, and 4. The form contains the following field:

- Deployment Profile: ?

At the bottom right, there are three buttons: '< Previous', 'Next >', and 'Apply'.



Configuration Templates

This chapter describes how to use the **Configuration** templates to manage appliances and appliance objects.

It acts as a reference and follows the order of the items in the **Configuration** menu.

In This Chapter

- **Using Configuration Templates** See page 41.
- **System Template** See page 42.
- **Tunnels Template** See page 44.
- **User Defined Apps Template** See page 50.
- **User Defined Apps Template** See page 50.
- **Application Groups Template** See page 52.
- **Access Lists Template** See page 53.
- **Shaper Template** See page 46.
- **QoS Policies Template** See page 57.
- **Optimization Policies Template** See page 61.
- **NAT Policies Template** See page 66.
- **SSL Certificates Template** See page 69.
- **SSL CA Certificates Template** See page 71.
- **SSL for SaaS Template** See page 72.
- **Threshold Crossing Alerts Template** See page 74.
- **Auth/Radius/TACACS+ Template** See page 76.
- **SNMP Template** See page 78.
- **NetFlow Template** See page 80.
- **DNS Template** See page 81.
- **Logging Template** See page 82.
- **Banner Messages Template** See page 84.
- **Cloud Portal Registration Template** See page 85.

- **SaaS Optimization Template** See page 86.
- **VRRP Template** See page 88.
- **CLI Template** See page 89.
- **Session Management Template** See page 90.
- **Default Users Template** See page 91.
- **Date/Time Template** See page 93.

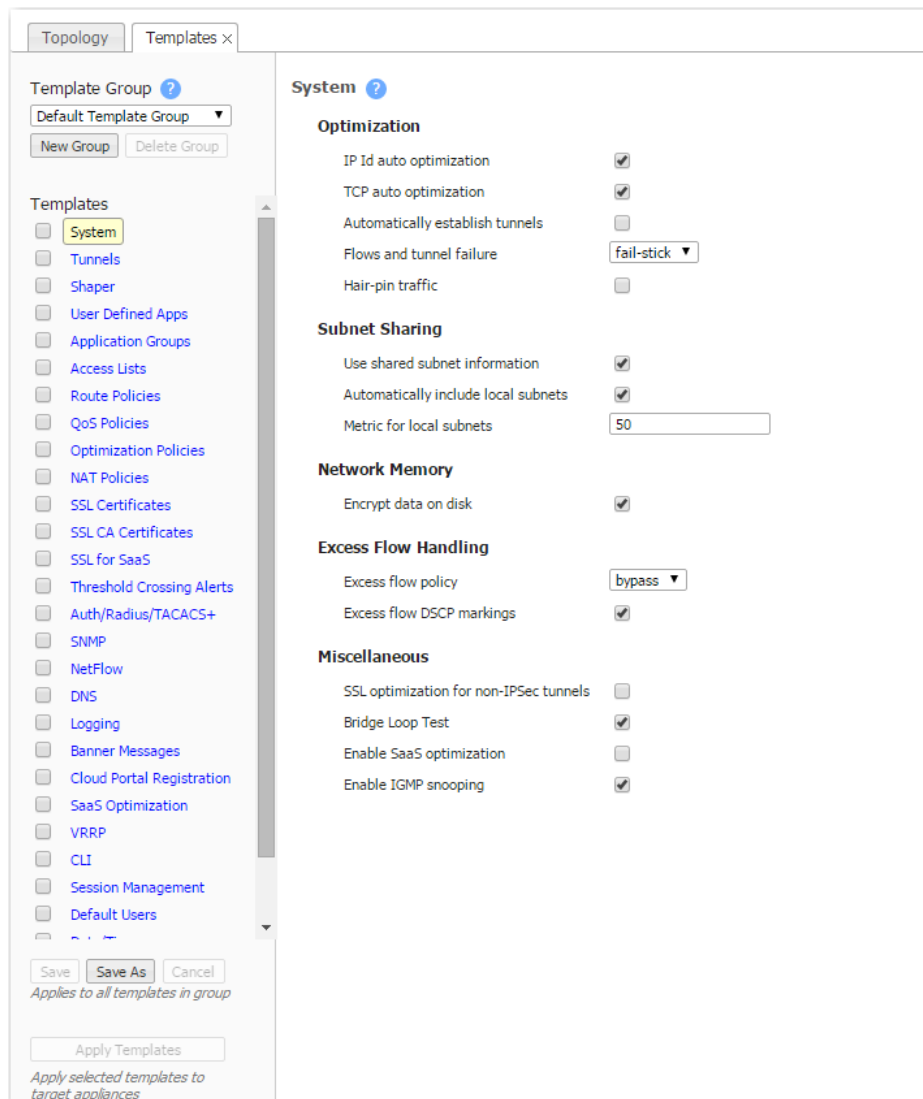
Using Configuration Templates

A *Template Group* is a collection of templates used to configure settings across multiple appliances.

- **IMPORTANT:** Templates will **REPLACE** all settings on the appliance with the template settings unless the template has a **MERGE** option and that option is selected.
- To edit a template, click the template label next to its checkbox.
- You **cannot** save changes to the **Default Template Group**. To save the edits as a new template group, click **Save As**.
- To apply templates to appliances selected in the tree, select the desired template checkbox(es) and click **Apply Templates**. A summary table appears, showing when the selected templates were last applied to the chosen appliances, by whom, and if there are any differences between the previously applied template and the current template. Then you can either click **Apply** or **Close**.
- There is no permanent association between a template and an appliance — it's a one-time, one-way action.
- When returning to the **Templates** page, the **Template Group** field defaults to showing the last template group viewed.
- Unsaved changes display as an icon to the right of the template label.

System Template

Use this page to configure system-level features.



Optimization

- **Optimize traffic** is a global setting for turning optimization on or off. Useful for comparing statistics before and after.
- **IP Id auto optimization** enables any IP flow to automatically identify the outbound tunnel and gain optimization benefits. Enabling this option reduces the number of required static routing rules (route map policies).
- **TCP auto optimization** enables any TCP flow to automatically identify the outbound tunnel and gain optimization benefits. Enabling this option reduces the number of required static routing rules (route map policies).
- **Automatically establish tunnels** reduces configuration overhead by removing the need to manually create tunnels.

- **Flows and tunnel failure.** If there are parallel tunnels and one fails, then *Dynamic Path Control* determines where to send the flows. There are three options:
 - **fail-stick.** When the failed tunnel comes back up, the flows don't return to the original tunnel. They stay where they are.
 - **fail-back.** When the failed tunnel comes back up, the flows return to the original tunnel.
 - **disable.** When the original tunnel fails, the flows aren't routed to another tunnel.

Subnet Sharing

- **Use shared subnet information** enables Silver Peak appliances to use the shared subnet information to route traffic to the appropriate tunnel. Subnet sharing eliminates the need to set up route maps in order to optimize traffic.
- **Automatically include local subnets** adds the local subnet(s) to the appliance subnet information.
- **Metric for local subnets** is a weight that is used for subnets of local interfaces. When a peer has more than one tunnel with a matching subnet, it chooses the tunnel with the greater numerical value.

Network Memory

- **Encrypt data on disk** enables encryption of all the cached data on the disks. Disabling this option is not recommended.

Excess Flow Handling

- **Excess flow policy** specifies what happens to flows when the appliance reaches its maximum capacity for optimizing flows. The default is to **bypass** flows. Or, you can choose to **drop** the packets.
- **Excess flow DSCP markings** specifies whether the appliance should continue to set DSCP markings for flows that are beyond appliance's capacity to optimize.

Miscellaneous

- **SSL optimization for non-IPSec tunnels** specifies if the appliance should perform SSL optimization when the outbound tunnel for SSL packets is not encrypted (for example, a GRE or UDP tunnel). To enable Network Memory for encrypted SSL-based applications, you must provision server certificates via the Unity Orchestrator. This activity can apply to the entire distributed network of Silver Peak appliances, or just to a specified group of appliances.
- **Bridge Loop Test** is only valid for virtual appliances. When enabled, the appliance can detect bridge loops. If it does detect a loop, the appliance stops forwarding traffic and raises an alarm. Appliance alarms include recommended actions.
- **Enable SaaS optimization** enables the appliance to determine what SaaS applications/services it can optimize. It does this by contacting Silver Peak's portal and downloading SaaS IP address and subnet information.
- **Enable IGMP Snooping.** IGMP snooping is a common layer-2 LAN optimization that filters the transmit of multicast frames only to ports where multicast streams have been detected. Disabling this feature floods multicast packets to all ports. IGMP snooping is recommended and enabled by default.

Tunnels Template

Use this template to assign and manage **tunnel properties**.

- Tunnel templates can be applied to any appliances (with or without tunnels). However, only existing tunnels can accept the template settings. To enable an appliance to apply these same settings to future tunnels, select **Make these the Defaults for New Tunnels**.
- Applying tunnel templates **does not** create new tunnels. To create tunnels, use the **Tunnel Builder** tab.
- To **view**, **edit**, and **delete** tunnels, use the **Tunnels** tab.

The screenshot displays the 'Tunnel' configuration page in the Silver Peak Unity Orchestrator. The interface is split into two main sections: a left sidebar for template management and a main configuration area for a specific tunnel template.

Left Sidebar (Template Group and Templates):

- Template Group:** A dropdown menu is set to 'Default Template Group'. Below it are 'New Group' and 'Delete Group' buttons.
- Templates:** A list of template categories with checkboxes: System, Tunnels (highlighted), Shaper, User Defined Apps, Application Groups, Access Lists, Route Policies, QoS Policies, Optimization Policies, NAT Policies, SSL Certificates, SSL CA Certificates, SSL for SaaS, Threshold Crossing Alerts, Auth/Radius/TACACS+, SNMP, NetFlow, DNS, Logging, Banner Messages, Cloud Portal Registration, SaaS Optimization, and VRRP.
- Buttons:** 'Save', 'Save As', and 'Cancel' buttons are present. Below them is the text 'Applies to all templates in group'. At the bottom is an 'Apply Templates' button with the text 'Apply selected templates to target appliances'.

Main Configuration Area (Tunnel ?):

- Make these the Defaults for New Tunnels:** A checked checkbox.
- General:**
 - Admin State: dropdown menu set to 'up'.
 - Mode: dropdown menu set to 'udp'.
 - IPsec Preshared Key: radio button for 'Default' and a password field with '*****'.
 - IPsec Anti-replay Window: dropdown menu set to '1024'.
 - UDP Destination Port: text input field with '4163'.
 - UDP Flows: text input field with '256'.
 - Auto Max BW Enabled: checked checkbox.
- Packet:**
 - Coalescing Enabled: checked checkbox.
 - Coalescing Wait (ms): text input field with '0'.
 - Reorder Wait (ms): text input field with '100' and '(0..500) ms' range.
 - FEC: dropdown menu set to 'disable'.
 - FEC Ratio: dropdown menu set to '1:10'.
 - Auto Discover MTU Enabled: checked checkbox.
 - MTU (bytes): text input field with '1500' and '(700..9000) Bytes' range.
- Tunnel Health:**
 - Retry Count: text input field with '30'.
 - DSCP: dropdown menu set to 'be'.
- FastFail Thresholds:**
 - Fastfail Enabled: dropdown menu set to 'disable'.
 - Latency (ms): text input field with '0' and '(0..65535)' range.
 - Loss (%): text input field with '0' and '(0..10)' range.
 - Jitter (ms): text input field with '0' and '(0..65535)' range.

Definitions (alphabetically)

- **Admin State** brings the tunnel **Up** or **Down**.
- **Auto Discover MTU Enabled** allows an appliance to determine the best MTU to use.
- **Auto Max BW Enabled** allows the appliances to auto-negotiate the maximum tunnel bandwidth.
- **Coalescing Enabled** allows the appliance to coalesce smaller packets into larger packets.
- **Coalescing Wait (ms)** is the number of milliseconds that the appliance should hold packets while attempting to coalesce smaller packets into larger ones.

- **DSCP** determines which DSCP marking the keep-alive messages should use.
- **Fastfail Thresholds** – When multiple tunnels are carrying data between two appliances, this feature determines how quickly to disqualify a tunnel from carrying data.
 - **Fastfail Enabled** – This option is triggered when a tunnel's keepalive signal doesn't receive a reply. The options are **disable**, **enable**, and **continuous**. If the disqualified tunnel subsequently receives a keepalive reply, its recovery is instantaneous.
 - If set to **disable**, keepalives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost.
 - If set to **enable**, keepalives are sent every second, and a missed reply increases the rate at which keepalives are sent from 1 per second to 10 per second. Failover occurs after 1 second.
 - When set to **continuous**, keepalives are continuously sent at 10 per second. Therefore, failover occurs after one tenth of a second.
 - Thresholds for **Latency**, **Loss**, or **Jitter** are checked once every second.
 - Receiving 3 successive measurements in a row that exceed the threshold puts the tunnel into a brownout situation and flows will attempt to fail over to another tunnel within the next 100mS.
 - Receiving 3 successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.
- **FEC** (Forward Error Correction) can be set to **enable**, **disable**, and **auto**.
- **FEC Ratio** is an option when FEC is set to **auto**, that specifies the maximum ratio. The options are 1:2, 1:5, 1:10, or 1:20.
- **IPSec Anti-replay window** provides protection against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet. The decryptor keeps track of which packets it has seen on the basis of these numbers. The default window size is 64 packets.
- **IPSec Preshared Key** is a shared, secret string of Unicode characters that is used for authentication of an IPSec connection between two parties.
- **Mode** determines whether the tunnel is **udp**, **gre**, or **ipsec**. If used, **IPSec** must be enabled at both ends of the tunnel.
- **MTU (bytes)** (Maximum Transmission Unit) is the largest possible unit of data that can be sent on a given physical medium. For example, the default MTU of Ethernet is 1500 bytes. Silver Peak provides support for MTUs up to 9000 bytes.
- **Reorder Wait (ms)** is the number of milliseconds to allow for out-of-order packets to reorder. The default value is 100 ms.
- **Retry Count** is the number of failed keep-alive messages that are allowed before the appliance brings the tunnel down.
- **UDP destination port** is used in UDP mode. Accept the default value unless the port is blocked by a firewall.
- **UDP flows** is the number of flows over which to distribute tunnel data. Accept the default.

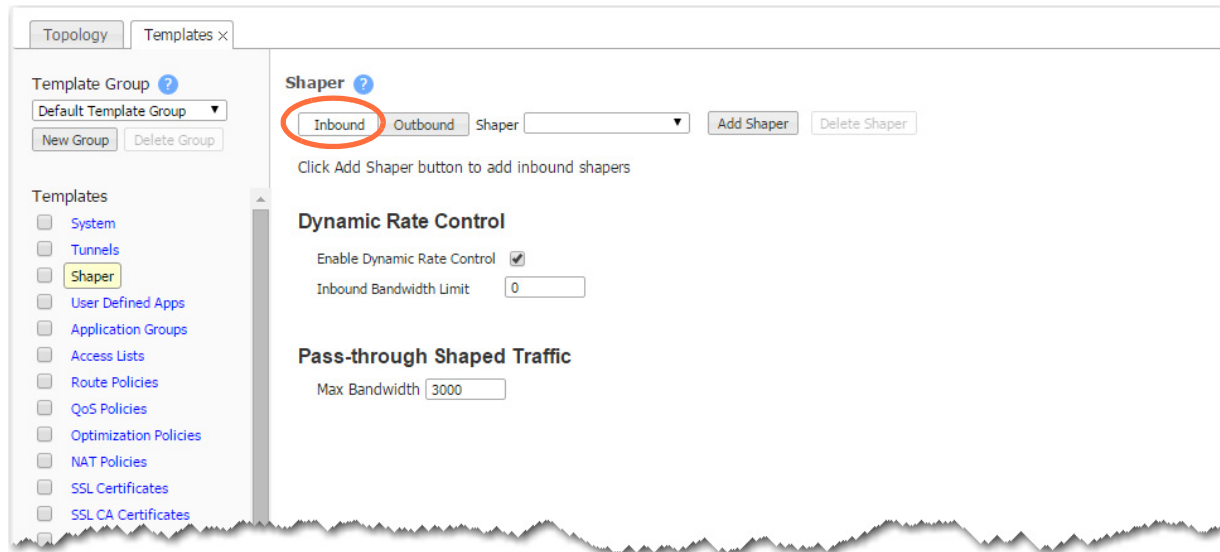
Shaper Template

The **Shaper** template is a simplified way of globally configuring QoS (Quality of Service) on the appliances:

- The Shaper shapes traffic by allocating bandwidth as a percentage of the **system bandwidth**.
- The Shaper's parameters are organized into ten traffic classes. Four traffic classes are preconfigured and named --- **real-time**, **interactive**, **default**, and **best effort**.
- The system applies these QoS settings globally after compressing (deduplicating) all the outbound tunnelized and pass-through-shaped traffic --- shaping it as it exits to the WAN.
- Applying the template to an appliance updates its system-level **wan** Shaper. If the appliance has any added, interface-specific Shapers, they are preserved.
- If you have more than one WAN-side interface, create a shaper for each of those interfaces.
- You can rename or edit any traffic class.
- To view any applied configurations, access the **Configuration > Shaper** page.

The screenshot displays the configuration page for a Shaper template. On the left, a sidebar lists various templates, with 'Shaper' highlighted. The main content area is titled 'Shaper' and includes a dropdown menu for 'Inbound' and 'Outbound' traffic, with 'Outbound' selected and circled in red. Below this, there are buttons for 'Add Shaper', 'Delete Shaper', and 'Enable Shaper'. The 'wan' shaper is selected. A table titled 'wan Traffic Classes' shows 10 traffic classes with columns for ID, Traffic Name, Priority, Min Bandwidth (percentage, kbps), Excess Weighting, Max Bandwidth (percentage, kbps), and Max Wait Time (ms). Below the table is a 'Pass-through Shaped Traffic' section with a 'Max Bandwidth' input field set to 3000.


ID	Traffic Name	Priority	Min Bandwidth		Excess Weighting	Max Bandwidth		Max Wait Time (ms)
			%	kbps		%	kbps	
1	default	5	30	1200	100	100	4000	500
2	real-time	1	30	1200	1000	100	4000	100
3	interactive	2	20	800	1000	100	4000	200
4	best-effort	8	20	800	100	100	4000	500
5		5	30	1200	100	100	4000	500
6		5	30	1200	100	100	4000	500
7		5	30	1200	100	100	4000	500
8		5	30	1200	100	100	4000	500
9		5	30	1200	100	100	4000	500
10		5	30	1200	100	100	4000	500



Dynamic Rate Control

Tunnel Max Bandwidth is the maximum rate at which an appliance can transmit.

Auto BW negotiates the link between a pair of appliances. In this example, the appliances negotiate each link down to the lower value, 100 Mbps.



However, if **A** and **B** transmit at the same time, **Hub** could easily be overrun.

If **Hub** experiences congestion:

- **Enable Dynamic Rate Control.** That allows Hub to regulate the tunnel traffic by lowering each remote appliance's **Tunnel Max Bandwidth**. The smallest possible value is that appliance's **Tunnel Min(imum) Bandwidth**.
- **Inbound BW Limit** caps how much the appliance can receive.

Configuring Max Bandwidth for Pass-through Shaped Traffic

- By default, the values are the same for **Max [WAN] Bandwidth** (for tunnelized traffic) and the **Max Bandwidth** for pass-through shaped traffic.
- However, you can cap the maximum amount of bandwidth allocated to pass-through shaped traffic by entering an upper limit in the **Pass-through Shaped Traffic Max Bandwidth** field.

It's important to note that this is not the same as configuring a percentage of Max WAN BW. This calculation is done after exiting the Shaper, so until that point, **all** shaped packets have queued through the traffic classes as they arrived. As a result, pass-through packets in a higher priority traffic class have a better chance of getting through in the event that the max is exceeded, or if congestion occurs.

Definitions

- **Priority:** Determines the order in which to allocate each class's minimum bandwidth - **1** is first, **10** is last.
- **Min Bandwidth:** Refers to the percentage of bandwidth guaranteed to each traffic class, allocated by priority. However, if the sum of the percentages is greater than 100%, then lower-priority traffic classes might not receive their guaranteed bandwidth if it's all consumed by higher-priority traffic.
If you set **Min Bandwidth** to a value greater than **Max Bandwidth**, then **Max** overrides **Min**.
- **Excess Weighting:** If there is bandwidth left over after satisfying the minimum bandwidth percentages, then the excess is distributed among the traffic classes, in proportion to the weightings specified in the **Excess Weighting** column. Values range from 1 to 10,000.
- **Max Bandwidth:** You can limit the maximum bandwidth that a traffic class uses by specifying a percentage in the **Max Bandwidth** column. The bandwidth usage for the traffic class will never exceed this value.
- **Max Wait Time:** Any packets waiting longer than the specified **Max Wait Time** are dropped.

The Paths Through Policies and Shaping

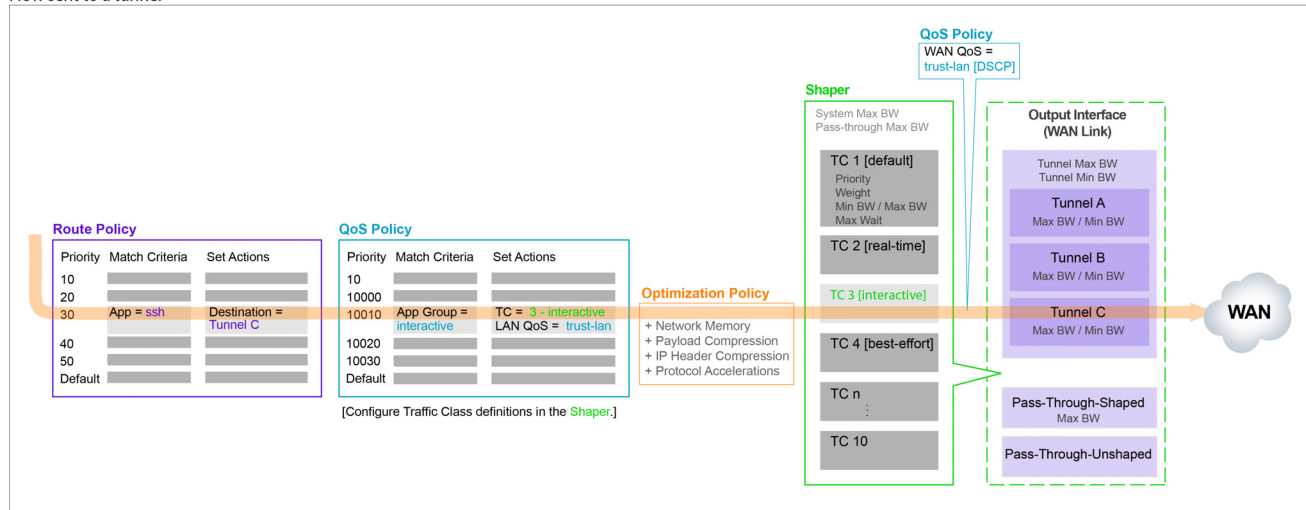
The following diagram illustrates a flow's progress through the policies and the Shaper when the Route Policy Set Action, **Destination**, is:

- a specific tunnel
- pass-through shaped
- pass-through unshaped

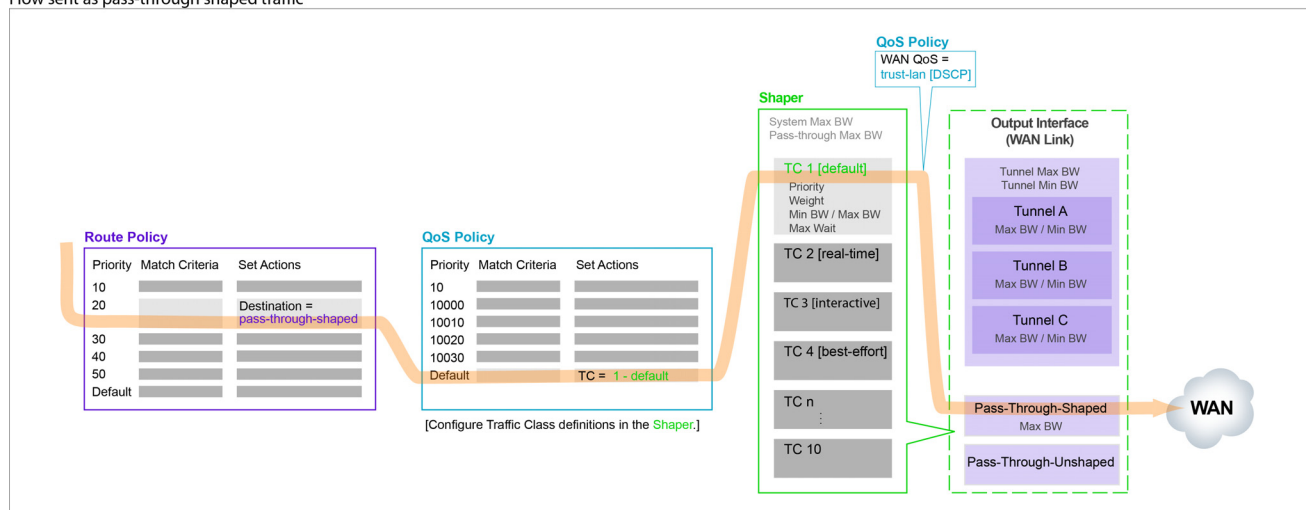


Note If the Route Policy's Set Action is *auto-optimized* and the local appliance initiates either TCP-based or IP-based handshaking, then the remote appliance determines which tunnel to use, based on information it receives in the first packets from the local appliance. (For more information about auto-optimization, see the *Appliance Manager Operator's Guide*.)

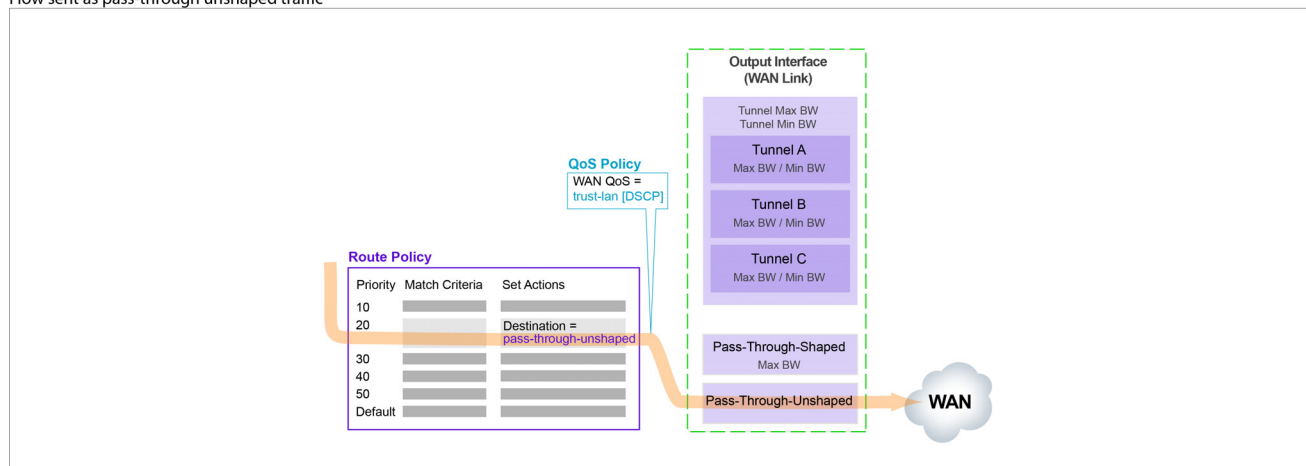
Flow sent to a tunnel



Flow sent as pass-through shaped traffic



Flow sent as pass-through unshaped traffic



User Defined Apps Template

Use this template to create user-defined applications (UDA).

The screenshot displays the 'User Defined Applications' configuration page. On the left, a sidebar lists various templates, with 'User Defined Apps' highlighted. The main area shows a table with one rule defined. The table has columns for Priority, Application, Protocol, Source IP/Subnet, Dest IP/Subnet, Port/Range, DSCP, and Interface. The rule has a priority of 1000, application name 'app1000', protocol 'ip', source IP/Subnet '0.0.0.0/0', destination IP/Subnet '0.0.0.0/0', port/range '0', DSCP 'any', and interface 'any'. Below the table, there are buttons for 'Save', 'Save As', and 'Cancel', and an 'Apply Templates' button.

Priority	Application	Protocol	Source IP/Su...	Dest IP/Subnet	Port/Range	DSCP	Interface
1000	app1000	ip	0.0.0.0/0	0.0.0.0/0	0	any	any

Where can you use them?

- Route Policy
- QoS Policy
- Optimization Policy
- NAT Policy
- Access Lists (ACL)
- Application Groups

Behavior

- For reporting symmetry, you must define the same application(s) on peer appliances. Otherwise, the application may be a UDA on one appliance, and yet be categorized as an **unassigned application** on another, paired appliance.
- Each application consists of at least one rule.
- A warning displays if you reach the maximum number of rules, ports, or addresses allowed.
- If a UDA is in use, deleting it deletes **all** the dependent entries. A warning message appears before deletion.

- Multiple UDAs can have the same name. Whenever that name is referenced, the software sequentially matches against each UDA definition having that name. So, dependent entries are only deleted when you delete the **last** definition of that UDA.



Note When it comes to flow and application statistics reports, user-defined applications are always checked before built-in applications.

Ports are unique. If a port or a range includes a built-in port, then the custom application is the one that lays claim to it.

If two distinctly named user-defined applications have a port number in common, then report results will be skewed, depending on the priority assigned to the custom applications. A port is only counted once.

Priority

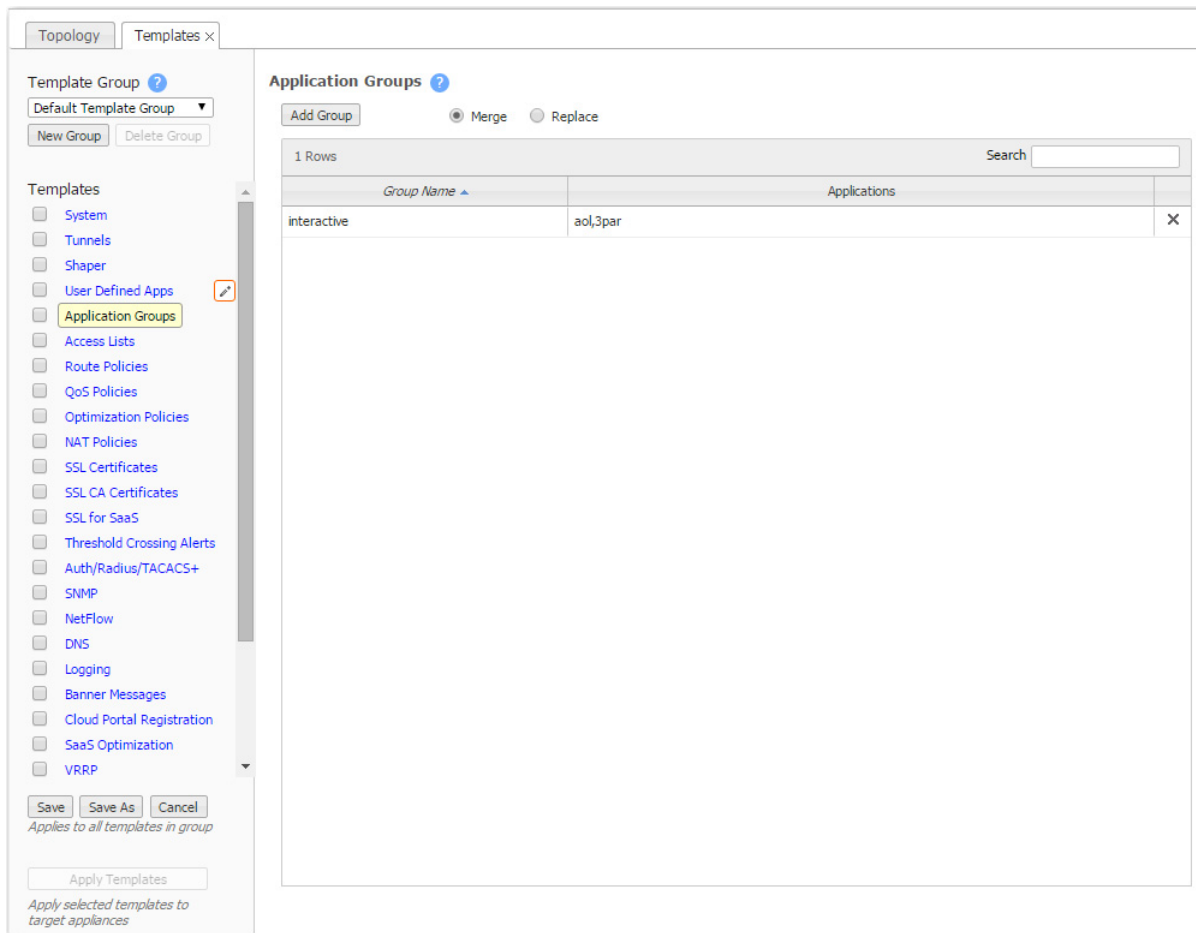
- Range = 1000 – 50000
- Templates won't overwrite or delete applications on the appliances that have priorities in the range, 1 – 999.
- By default, adding a rule/application increments the last Priority by 10.

Source or Destination

- An IP address can specify a subnet - for example: 10.10.10.0/24.
- An IP address can specify a range - for example: 10.10.10.20-30.
- To allow **any IP address**, use 0.0.0.0/0.
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- Specify either a single port or a range of ports - for example: 1234-1250.
- To allow **any port**, use **0**.
- Separate multiple items with any of the following: a line break, a comma, or a single space.

Application Groups Template

Application groups associate applications into a common group that you can use as a MATCH criteria. The applications can be built-in, user-defined, or a combination of both.



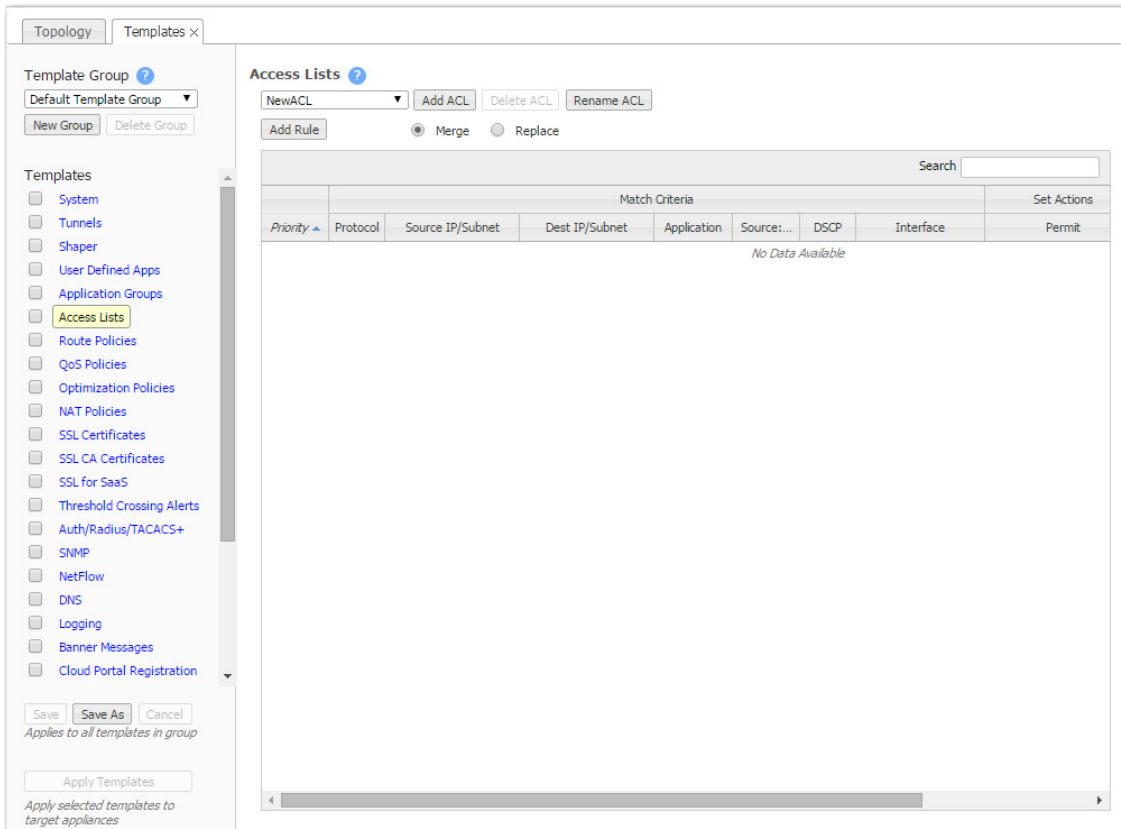
- The **Group Name** cannot be empty or have more than 64 characters.
- Group names are not case-sensitive.
- A group can be empty or contain up to 128 applications.
- An application group cannot contain an application group.
- For reporting symmetry, you must define the same application groups on peer appliances. Otherwise, the application group may be named on one appliance, and yet be categorized as an **unassigned application** on another, paired appliance.

By default, applying the template to an appliance completely deletes and replaces the appliance's application groups.

If you would rather append the template's groups to the appliance's application groups, then select **Merge** before applying the template. If both have a group with the same name, the content will be combined on the appliance.

Access Lists Template

Use this page to create, modify, delete, and rename **Access Control Lists (ACL)**.



An **ACL** is a reusable **MATCH** criteria for filtering flows, and is associated with an action, **permit** or **deny**: You can use the same ACL as the **MATCH** condition in more than one policy --- Route, QoS, Optimization, or NAT.

- An Access Control List (ACL) consists of one or more ordered access control rules.
- An ACL only becomes active when it's used in a policy.
- **Deny** prevents further processing of the flow by *that ACL, specifically*. The appliance continues to the next entry in the policy.
- **Permit** allows the matching traffic flow to proceed on to the policy entry's associated SET action(s). The default is **permit**.
- When creating ACL rules, list **deny** statements first, and prioritize less restrictive rules ahead of more restrictive rules.

Priority

- With this template, you can create rules with priority from **1000 – 9999**, inclusive. When you apply the template to an appliance, the Orchestrator deletes all appliance entries in that range before applying its policies.
- If you access an appliance directly (via the WebUI or the command line interface), you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 65534**).
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- 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**.

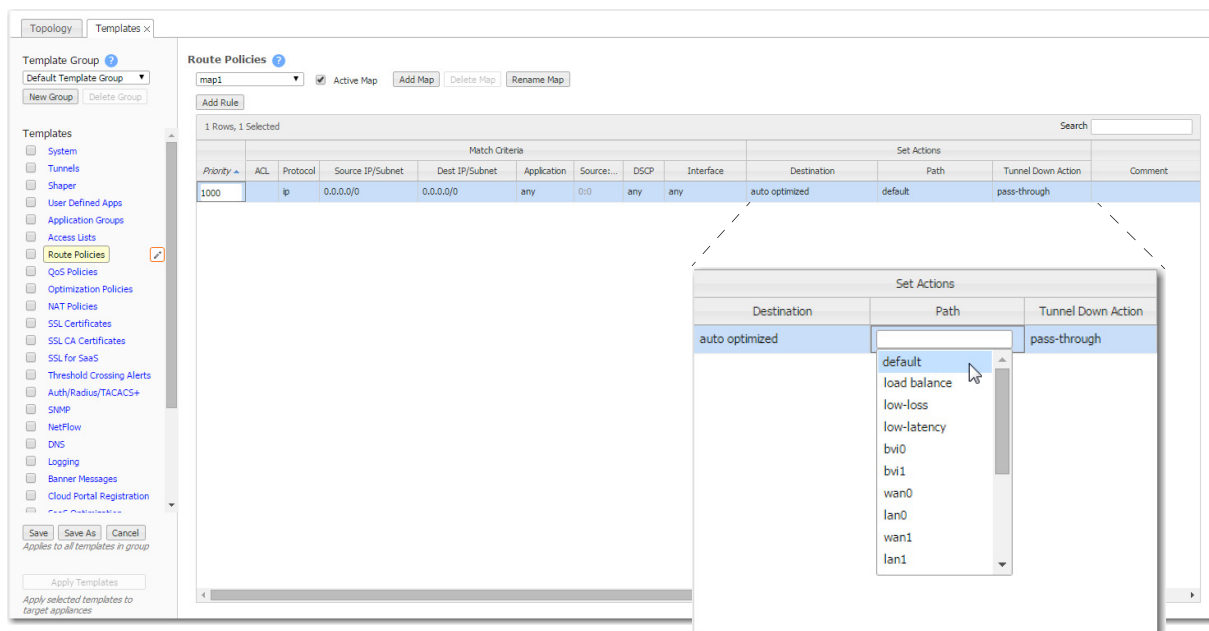
Route Policies Template

Only use the Route Policy template to create (and apply) rules for flows that are to be:

- sent pass-through (shaped or unshaped)
- dropped
- configured for a specific high-availability deployment
- routed based on application, ports, VLAN, DSCP, or ACL (Access Control List)

You may also want to create a Route Policy entry when multiple tunnels exist to the remote *peer*, and you want the appliance to dynamically select the best path based on one of these criteria:

- load balancing
- lowest loss
- lowest latency
- a preferred interface
- a specific tunnel



Why?

Each appliance's default routing behavior is to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. **Auto-optimization** strategies reduce the need to create explicit route map entries for optimization. The three strategies that Silver Peak uses are **TCP-based** auto-opt, **IP-based** auto-opt, and **subnet sharing**. By default, all three are enabled on the **System** template.

Priority

- With this template, you can create rules with priority from **1000 – 9999**, inclusive. When you apply the template to an appliance, the Orchestrator deletes all appliance Route Policy entries in that range before applying its policies.

- If you access an appliance directly (via the WebUI or the command line interface), you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 65534**).
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- 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**.

Set Actions Definitions

The Route Policy template's SET actions determine:

- where the appliance directs traffic
 - In the **Destination** column, you specify how to characterize the flow. The options are **auto-optimized**, **pass-through [shaped]**, **pass-through-unshaped**, or **dropped**.
 - When **auto-optimized**, a flow is directed to the appropriate tunnel. If you choose, you can specify that the appliance use metrics to dynamically select the best path based on one of these criteria:
 - load balancing
 - lowest loss
 - lowest latency



Note When configuring the Route Policy for an **individual** appliance when multiple tunnels exist to the remote *peer*, you can also select the path based on a preferred interface or a specific tunnel. For further information, see the *Appliance Manager Operator's Guide*.

- how traffic is managed if a tunnel is down
 - A **Tunnel Down Action** can be **pass-through [shaped]**, **pass-through-unshaped**, or **dropped**.



Note When configuring the Route Policy for an **individual** appliance, the **continue** option is available if a specific tunnel is named in the **Tunnel** column. That option enables the appliance to read subsequent entries in the individual Route Policy in the event that the tunnel used in a previous entry goes down. For further information, see the *Appliance Manager Operator's Guide*.

QoS Policies Template

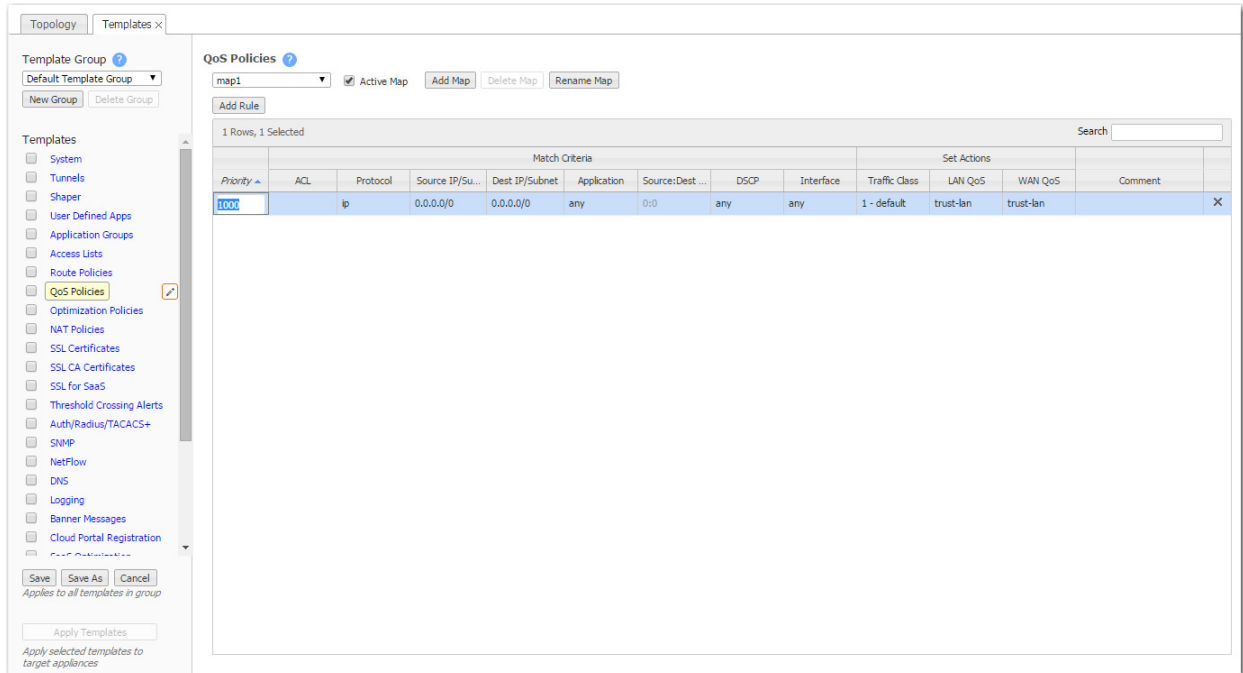
The **QoS Policy** determines how flows are queued and marked.

The QoS Policy's SET actions determine two things:

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

Use the **Shaper** to define, prioritize, and name traffic classes.

Think of it as the Shaper **defines** and the QoS Policy **assigns**.



Priority

- You can create rules with any priority between 1 and 65534.
 - If you are using Orchestrator templates to add route map entries, the Orchestrator will delete all entries from **1000 – 9999**, inclusive, before applying its policies.
 - You can create rules from **1 – 999**, which have higher priority than Orchestrator rules.
 - Similarly, you can create rules from **10000 – 65534** which have lower priority than Orchestrator rules.
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- 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**.

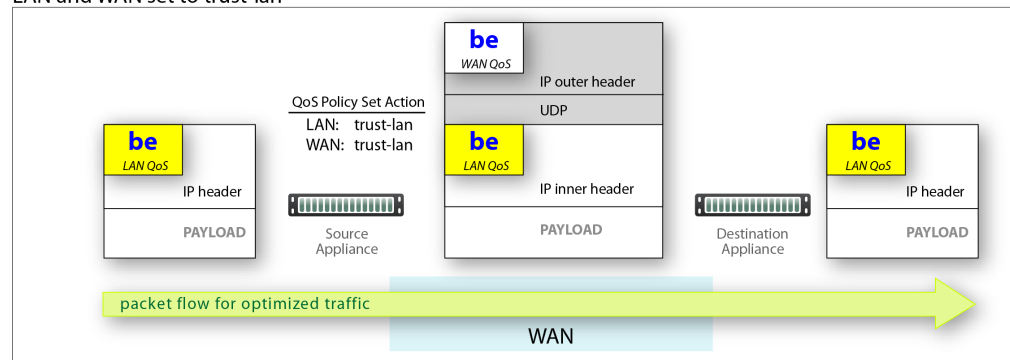
Handling and Marking DSCP Packets

- DSCP markings specify end-to-end QoS policies throughout a network.
- The default values for **LAN QoS** and **WAN QoS** are **trust-lan**.

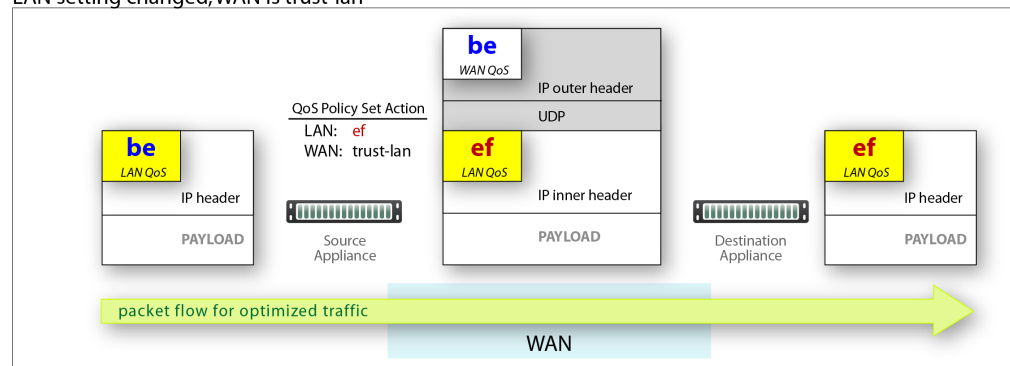
Applying DSCP Markings to Optimized (Tunnelized) Traffic

- The appliance encapsulates optimized traffic. This adds an IP outer header to packets for travel across the WAN. This outer header contains the **WAN QoS** DSCP marking.
- **LAN QoS** – the DSCP marking applied to the IP header before encapsulation
- **WAN QoS** – the DSCP marking in the encapsulating outer IP header. The remote appliance removes the outer IP header.

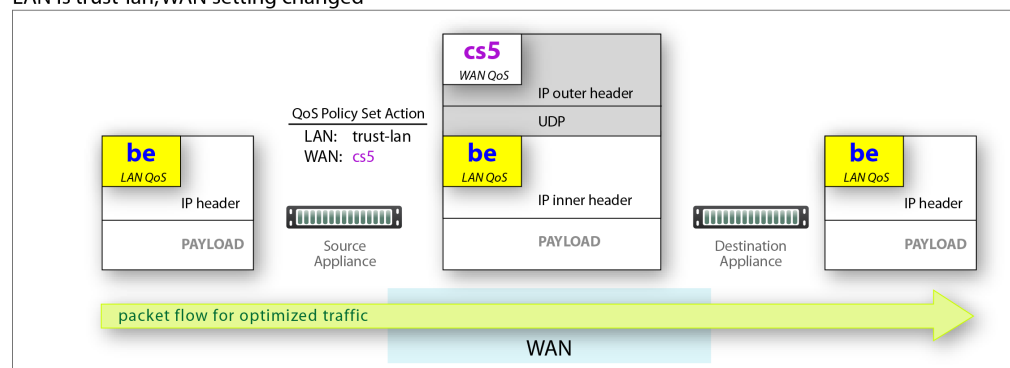
LAN and WAN set to trust-lan



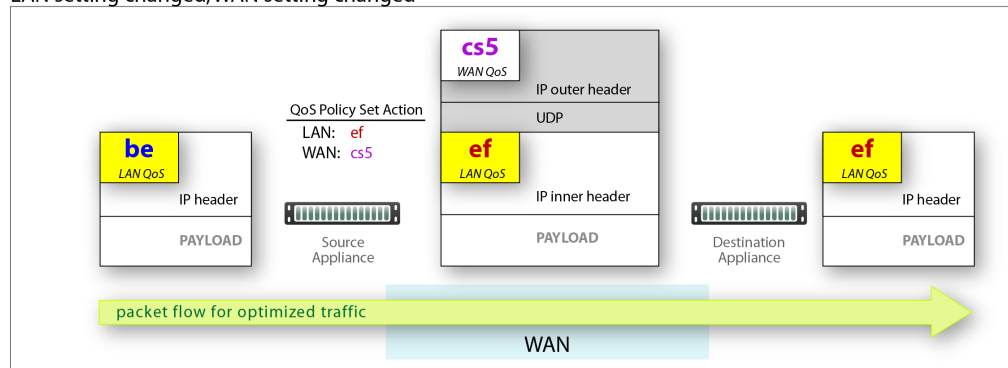
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



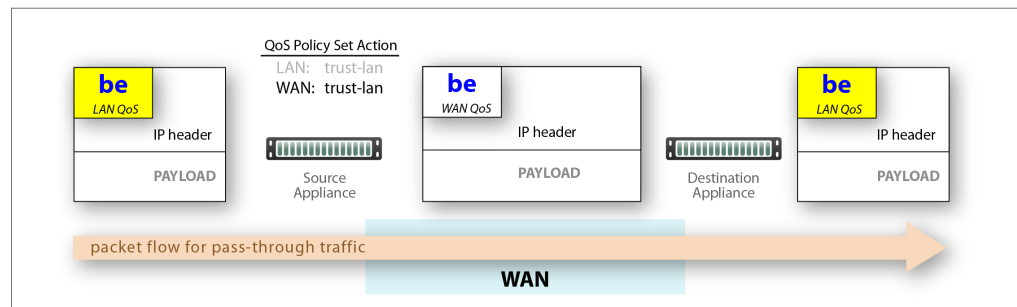
LAN setting changed, WAN setting changed




Applying DSCP Markings to Pass-through Traffic

- The appliance applies the QoS Policy's DSCP markings to all pass-through flows -- shaped and unshaped.
- Pass-through traffic doesn't receive an additional header, so it's handled differently:
 - The Optimization Policy's LAN QoS Set Action is ignored.
 - The specified WAN QoS marking replaces the packet's existing LAN QoS DSCP marking.
 - When the packet reaches the remote appliance, it retains the modified QoS setting as it travels to its destination.

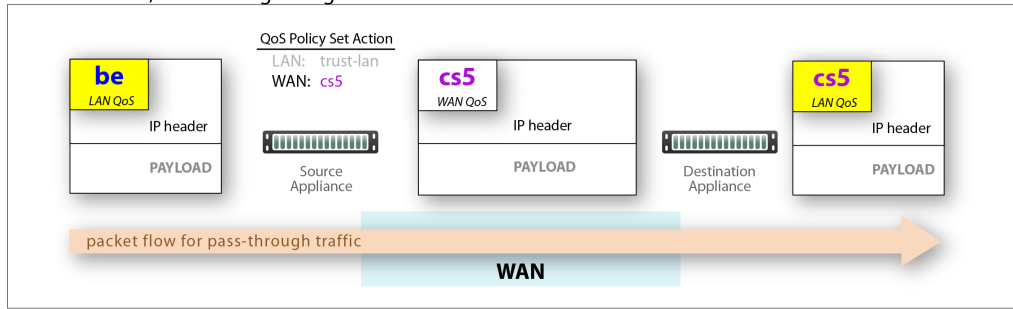
LAN and WAN set to trust-lan



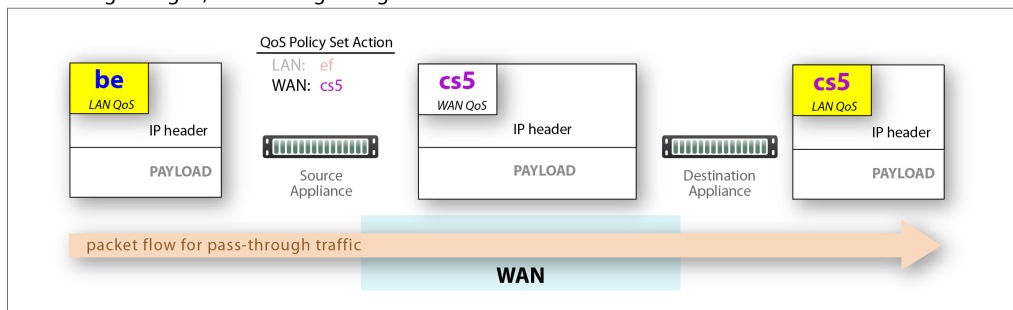
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed

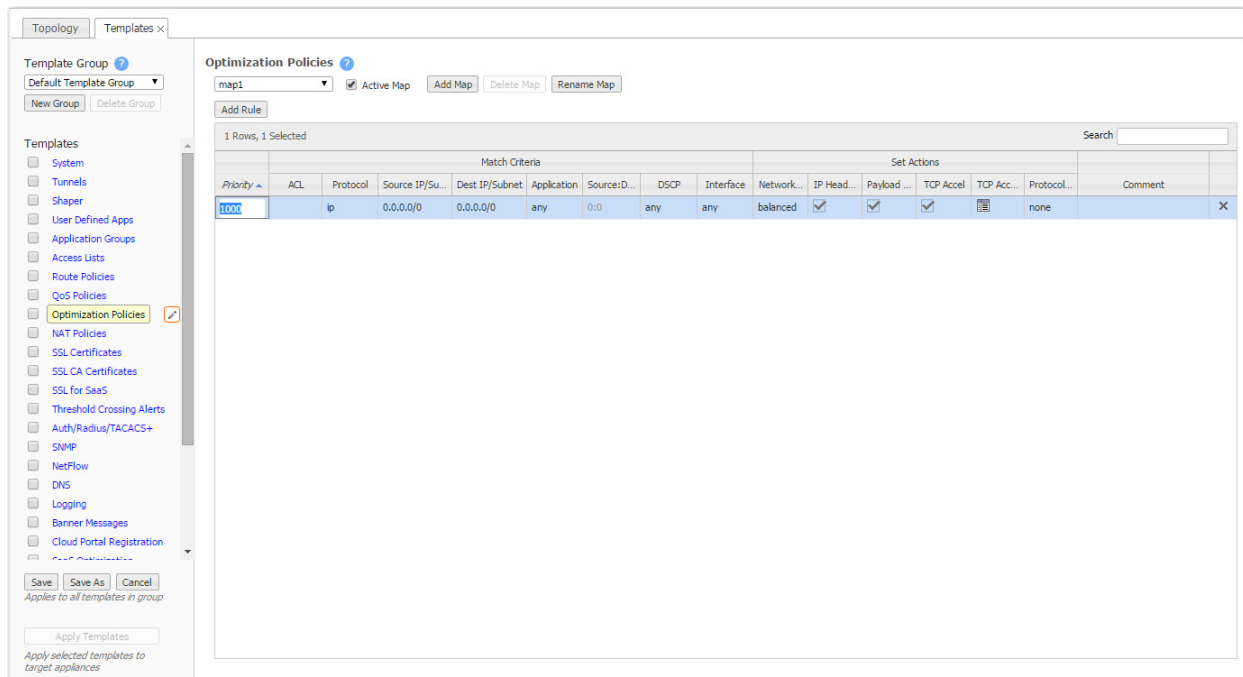


LAN setting changed, WAN setting changed



Optimization Policies Template

Optimization templates apply Optimization policies to appliances.



Priority

- With this template, you can create rules with priority from **1000 – 9999**, inclusive. When you apply the template to an appliance, the Orchestrator deletes all appliance entries in that range before applying its policies.
- If you access an appliance directly (via the WebUI or the command line interface), you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 65534**).
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- 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**.

Set Actions Definitions

- **Network Memory** addresses limited bandwidth. This technology uses advanced fingerprinting algorithms to examine all incoming and outgoing WAN traffic. Network Memory localizes information and transmits only modifications between locations.
 - **Maximize Reduction** optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP, where bandwidth savings are the primary concern.
 - **Minimize Latency** ensures that Network Memory processing adds no latency. This may come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It's also appropriate when the primary objective is to fully utilize the WAN pipe to increase the LAN-side throughput, as opposed to conserving WAN bandwidth.
 - **Balanced** is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types.
 - **Disabled** turns off Network Memory.
- **IP Header Compression** is the process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end. It's possible to compress the protocol headers due to the redundancy in header fields of the same packet, as well as in consecutive packets of a packet stream.
- **Payload Compression** uses algorithms to identify relatively short byte sequences that are repeated frequently. These are then replaced with shorter segments of code to reduce the size of transmitted data. Simple algorithms can find repeated bytes within a single packet; more sophisticated algorithms can find duplication across packets and even across flows.
- **TCP Acceleration** uses techniques such as selective acknowledgements, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links.
- **Protocol Acceleration** provides explicit configuration for optimizing CIFS, SSL, SRDF, Citrix, and iSCSI protocols. In a network environment, it's possible that not every appliance has the same optimization configurations enabled. Therefore, the site that initiates the flow (the *client*) determines the state of the protocol-specific optimization.

TCP Acceleration Options

TCP acceleration uses techniques such as selective acknowledgement, window scaling, and message segment size adjustment to compensate for poor performance on high latency links.

This feature has a set of advanced options with default values.



CAUTION Because changing these settings can affect service, Silver Peak recommends that you **do not modify** these without direction from Customer Support.

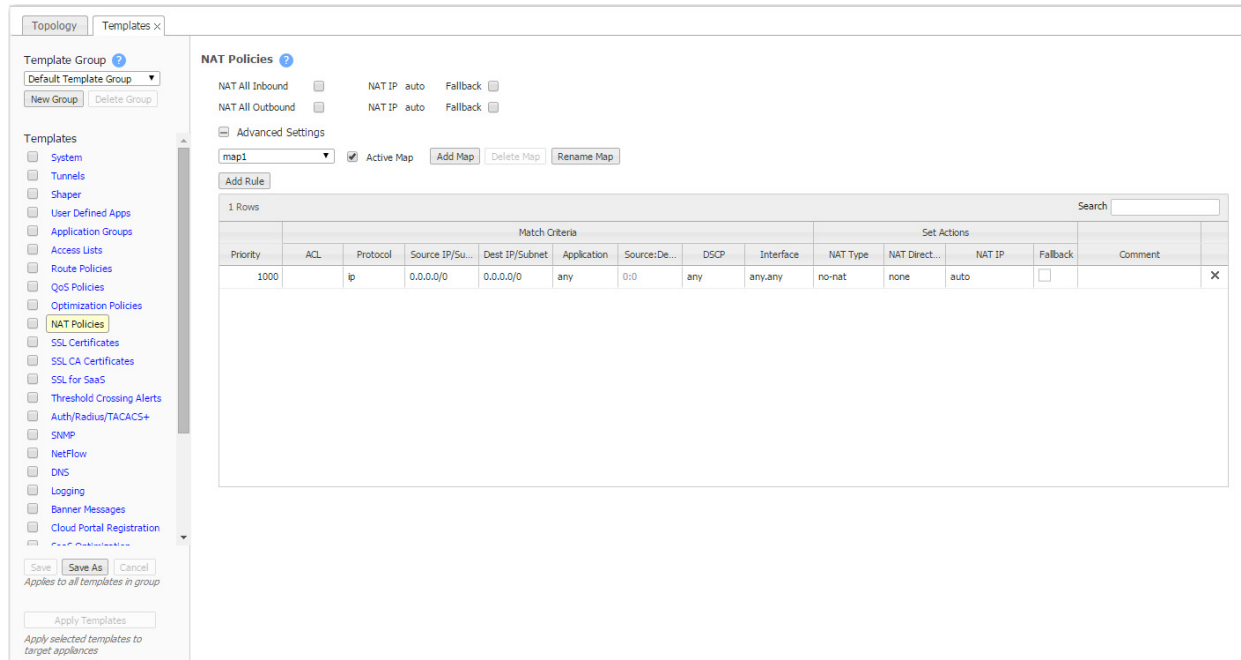
Option	Explanation
Adjust MSS to Tunnel MTU	Limits the TCP MSS (Maximum Segment Size) advertised by the end hosts in the SYN segment to a value derived from the Tunnel MTU (Maximum Transmission Unit). This is $TCP\ MSS = Tunnel\ MTU - Tunnel\ Packet\ Overhead$. This feature is enabled by default so that the maximum value of the end host MSS is always coupled to the Tunnel MSS. If the end host MSS is smaller than the tunnel MSS, then the end host MSS is used instead. A use case for disabling this feature is when the end host uses Jumbo frames.
Preserve Packet Boundaries	Preserves the packet boundaries end to end. If this feature is disabled, then the appliances in the path can coalesce consecutive packets of a flow to use bandwidth more efficiently. It's enabled by default so that applications that require packet boundaries to match don't fail.

Option	Explanation (Continued)
Enable Silver Peak TCP SYN option exchange	<p>Controls whether or not Silver Peak forwards its proprietary TCP SYN option on the LAN side. Enabled by default, this feature detects if there are more than two Silver Peak appliances in the flow's data path, and optimizes accordingly.</p> <p>Disable this feature if there's a LAN-side firewall or a third-party appliance that would drop a SYN packet when it encounters an unfamiliar TCP option.</p>
Route Policy Override	<p>Tries to override asymmetric route policy settings. It emulates auto-opt behavior by using the same tunnel for the returning SYN+ACK as it did for the original SYN packet.</p> <p>Disable this feature if the asymmetric route policy setting is necessary to correctly route packets. In that case, you may need to configure flow redirection to ensure optimization of TCP flows.</p>
Auto Reset Flows	<p>NOTE: Whether this feature is enabled or not, the default behavior when a tunnel goes Down is to automatically reset the flows.</p> <p>If enabled, it resets all TCP flows that aren't accelerated but should be (based on policy and on internal criteria like a Tunnel Up event).</p> <p>The internal criteria can also include:</p> <ul style="list-style-type: none"> • Resetting all TCP accelerated flows on a Tunnel Down event. • Resetting all unaccelerated TCP flows that are associated with a normally operating Tunnel, where: <ul style="list-style-type: none"> - TCP acceleration is enabled - SYN packet was not seen (so this flow was either part of WCCP redirection, or it already existed when the appliance was inserted in the data path).
IP Black Listing	<p>If selected and if the appliance doesn't receive a TCP SYN-ACK from the remote end within 5 seconds, the flow proceeds without acceleration and the destination IP address is blacklisted for one minute.</p>
End to End FIN Handling	<p>This feature helps to fine tune TCP behavior during a connection's graceful shutdown event. When this feature is ON (Default), TCP on the local appliance synchronizes this graceful shutdown of the local LAN side with the remote Silver Peak's LAN side. When this feature is OFF (Default TCP), no such synchronization happens and the two LAN segments at the ends gracefully shutdown independently.</p>
WAN Window Scale	<p>This is the WAN-side TCP Window scale factor that Silver Peak uses internally for its WAN-side traffic. This is independent of the WAN-side factor advertised by the end hosts.</p>
Slow LAN Defense	<p>Resets all flows that consume a disproportionate amount of buffer and have a very slow throughput on the LAN side. Owing to a few slower end hosts or a lossy LAN, these flows affect the performance of all other flows such that no flows see the customary throughput improvement gained through TCP acceleration.</p> <p>This feature is enabled by default. The number relates indirectly to the amount of time the system waits before resetting such slow flows.</p>
WAN Congestion Control	<p>Selects the internal Congestion Control parameter:</p> <ul style="list-style-type: none"> • Optimized - This is the default setting. This mode offers optimized performance in almost all scenarios. • Standard - In some unique cases it may be necessary to downgrade to Standard performance to better interoperate with other flows on the WAN link. • Aggressive - Provides aggressive performance and should be used with caution. Recommended mostly for Data Replication scenarios.
Per-Flow Buffer	<p>(Max LAN to WAN Buffer and Max WAN to LAN Buffer)</p> <p>This setting clamps the maximum buffer space that can be allocated to a flow, in each direction.</p>

Option	Explanation (Continued)
Slow LAN Window Penalty	This setting (OFF by default) penalizes flows that are slow to send data on the LAN side by artificially reducing their TCP receive window. This causes less data to be received and helps to reach a balance with the data sending rate on the LAN side.
LAN Side Window Scale Factor Clamp	This setting allows the appliance to present an artificially lowered WSF to the end host. This reduces the need for memory in scenarios where there are a lot of out-of-order packets being received from the LAN side. These out-of-order packets cause a lot of buffer utilization and maintenance.
Persist timer Timeout	Allows the TCP to terminate connections that are in Persist timeout stage after the configured number of seconds.
Keep Alive Timer	Allows us to change the Keep Alive timer for the TCP connections. <ul style="list-style-type: none"><li data-bbox="591 621 1414 642">• Probe Interval - Time interval in seconds between two consecutive Keep Alive Probes<li data-bbox="591 663 1414 684">• Probe Count - Maximum number of Keep Alive probes to send<li data-bbox="591 705 1414 726">• First Timeout (Idle) - Time interval until the first Keep Alive timeout

NAT Policies Template


Use this template to add NAT map rules to all the appliances that support **Network Address Translation**.



Two use cases illustrate the need for NAT:

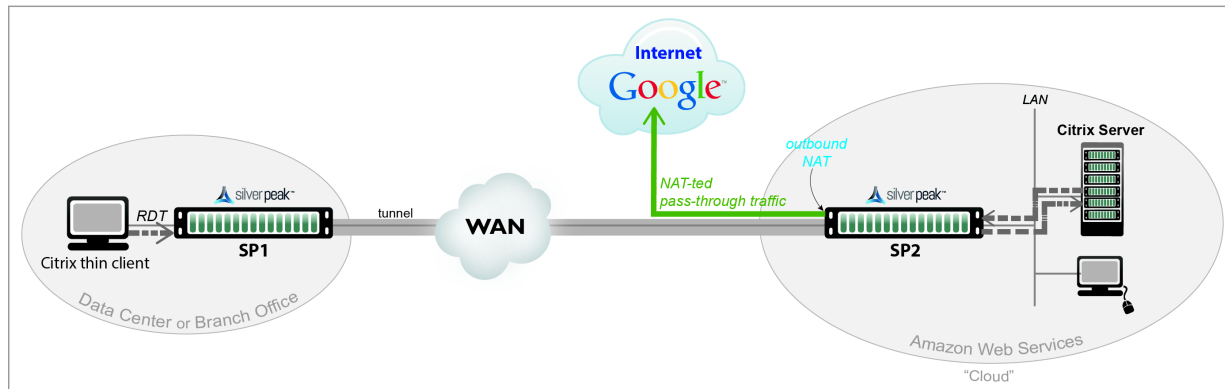
- 1 **Inbound NAT**. The appliance automatically creates a source NAT (Network Address Translation) map when retrieving subnet information from the Silver Peak Cloud portal. This ensures that traffic destined to SaaS servers has a return path to the appliance from which that traffic originated.

NAT with a SaaS Service



- 2 **Outbound NAT.** The appliance and server are in the cloud, and the server accesses the internet. As in the example below, a Citrix thin client accesses its cloud-based server, and the server accesses the internet.

NAT with the Internet



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 addressing requirements.

- Enabling **NAT all** applies NAT policies to pass-through traffic as well as optimized traffic, ensuring that black-holing doesn't occur. **NAT all** on outbound only applies 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. You can do this by deploying the appliance in Router mode in-path with two (or four) interfaces.

Advanced Settings

The appliance can perform **source network address translation** (Source NAT or SNAT) on inbound or outbound traffic.

There are two types of NAT policies:

- **Dynamic** – created automatically by the system for inbound NAT when the **SaaS Optimization** feature is enabled and SaaS service(s) are selected for optimization. The appliance polls the *Silver Peak Unity Cloud Intelligence* service for a directory of SaaS services, and NAT policies are created for each of the subnets associated with selected SaaS service(s), ensuring that traffic destined for servers in use by those SaaS services has a return path to the appliance.
- **Manual** – created by the administrator for specific IP addresses / ranges or subnets. When assigning priority numbers to individual policies within a NAT map, first view **dynamic policies** to ensure that the manual numbering scheme doesn't interfere with dynamic policy numbering (that is, the manually assigned priority numbers cannot be in the range: 4000-5000). The default (**no-NAT**) policy is numbered 65535.

The NAT policy map has the following criteria and **Set Actions**:

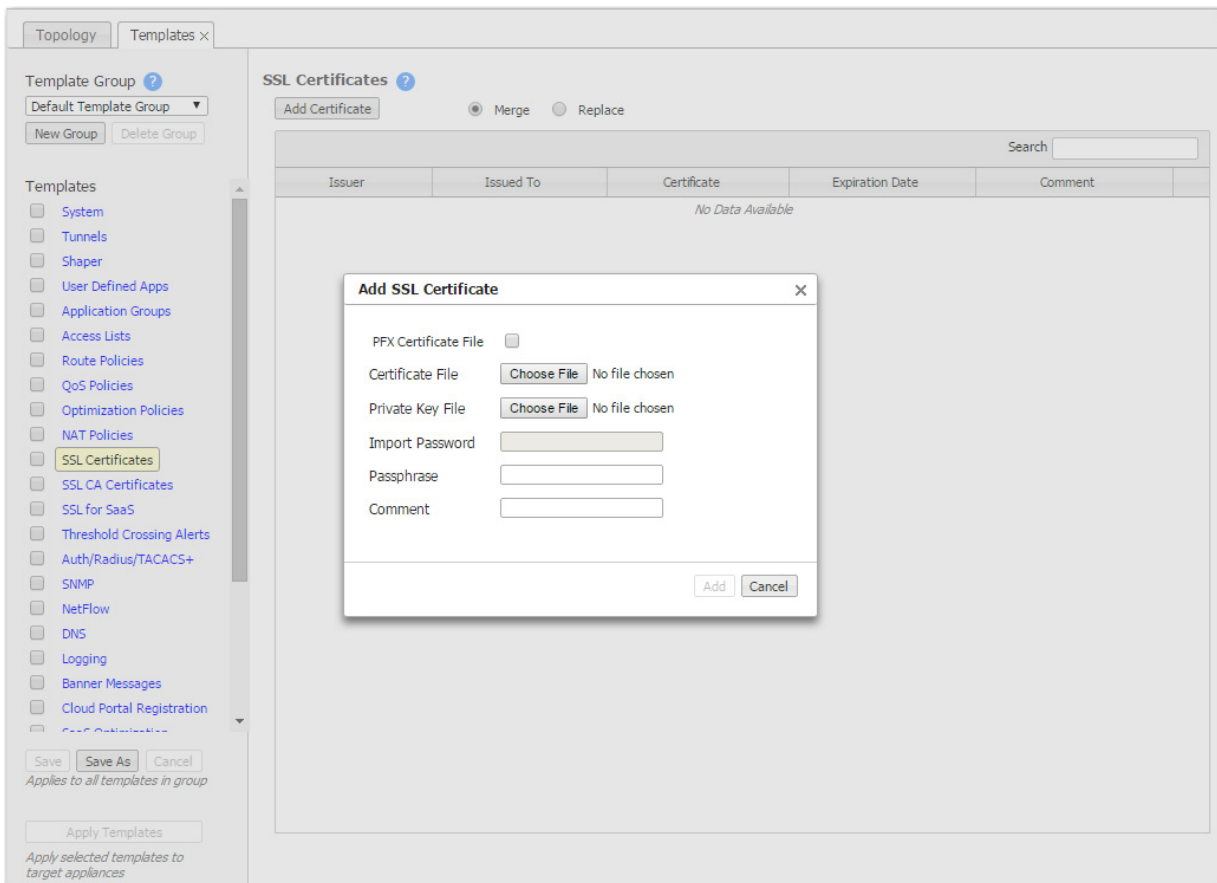
- **Source or Destination**
 - An IP address can specify a subnet - for example: 10.10.10.0/24.
 - To allow **any IP address**, use 0.0.0.0/0.

- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
 - To allow **any port**, use **0**.
- **NAT Type**
 - **no-nat** is the *default*. No IP addresses are changed.
 - **source-nat** changes the source address and the source port in the IP header of a packet.
 - **NAT Direction**
 - **inbound** NAT is on the LAN interface.
 - **outbound** NAT is on the WAN interface.
 - **none** -- the only option if the NAT Type is **no-nat**.
 - **NAT IP**
 - **auto** -- Select if you want to NAT **all** traffic. The appliance then picks the first available NAT IP/Port.
 - **tunnel** -- Select if you only want to NAT **tunnel** traffic. Applicable only for inbound NAT, as outbound doesn't support NAT on tunnel traffic.
 - **[IP address]** -- Select if you want to make NAT use this IP address during address translation.
 - **Fallback** -- If the IP address is full, the appliance uses the next available IP address.

When you select a specific IP, then ensure that the routing is in place for NAT-ted return traffic.

SSL Certificates Template

By supporting the use of SSL certificates and keys, Silver Peak provides deduplication for Secure Socket Layer (SSL) encrypted WAN traffic



- Silver Peak decrypts SSL data using the configured certificates and keys, optimizes the data, and transmits data over an IPsec tunnel. The peer Silver Peak appliance uses configured SSL certificates to re-encrypt data before transmitting.
- Peers that exchange and optimize SSL traffic must use the same certificate and key.
- Use this template to provision a certificate and its associated key across multiple appliances.
 - You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate.
 - The default is PEM when PFX Certificate File is deselected.
 - If the key file has an encrypted key, enter the passphrase needed to decrypt it.
- Silver Peak supports
 - X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
 - SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.
- Silver Peak appliances support:
 - **Protocol versions:** SSLv3, SSLv3.3, TLS1.0, TLS1.1, TLS1.2
 - **Key exchanges:** RSA, DHE, ECDHE

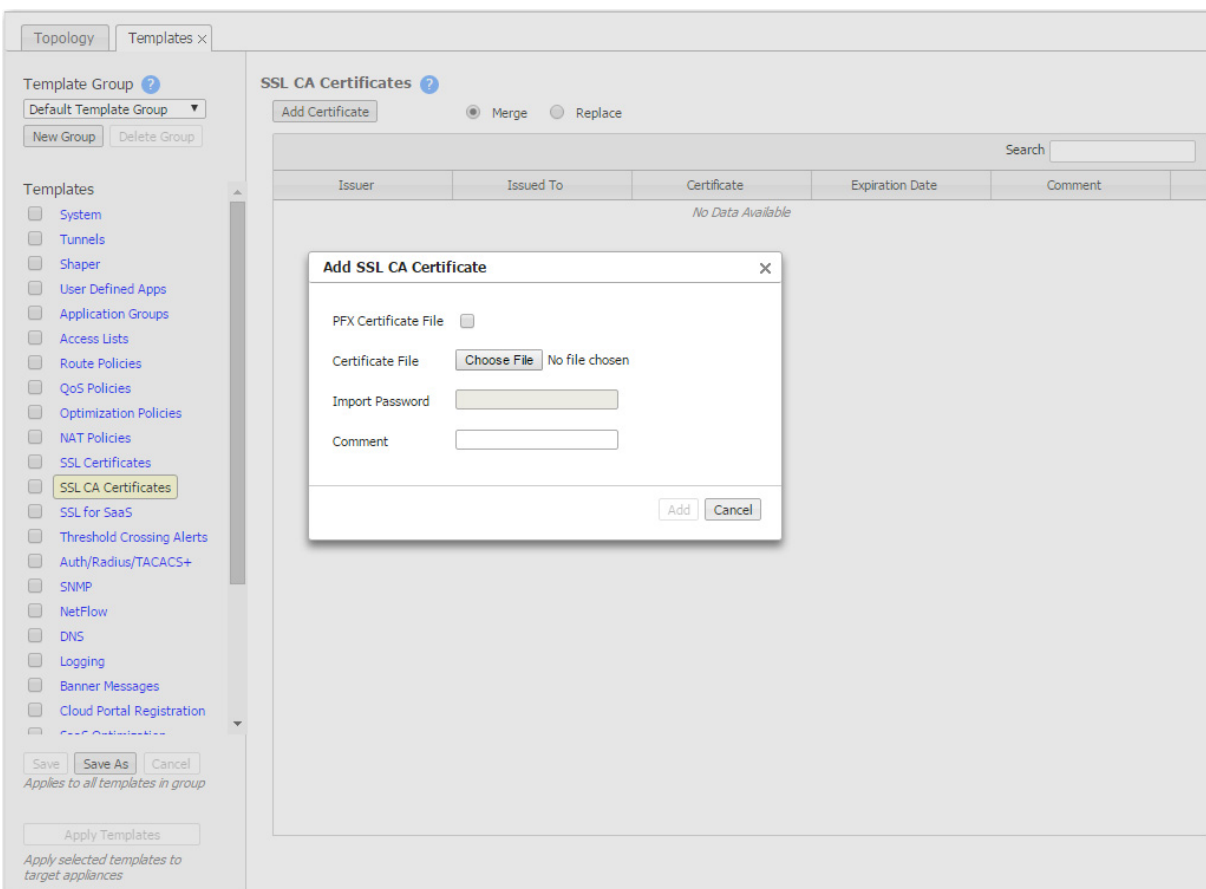
- **Authentication:** RSA
- **Cipher algorithms:** RC4, 3DES, AES128, AES256, AES128-GCM, AES256-GCM
- **Message Digests:** MD5, SHA, SHA256, SHA284
- Before installing the certificates, you must do the following:
 - Configure the tunnels bilaterally for **IPSec** mode.
To do so, access the **Tunnels** template and for **Mode**, select **ipsec**.
 - Verify that **TCP acceleration** and **SSL acceleration** are enabled.
To do so, access the **Configuration > Optimization Policies** page, and review the **Set Actions**.
- If you choose to be able to decrypt the flow, optimize it, and send it in the clear between appliances, then access the **System** template and select **SSL optimization for non-IPsec tunnels**.



Tip For a historical matrix of SSL/TLS versions and ciphers for VXOA releases, click [here](#).

SSL CA Certificates Template

If the enterprise certificate that you used for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, then you must add those CA certificates here. If the browser can't validate up the chain to the root CA, it will warn you that it can't trust the certificate.



- Use this page to directly load the certificate in the Orchestrator.
 - You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate.
 - The default is PEM when PFX Certificate File is deselected.

- Silver Peak supports:
 - X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
 - SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.

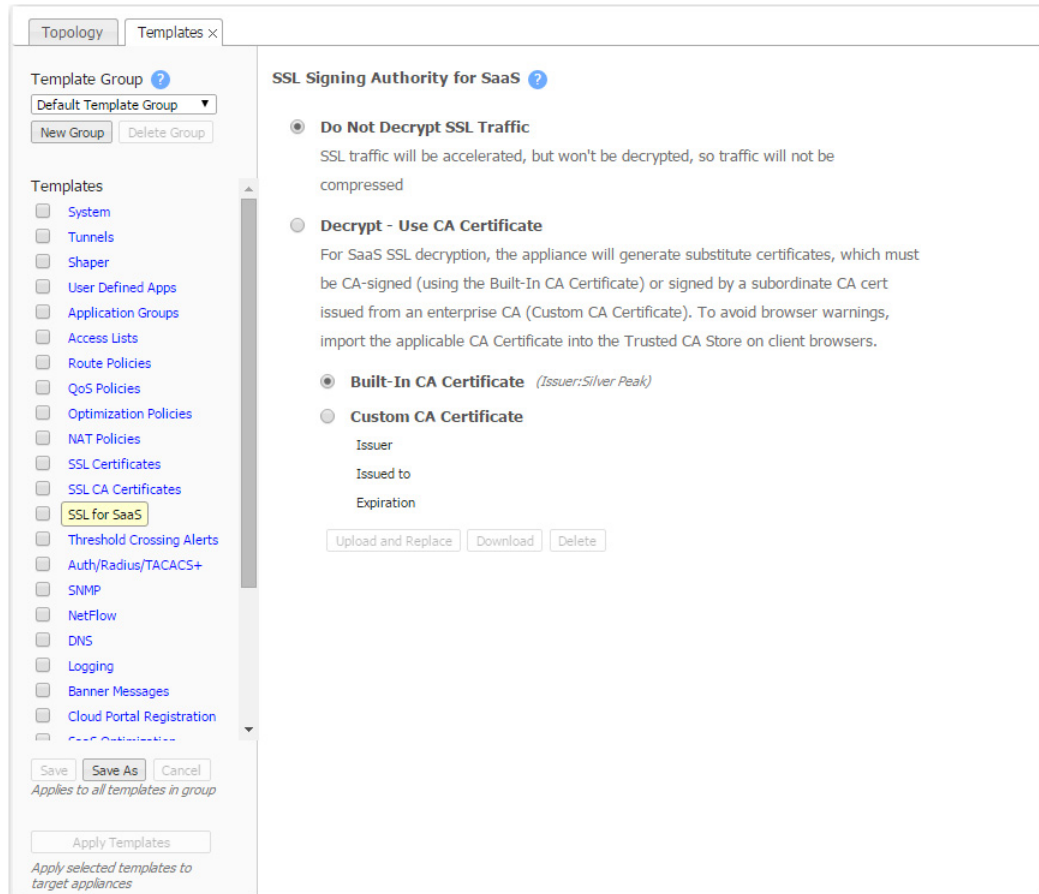


Tip For a historical matrix of SSL/TLS versions and ciphers for VXOA releases, click [here](#).

SSL for SaaS Template

To fully compress SSL traffic for a SaaS service, the appliance must decrypt it and then re-encrypt it.

To do so, the appliance generates a substitute certificate that must then be signed by a Certificate Authority (CA).



There are two possible signers:

- For a *Built-In CA Certificate*, the signing authority is Silver Peak.
 - The appliance generates it locally, and each certificate is unique. This is an ideal option for Proof of Concept (POC) and when compliance is not a big concern.
 - To avoid browser warnings, follow up by importing the certificate into the browser from the client-side appliance.
- For a *Custom CA Certificate*, the signing authority is the Enterprise CA.
 - If you already have a subordinate CA certificate (for example, an SSL proxy), you can upload it to the Orchestrator and push it out to the appliances. If you need a copy of it later, just download it from here.
 - If this substitute certificate is subordinate to a root CA certificate, then also install the higher-level **SSL CA certificates** (into the **SSL CA Certificates** template) so that the browser can validate up the chain to the root CA.
 - If you **don't** already have a subordinate CA certificate, you can access any appliance's **Configuration > SaaS Optimization** page and generate a Certificate Signing Request (CSR).

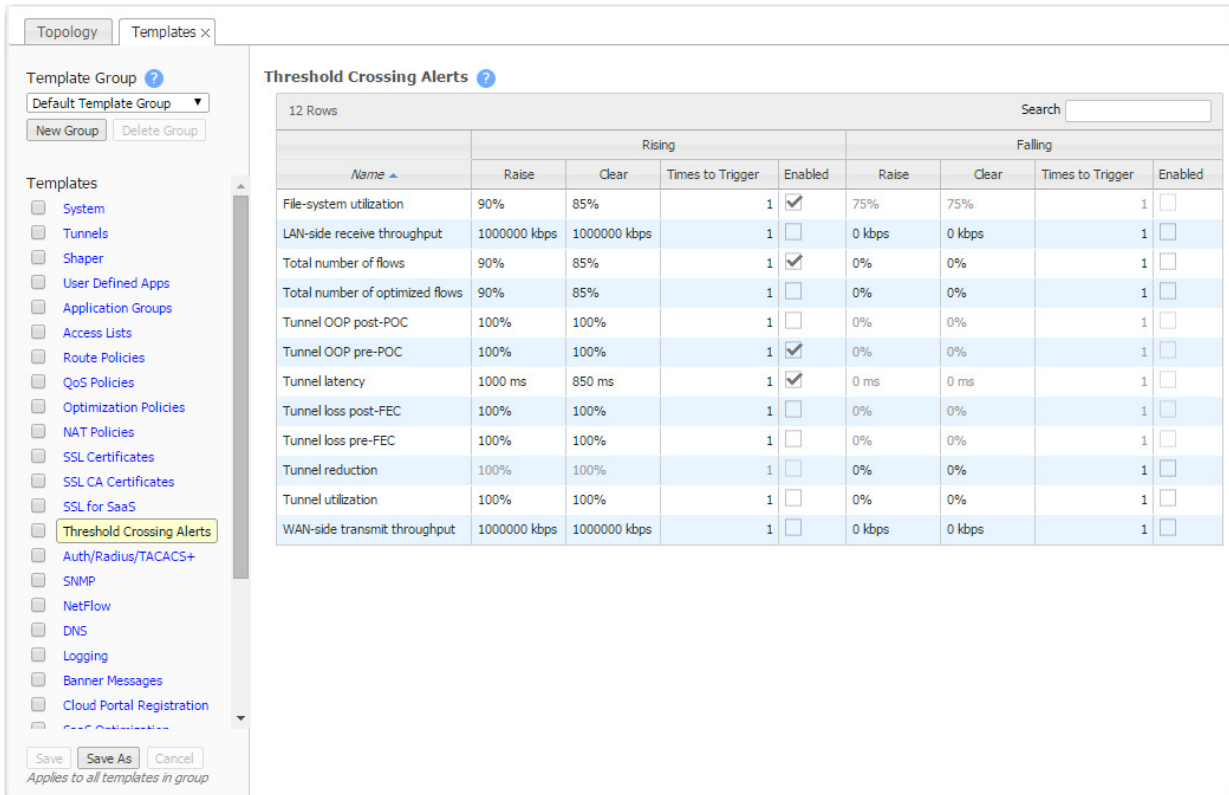
- Silver Peak appliances support:
 - **Protocol versions:** SSLv3, SSLv3.3, TLS1.0, TLS1.1, TLS1.2
 - **Key exchanges:** RSA, DHE, ECDHE
 - **Authentication:** RSA
 - **Cipher algorithms:** RC4, 3DES, AES128, AES256, AES128-GCM, AES256-GCM
 - **Message Digests:** MD5, SHA, SHA256, SHA284



Tip For a historical matrix of SSL/TLS versions and ciphers for VXOA releases, click [here](#).

Threshold Crossing Alerts Template


Threshold Crossing Alerts (TCAs) are preemptive, user-configurable alarms triggered when the specific thresholds are crossed.



The screenshot shows the 'Threshold Crossing Alerts' configuration page. On the left, there is a 'Templates' list with 'Threshold Crossing Alerts' selected. The main area displays a table with 12 rows of alerts. The table is organized into two sections: 'Rising' and 'Falling'. Each section has columns for 'Name', 'Raise', 'Clear', 'Times to Trigger', and 'Enabled'.

Name	Rising				Falling			
	Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
File-system utilization	90%	85%	1	<input checked="" type="checkbox"/>	75%	75%	1	<input type="checkbox"/>
LAN-side receive throughput	1000000 kbps	1000000 kbps	1	<input type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>
Total number of flows	90%	85%	1	<input checked="" type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Total number of optimized flows	90%	85%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel OOP post-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel OOP pre-POC	100%	100%	1	<input checked="" type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel latency	1000 ms	850 ms	1	<input checked="" type="checkbox"/>	0 ms	0 ms	1	<input type="checkbox"/>
Tunnel loss post-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel loss pre-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel reduction	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel utilization	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
WAN-side transmit throughput	1000000 kbps	1000000 kbps	1	<input type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>

They alarm on both rising and falling threshold crossing events (i.e., floor and ceiling levels). For both levels, one value raises the alarm, while another value clears it.



Rules:

- High raise threshold is greater than high clear threshold
- Low raise threshold is less than low clear threshold

Metrics and Defaults

Times to Trigger – A value of 1 triggers an alarm on the first threshold crossing instance. The default sampling granularity (or *rate* or *interval*) is one minute.

This table lists the **metrics** of each type of threshold crossing alert:

Table 3-1 Metrics for Threshold Crossing Alerts

	TCA Name	Unit	Metric
Appliance Level	WAN-side transmit throughput	kbps	Minute average WAN-side transmit TOTAL for all interfaces
	LAN-side receive throughput	kbps	Minute average LAN-side receive TOTAL for all interfaces
	Total number of optimized flows	flows	End of minute count
	Total number of flows	flows	End of minute count
	File-system-utilization	% (non-Network Memory)	End of minute count
Tunnel Level	Tunnel latency	msec	Second-sampled maximum latency during the minute
	Tunnel loss pre-FEC	1/10 th %	Minute average
	Tunnel loss post-FEC	1/10 th %	Minute average
	Tunnel OOP pre-POC	1/10 th %	Minute average
	Tunnel OOP post-POC	1/10 th %	Minute average
	Tunnel utilization	% of configured bandwidth	Minute average
	Tunnel reduction	%	Minute average



Note Enabled by default, there is also an **Appliance Capacity** TCA that triggers when an appliance reaches 95% of its total flow capacity. It doesn't automatically clear, but can be cleared by an operator. It is also not configurable.

Auth/Radius/TACACS+ Template

Silver Peak appliances support user **authentication** and **authorization** as a condition of providing access rights.

- **Authentication** is the process of validating that the end user, or a device, is who they claim to be.
- **Authorization** is the action of determining what a user is allowed to do. Generally, authentication precedes authorization.
- **Map order** refers to the order in which the authentication databases are queried.
- The configuration specified for authentication and authorization **applies globally** to all users accessing that appliance.
- If a logged-in user is inactive for an interval that exceeds the inactivity time-out, the appliance logs them out and returns them to the login page. You can change that value, as well as the maximum number of sessions, in the **Session Management template**.

Authentication and Authorization

To provide authentication and authorization services, Silver Peak appliances:

- support a built-in, **local database**
- can be linked to a **RADIUS** (Remote Address Dial-In User Service) server
- can be linked to a **TACACS+** (Terminal Access Controller Access Control System) server.

Both RADIUS and TACACS+ are client-server protocols.

Appliance-based User Database

- The local, built-in user database supports user names, groups, and passwords.
- The two user groups are **admin** and **monitor**. You must associate each user name with one or the other. Neither group can be modified or deleted.
- The **monitor** group supports reading and monitoring of all data, in addition to performing all actions. This is equivalent to the Command Line Interface's (CLI) **enable** mode privileges.
- The **admin** group supports full privileges, along with permission to add, modify, and delete. This is equivalent to the Command Line Interface's (CLI) **configuration** mode privileges.

RADIUS

- RADIUS uses UDP as its transport.
- With RADIUS, the authentication and authorization functions are coupled together.
- RADIUS authentication requests must be accompanied by a shared secret. The shared secret must be the same as defined in the RADIUS setup. Please see your RADIUS documentation for details.
- **Important:** Configure your RADIUS server's **priv levels** within the following ranges:
 - **admin** = 7 - 15
 - **monitor** = 1 - 6

TACACS+

- TACACS+ uses TCP as its transport.
- TACACS+ provides separated authentication, authorization, and accounting services.
- Transactions between the TACACS+ client and TACACS+ servers are also authenticated through the use of a shared secret. Please see your TACACS+ documentation for details.
- **Important:** Configure your TACACS+ server's roles to be **admin** and **monitor**.

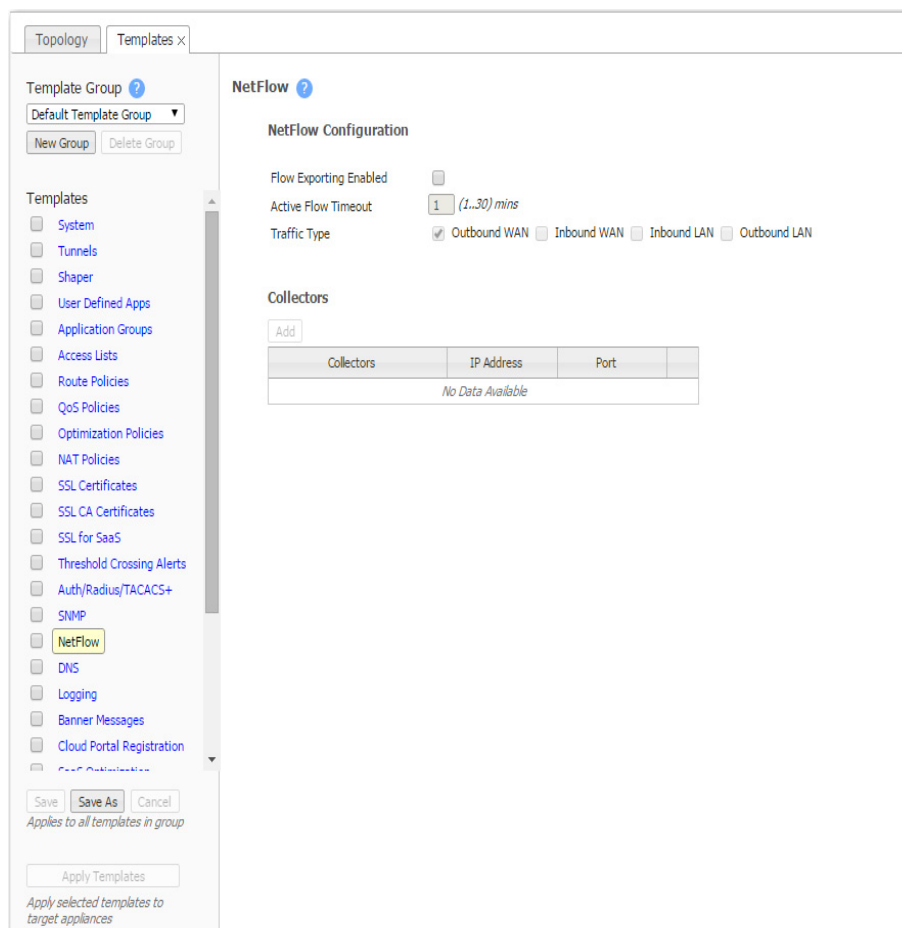
What Silver Peak recommends

- Use either RADIUS or TACACS+, but not both.
- For **Authentication Order**, configure the following:
 - **First** = Local
 - **Second** = either RADIUS or TACACS+. If not using either, then None.
 - **Third** = None
- When using RADIUS or TACACS+ to authenticate users, configure **Authorization Information** as follows:
 - **Map Order** = Remote First
 - **Default User** = admin

SNMP Template

Use this page to configure the appliance's **SNMP** agent, the trap receiver(s), and how to forward appliance alarms as SNMP traps to the receivers.

- The Silver Peak appliance supports the Management Information Base (MIB) II, as described in RFC 1213, for cold start traps and warm start traps, as well as Silver Peak proprietary MIBs.
- The appliance issues an SNMP trap during reset—that is, when loading a new image, recovering from a crash, or rebooting.
- The appliance sends a trap every time an alarm is raised or cleared. Traps contain additional information about the alarm, including severity, sequence number, a text-based description of the alarm, and the time the alarm was created. For additional information, see SILVERPEAK-MGMT-MIB.TXT in the [MIBS directory](#).



For **SNMP v1** and **SNMP v2c**, you only need configure the following:

- **Enable SNMP** = Allows the SNMP application to poll this Silver Peak appliance.
- **Enable SNMP Traps** = Allows the SNMP agent (in the appliance) to send traps to the receiver(s).
- **Read-Only Community** = The SNMP application needs to present this text string (secret) in order to poll this appliance's SNMP agent. The default value is **public**, but you can change it.
- **Default Trap Community** = The trap receiver needs to receive this string in order to accept the traps being sent to it. The default value is **public**, but you can change it.

For additional security *when the SNMP application polls the appliance*, you can select **Enable Admin User** for **SNMP v3**, instead of using **v1** or **v2c**. This provides a way to authenticate without using clear text:

- To configure SNMP v3 **admin** privileges, you must be logged in as **admin** in Appliance Manager.
- For SNMP v3, **authentication** between the user and the server acting as the SNMP agent is bilateral and **required**. You can use either the MD5 or SHA-1 hash algorithm.
- Using DES or AES-128 to encrypt for **privacy** is optional. If you don't specify a password, the appliance uses the default privacy algorithm (AES-128) and the same password you specified for authentication.

You can configure up to 3 **trap receivers**:

- **Host** = IP address where you want the traps sent
- **Community** = The trap receiver needs to receive a specific string in order to accept the traps being sent to it. By default, this field is blank because it uses the Default Trap Community string, which has the value, **public**. If the trap receiver you're adding has a different Community string, enter the community string that's configured on the trap receiver.
- **Version** = Select either **v1** (RFC 1157) or **v2c** (RFC 1901) standards. For both, authentication is based on a community string that represents an unencrypted password.
- **Enabled** = When selected, enables this specific trap receiver.

NetFlow Template

You can configure your appliance to export statistical data to NetFlow collectors.

- The appliance exports flows against two virtual interfaces -- **sp_lan** and **sp_wan** -- that accumulate the total of LAN-side and WAN-side traffic, regardless of physical interface.
- These interfaces appear in SNMP and are therefore "discoverable" by NetFlow collectors.
- **Flow Exporting Enabled** allows the appliance to export the data to collectors (and makes the configuration fields accessible).
- The Collector's **IP Address** is the IP address of the device to which you're exporting the NetFlow statistics. The default Collector Port is 2055.
- In **Traffic Type**, you can select as many of the traffic types as you wish. The default is **Outbound WAN**.

DNS Template

A Domain Name Server (DNS) keeps a table of the IP addresses associated with domain names. It allows you to reference locations by domain name, such as **mycompany.com**, instead of using the routable IP address.

- You can configure up to three name servers.
- Under Domain Names, add the network domains to which your appliances belong.

The screenshot shows a web-based configuration interface for DNS. At the top, there are tabs for 'Topology' and 'Templates x'. Below the tabs, there is a 'Template Group' section with a dropdown menu set to 'Default Template Group' and buttons for 'New Group' and 'Delete Group'. A list of templates is shown on the left, with 'DNS' highlighted. The main configuration area is titled 'DNS' and contains two sections: 'Name Servers' and 'Domain Names'. The 'Name Servers' section has three input fields for 'Primary DNS IP address', 'Secondary DNS IP address', and 'Tertiary DNS IP address'. The 'Domain Names' section has an 'Add' button and a table with a header 'Domain Name' and a row containing 'No Data Available'. At the bottom, there are 'Save', 'Save As', and 'Cancel' buttons, and an 'Apply Templates' button with a note: 'Apply selected templates to target appliances'.

Logging Template

Use this template to configure local and remote logging parameters.

Each requires that you specify the minimum severity level of 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.

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.
INFORMATIONAL	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

- You can 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 may 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.

Banner Messages Template

- The **Login Message** appears before the login prompt.
- The **Message of the Day** appears after a successful login.

The screenshot shows the 'Banner Messages' configuration page in the Unity Orchestrator. The interface is divided into two main sections: a left sidebar for template management and a main content area for editing messages.

Left Sidebar (Template Group):

- Buttons: **New Group**, **Delete Group**
- Dropdown: **Default Template Group**
- List of template categories with checkboxes:
 - SSL for SaaS
 - Threshold Crossing Alerts
 - Auth/Radius/TACACS+
 - SNMP
 - NetFlow
 - DNS
 - Logging
 - Banner Messages** (highlighted)
 - Cloud Portal Registration
 - SaaS Optimization
 - VRRP
 - CLI
 - Session Management
 - Default Users
 - Date/Time
- Instructions: *Check to apply* (with up arrow), *Click link to edit* (with up arrow)
- Buttons: **Save**, **Save As**, **Cancel**
- Note: *Applies to all templates in group*
- Button: **Apply Templates**
- Note: *Apply selected templates to target appliances*

Main Content Area (Banner Messages):

- Section: **Login Message**
- Text area: **Your Login Message**
- Section: **Message of the Day**
- Text area: **Your Message of the Day**

Cloud Portal Registration Template

Each Silver Peak appliance that uses Cloud-based features or products must register with the portal.

The screenshot shows a web-based configuration interface for a Cloud Portal Registration Template. The interface is divided into two main sections: a left sidebar for template selection and a right main area for configuration details.

Left Sidebar (Template Group):

- Buttons: **New Group**, **Delete Group**
- Dropdown: **Default Template Group**
- List of templates (each with a checkbox):
 - SSL Certificates
 - SSL CA Certificates
 - SSL for SaaS
 - Threshold Crossing Alerts
 - Auth/Radius/TACACS+
 - SNMP
 - NetFlow
 - DNS
 - Logging
 - Banner Messages
 - Cloud Portal Registration** (highlighted)
 - SaaS Optimization
 - VRRP
 - CLI
 - Session Management
 - Default Users
 - Date/Time
- Instructions: *Check to apply* (with an up arrow) and *Click link to edit* (with an up arrow)
- Buttons: **Save**, **Save As**, **Cancel**
- Note: *Applies to all templates in group*
- Button: **Apply Templates**
- Note: *Apply selected templates to target appliances*

Right Main Area (Cloud Portal):

- Section: **Cloud Portal**
- Fields:
 - Host:**
 - Port:**

- An enterprise or company has a single **Account Key** and **Account Name** for all its appliances.
- Those appliances must have connectivity to the portal.

SaaS Optimization Template

Use this template to select the SaaS applications/services you want to optimize.

To use this template, your Silver Peak appliance must be registered with an **Account Name** and **Account Key** for the SaaS optimization feature.

The screenshot displays the SaaS Optimization configuration page. On the left, a sidebar lists various templates, with 'SaaS Optimization' selected and highlighted in yellow. The main area shows the configuration for this template, including a checked 'Enable SaaS Optimization' checkbox and an 'RTT Calculation Interval' of 1440 minutes. Below this is a table with 29 rows of SaaS applications. Each row has columns for 'Application Name', 'Optimize', 'Adverti...', 'RTT Threshold', and 'Domains'. The 'Optimize' and 'Adverti...' columns contain checkboxes, and the 'Domains' column lists the domains for each application.

Application Name	Optimize	Adverti...	RTT Threshold	Domains
Adobe	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	adobe.com
Box	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.app.box.com, *.box.com, *.box.net, *.boxcdn.net, *.boxcloud.com
ConstantContact	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	constantcontact.com
CornerstoneOnDemand	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	cornerstoneondemand.com
Dropbox	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	dropbox.com, *.dropbox.com
Eloqua	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	eloqua.com, eloquatrainingcenter.com
GoToAssist	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com
GoToMeeting	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotomeeting.com
GoToTraining	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com
GoToWebinar	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotowebinar.com, gotoassist.com
Intuit	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	intuit.com
Jobvite	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	careers.jobvite.com, www.jobvite.com, hire.jobvite.com
Lithium	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	lithium.com
LiveOps	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	liveops.com
Marketo	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	marketo.com
NetSuite	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	netsuite.com
Office365	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.officeapps.live.com, *.microsoftonline-p.net, *.microsoftonlinesupport.ne...
Parature	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	parature.com
PardotExactTarget	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	pardot.com
PlexSystems	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	plex.com

SaaS optimization requires three things to work in tandem: **SSL** (Secure Socket Layer), **subnet sharing**, and **Source NAT** (Network Address Translation).

Enable SaaS optimization enables the appliance to contact Silver Peak's *Unity Cloud Intelligence Service* and download information about SaaS services.

- If **Advertise** is *selected* for a service (for example, SFDC), the appliance will:
 - Ping active SaaS subnets to determine RTT/metric
 - Add subnet sharing entries locally for subnets within RTT threshold
 - Advertise subnets and their metric (within threshold) via subnet sharing to client-side appliances
 - Upon seeing an SFDC flow, generate a substitute certificate for an SFDC SSL domain (one substitute certificate per domain)
 - Auto-generate dynamic NAT rules for SFDC (but not for unchecked services)
- When **Optimize** is *selected* for a service (for example, SFDC), the appliance will:

- Ping active SFDC subnets to determine the RTT (metric)
 - Does not advertise metric via subnet sharing (unless **Advertise** is also selected)
 - Receives subnet sharing metric (RTT) from associated appliances
 - Compares its own RTT (local metric) with advertised metric
 - If its own RTT is lower, then the packet is sent pass-through (direct to the SaaS server).
 - If an advertised RTT is lower, then the packet is tunnelized.
 - Generate a substitute certificate for an SFDC SSL domain (one sub cert per domain)
 - No NAT rules created
- When **Optimize** is *not selected* for a service (for example, SFDC), the appliance:
- Receives subnet sharing advertisements for SFDC but doesn't use them
 - Does no RTT calc pinging
 - Does not participate in SSL
 - Creates no NAT rules
 - Sends all SFDC traffic as pass-through

TIPS

- Initially, you may want to set a higher **RTT Threshold** value so that you can see a broader scope of reachable data centers/servers for any given SaaS application/service.
- If the **Monitoring** page shows no results at **50 ms**, you may want to reposition your SaaS gateway (advertising appliance) closer to the service.

VRRP Template

Use this template to distribute common parameters for appliances deployed with **Virtual Router Redundancy Protocol (VRRP)**.

The screenshot displays the VRRP configuration interface. On the left, a 'Template Group' dropdown is set to 'Default Template Group'. Below it are 'New Group' and 'Delete Group' buttons. A list of template categories is shown, with 'VRRP' highlighted. The main configuration area for VRRP includes:

- Admin:** A dropdown menu set to 'Up'.
- Advertisement Timer:** A text input field containing '(1..255)'.
- Priority:** A text input field containing '(1..254)'.
- Preemption:** An unchecked checkbox.
- Authentication String:** An empty text input field.

 At the bottom of the interface, there are 'Save', 'Save As', and 'Cancel' buttons. Below these is an 'Apply Templates' button with the text 'Apply selected templates to target appliances'.

In an out-of-path deployment, one method for redirecting traffic to the Silver Peak appliance is to configure 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.

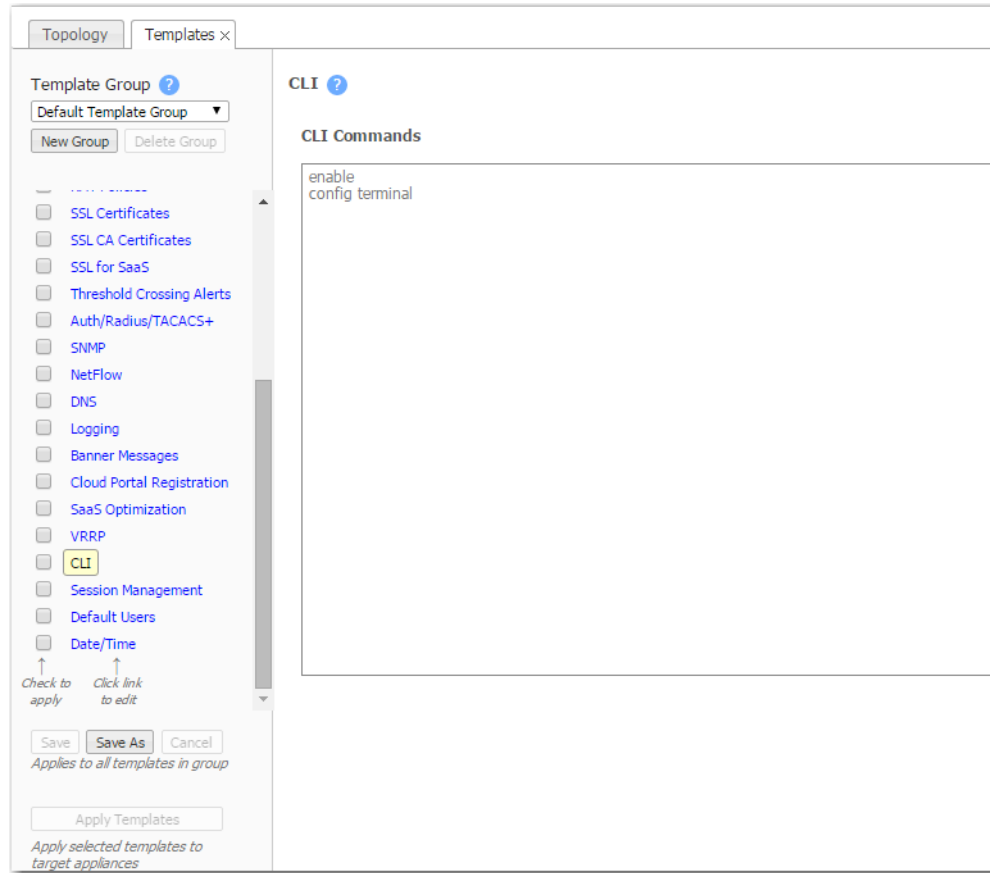
Definitions (alphabetically)

- **Admin** = The options are **up** (enable) and **down** (disable).
- **Advertisement Timer** = default is **1 second**.
- **Authentication String** = Clear text password for authenticating group members.
- **Preemption**. Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
- **Priority**. The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.

CLI Template

Use this template to enter any sequence of **Command Line Interface (CLI)** commands.

Enter each Command Line Interface (CLI) command in a new line.



The screenshot shows a web-based configuration interface for CLI templates. The interface is divided into two main sections: a left sidebar for template management and a main content area for editing the CLI commands.

Left Sidebar (Template Group):

- Template Group: [Default Template Group](#) (dropdown)
- Buttons: [New Group](#), [Delete Group](#)
- List of template categories (each with a checkbox):
 - SSL Certificates
 - SSL CA Certificates
 - SSL for SaaS
 - Threshold Crossing Alerts
 - Auth/Radius/TACACS+
 - SNMP
 - NetFlow
 - DNS
 - Logging
 - Banner Messages
 - Cloud Portal Registration
 - SaaS Optimization
 - VRRP
 - CLI** (highlighted)
 - Session Management
 - Default Users
 - Date/Time
- Instructions: [Check to apply](#), [Click link to edit](#)
- Buttons: [Save](#), [Save As](#), [Cancel](#)
- Text: *Applies to all templates in group*
- Button: [Apply Templates](#)
- Text: *Apply selected templates to target appliances*

Main Content Area (CLI Commands):

- Section: **CLI** (with a help icon)
- Section: **CLI Commands**
- Text area containing the commands:

```
enable
config terminal
```

Session Management Template

Use this page to configure access to the web server.

The screenshot shows the 'Session Management' configuration page. On the left, there is a 'Template Group' section with a dropdown menu set to 'Default Template Group' and buttons for 'New Group' and 'Delete Group'. Below this is a list of template categories with checkboxes, including 'Session Management' which is highlighted. At the bottom of the list are 'Save', 'Save As', and 'Cancel' buttons, with a note 'Applies to all templates in group'. Below these are 'Apply Templates' and 'Apply selected templates to target appliances' buttons. The main 'Session Management' section on the right contains three fields: 'Auto Logout' set to 15 (with a note '(0-60 minutes, 0 indicates no timeout)'), 'Max Session' set to 10 (with a note '(5-50)'), and 'Web Protocol' with radio buttons for 'HTTP', 'HTTPS', and 'Both' (which is selected).

- **Auto Logout** ends your web session after the specified minutes of inactivity.
- If the number of **Max Sessions** is exceeded, there are two possible consequences:
 - You'll get a message that the browser can't access the appliance.
 - Since the Orchestrator must create a session to communicate with the appliance, it won't be able to access the appliance.
- Although **Web Protocol** defaults to **Both** for legacy reasons, Silver Peak recommends that you select **HTTPS** for maximum security.

Default Users Template

Use this page to manage the default users and, if desired, require a password with the highest user privilege level when using the Command Line Interface.

The screenshot shows the 'Default Users' configuration page. On the left, a sidebar lists various template groups, with 'Default Users' highlighted. The main area is titled 'Default Users' and contains a table of user accounts. Below the table, there are options to require a password for CLI 'Enable' privilege and input fields for the password and its confirmation.

User Name	Capability	Password	Confirm Password
admin	admin		
monitor	monitor		

Below the table, the 'Password for CLI "Enable" privilege' section includes a 'Require Password' checkbox and two input fields for 'Password' and 'Confirm Password'.

Default User Accounts

- Each appliance has two default users, **admin** and **monitor**, who cannot be deleted.
- You can, however, assign a new password for either one, and apply it to any appliances you wish.

Command Line Interface privileges

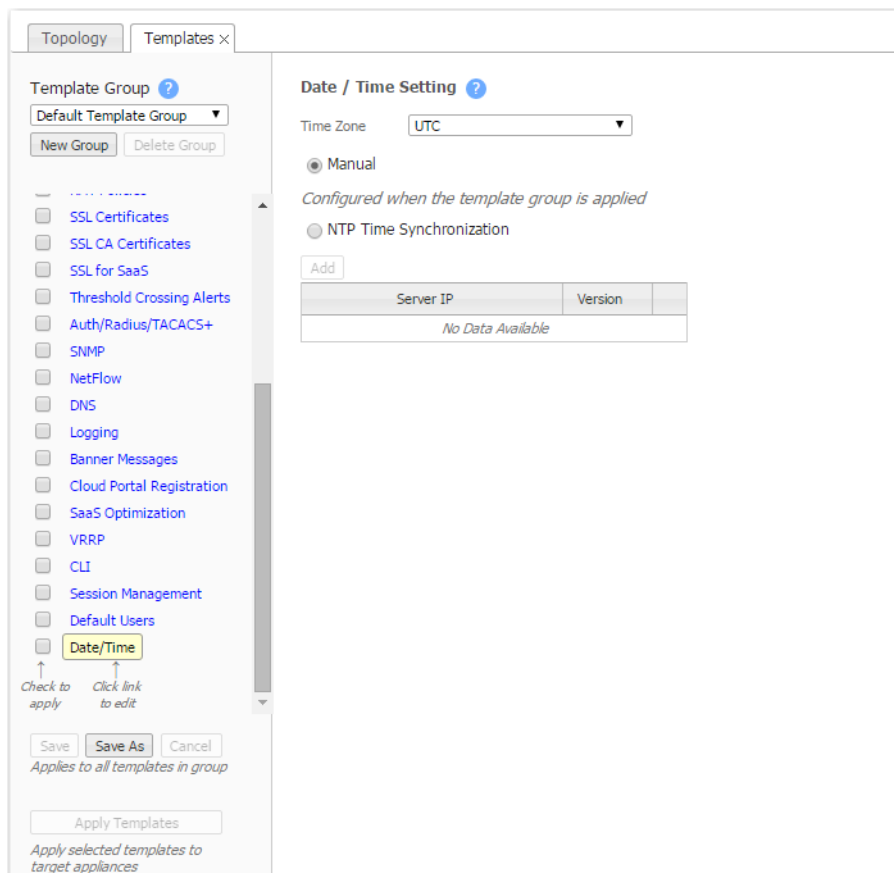
- The Command Line Interface (CLI) for Silver Peak physical (NX) appliances has three command modes. In order of increasing permissions, they are User EXEC Mode, Privileged EXEC Mode, and Global Configuration Mode.
- When you first log into a Silver Peak appliance via a console port, you are in User EXEC Mode. This provides access to commands for many non-configuration tasks, such as checking the appliance status.
- To access the next level, Privileged EXEC Mode, you would enter the *enable* command. With this template, you can choose to associate and enforce a password with the *enable* command.

Guidelines for Creating Passwords

- Passwords should be a minimum of 8 characters.
- There should be at least one lower case letter and one upper case letter.
- There should be at least one digit.
- There should be at least one special character.
- Consecutive letters in the password should not form words found in the dictionary.

Date/Time Template

Configure an appliance's **date and time** manually, or configure it to use an NTP (Network Time Protocol) server.



- From the **Time Zone** list, select the appliance's geographical location.
- Selecting **Manual** will match the appliance time to your web client system time when the template is applied. This is done to eliminate the delay between configuring time manually and applying the template.
- To use an NTP server, select **NTP Time Synchronization**.
 - Click **Add**.
 - Enter the IP address of the server, and select the version of NTP protocol to use.

When you list more than one NTP server, the Appliance Manager selects the servers in the order listed, always defaulting to the available server uppermost on the list.

Data Collection

- Silver Peak's GMS (Global Management System) collects and puts all stats in its own database in Coordinated Universal Time (UTC).
- When a user views stats, the appliance (or GMS server) returning the stats always presents the information relative to its own time zone.

System, Network, and Policy Configuration Tabs

This chapter describes the tabs for configuring network and appliance parameters.

In This Chapter

- **Deployment Tab** See page 96.
- **Interfaces Tab** See page 98.
- **Bridge Interfaces Tab** See page 99.
- **Tunnels Tab** See page 100.
- **Tunnel Groups Tab** See page 104.
- **Shaper Tab** See page 106.
- **Subnets Tab** See page 108.
- **SSL Certificates Tab** See page 110.
- **SSL CA Certificates Tab** See page 112.
- **SSL for SaaS Tab** See page 113.
- **VRRP Tab** See page 115.
- **WCCP Tab** See page 117.
- **Route Policies Tab** See page 120.
- **QoS Policies Tab** See page 122.
- **Optimization Policies Tab** See page 124.
- **Access Lists Tab** See page 126.
- **User Defined Applications Tab** See page 127.
- **Application Groups Tab** See page 128.
- **NAT Policies Tab** See page 129.
- **SaaS Optimization Tab** See page 132.
- **Threshold Crossing Alerts Tab** See page 134.

Deployment Tab

Configuration > Deployment

This report summarizes the appliance **Deployment** settings.

Summary view

The screenshot shows the 'Summary' view of the Deployment tab. The 'Summary' button is circled in red. A modal window titled 'Deployment Details' is open for the 'laine2-vxb' appliance. The modal contains an 'Export' button and a table with 2 rows of interface data.

Interface	Label	IP/Mask	WAN/LAN ...	Next Hop	Public IP	Shaping (Kbps)				
						Inbound	Outbound	NAT	Harden	DHCP
wan0		10.3.184.20/24	WAN	10.3.184.1		0	0	No	No	No
lan0			LAN			0	0	No	No	No

Details view

The screenshot shows the 'Details' view of the Deployment tab. The 'Details' button is circled in red. The main table displays configuration details for four appliances.

Edit	Appliance N...	Interface	Label	IP/Mask	WAN/LAN Si...	Next Hop	Public IP	Shaping (Kbps)				
								Inbound	Outbound	NAT	Harden	DHCP
	laine2-vxa	lan0			LAN			0	0	No	No	No
	laine2-vxa	wan0		10.3.183.20/24	WAN	10.3.183.1		0	0	No	No	No
	laine2-vxb	lan0			LAN			0	0	No	No	No
	laine2-vxb	wan0		10.3.184.20/24	WAN	10.3.184.1		0	0	No	No	No

- One of four deployment **Modes** displays:
 - **Router**: Single or dual WAN interfaces share LAN and WAN data traffic.
 - **InLine Router**: Uses separate LAN and WAN interfaces to route data traffic.
 - **Bridge**: Uses a virtual interface, *bvi*, created by binding the WAN and LAN interfaces
 - **Server**: Both management and data traffic use the **mgmt0** interface.

- **LAN labels** identify the data, such as *data*, *VoIP*, or *replication*.
- **WAN labels** identify the service, such as *MPLS* or *Internet*.

- The Shaper shapes traffic by allocating bandwidth as a *percentage* of the system bandwidth. This table displays the actual inbound or outbound **Shaping** in *kbps*.

- If **NAT** (Network Address Translation) is configured on the interface, it displays **Yes**. If not, then **No**.

- A **hardened** WAN-side interface provides additional security in Router mode and in Bridge modes. 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.
 - **Harden**: Displays as **Yes**, **No**, **No data** (not configured), or **Invalid data** (error condition).

- **DHCP**: Displays as **Yes**, **No**, **No data** (not configured), or **Invalid data** (error condition).

Interfaces Tab

Configuration > Interfaces

The **Interfaces** tab lists the appliance interfaces.

Edit	Appliance Na...	Name	Status	IP Address/Mask	Public IP	DHCP	Admin	Speed	Duplex	MTU	MAC Address
	laine-vxa	bvi0	Up	10.1.153.20/24		<input type="checkbox"/>	Up	N/A	N/A	1500	00:0C:29:1...
	laine-vxa	lan0	Up			<input type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:1...
	laine-vxa	wan1	Down			<input type="checkbox"/>	Down	1000Mb/s (a...	full (auto)	1500	Unassigned
	laine-vxa	wan0	Up			<input type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:1...
	laine-vxa	lan1	Down			<input type="checkbox"/>	Down	1000Mb/s (a...	full (auto)	1500	Unassigned
	laine-vxa	mgmt0	Up	10.0.238.71/26		<input checked="" type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:1...
	laine-vxa	mgmt1	Down	169.254.0.1/16		<input type="checkbox"/>	Down	(auto)	(auto)	1500	00:0C:29:1...
	laine-vxb	mgmt0	Up	10.0.238.69/26		<input checked="" type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:E...
	laine-vxb	wan1	Down			<input type="checkbox"/>	Down	1000Mb/s (a...	full (auto)	1500	Unassigned
	laine-vxb	lan1	Down			<input type="checkbox"/>	Down	1000Mb/s (a...	full (auto)	1500	Unassigned
	laine-vxb	bvi0	Up	10.1.154.20/24		<input type="checkbox"/>	Up	N/A	N/A	1500	00:0C:29:E...
	laine-vxb	mgmt1	Down	169.254.0.1/16		<input type="checkbox"/>	Down	(auto)	(auto)	1500	00:0C:29:E...
	laine-vxb	wan0	Up			<input type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:E...
	laine-vxb	lan0	Up			<input type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:E...
	laine2-vxa	mgmt0	Up	10.0.238.20/26	128.242.109.226	<input checked="" type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:F...
	laine2-vxa	bvi0	Up	10.3.183.20/24		<input type="checkbox"/>	Up	N/A	N/A	1500	00:0C:29:F...
	laine2-vxa	wan1	Down			<input type="checkbox"/>	Down	1000Mb/s (a...	full (auto)	1500	Unassigned
	laine2-vxa	wan0	Up			<input type="checkbox"/>	Up	10000Mb/s (...)	full (auto)	1500	00:0C:29:F...
	laine2-vxa	mgmt1	Down	169.254.0.1/16		<input type="checkbox"/>	Down	1000Mb/s (a...	full (auto)	1500	Unassigned

- As a best practice, assign static IP addresses to management interfaces to preserve their reachability.
- **Speed/Duplex** should never display as *half duplex* after auto-negotiation. If it does, the appliance will experience performance issues and dropped connections. To resolve, check the cabling on the appliance and the ports on the adjacent switch/router.
- To directly change interface parameters for a particular appliance, click **Edit**. It takes you to the Appliance Manager's **Configuration > Interfaces** page.
- To change the IP address for a **lan** or **wan** interface, either use the Appliance Manager's **Configuration > Deployment** page or the CLI (Command Line Interface).
- To change the IP address for **mgmt0**, either use the Appliance Manager's **Administration > Management IP/Hostname** page or the CLI.

Terminology

- **blan**: Bonded lan interfaces (as in **lan0** + **lan1**).
- **bvi0**: Bridge Virtual Interface. When the appliance is deployed in-line (Bridge mode), it's the routed interface that represents the bridging of **wan0** and **lan0**.
- **bwlan**: Bonded **wan** interfaces (as in **wan0** + **wan1**).
- **tlan**: 10-Gbps fiber **lan** interface.
- **twan**: 10-Gbps fiber **wan** interface.

Bridge Interfaces Tab

Configuration > Bridge Interfaces

For appliances in Bridge mode, this table lists how the data traffic spans the LAN and WAN interfaces.

Bridge Interfaces

Appliance Name	Interface	state	linkState	Pass-through Tx Interface
laine-vxa	lan0	active	up	wan0
laine-vxa	lan1	init	down (admin down)	wan1
laine-vxa	wan0	active	up	lan0
laine-vxa	wan1	init	down (admin down)	lan1
laine-vxb	lan0	active	up	wan0
laine-vxb	lan1	init	down (admin down)	wan1
laine-vxb	wan0	active	up	lan0
laine-vxb	wan1	init	down (admin down)	lan1
laine2-vxa	lan0	active	up	wan0
laine2-vxa	lan1	init	down (admin down)	wan1
laine2-vxa	wan0	active	up	lan0
laine2-vxa	wan1	init	down (admin down)	lan1
laine2-vxb	lan0	active	up	wan0
laine2-vxb	lan1	init	down (admin down)	wan1
laine2-vxb	wan0	active	up	lan0
laine2-vxb	wan1	init	down (admin down)	lan1

For pass-through traffic, ingress is at the **Interface** and egress is at the **Pass-through Tx [transmit] Interface**.

Tunnels Tab

Configuration > Tunnels

Use this page to **view**, **edit**, and **delete** tunnels.

- To manage tunnels and assign their properties, use the **Tunnels** section of the **Templates** tab.
- To create tunnels, use **Business Intent Overlays** or **Tunnel Groups**.
- Overlay tunnels consist of bonded underlay tunnels.

The screenshot displays the 'Tunnels' tab in the Silver Peak Unity Orchestrator. The main table lists 4 rows of tunnels. Below the table, two pop-up windows are visible: 'Appliance Site Info' and 'Tunnel Advanced Options'.

Edit	Appliance	Name	Status	MTU	Local IP	Remote IP	Max BW (Kb...)	Mode	Uptime	Advanced Options
	laine-vxa	auto_tun_10.1...	up - active	Auto	10.1.153.20	10.1.154.20	Auto	udp	19d 5h 50m 23s	
	laine-vxb	auto_tun_10.1...	up - active	Auto	10.1.154.20	10.1.153.20	Auto	udp	19d 5h 50m 8s	
	laine2-vxa	tun1	up - active	Auto	10.3.183.20	10.3.184.20	Auto	udp	4h 44m 7s	
	laine2-vxb	tun1	up - active	Auto	10.3.184.20	10.3.183.20	Auto	udp	4h 44m 7s	

Appliance	Network Role	Site	Priority
tallin	Mesh		0
laine2-vxb	Mesh		0
laine2-vxa	Mesh		0
laine-vxb	Mesh		0
laine-vxa	Mesh		0

General	
IPsec Anti-replay Window	1024
UDP Destination Port	4163
UDP Flows	256
Min BW (Kbps)	32
Packet	
Coalescing Enabled	true
Coalescing Wait (ms)	0
Reorder Wait (ms)	100
FEC	disable
FEC Ratio	1:10
Tunnel Health	
Retry Count	30
DSCP	
FastFail Thresholds	
Fastfail Enabled	disable
Latency (ms)	0
Loss (%)	0
Jitter (ms)	0

Troubleshooting

What to check if you're using Business Intent Overlays, and you don't see the tunnels you expect to see in this table:

- 1 *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.

- 2 *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**.

- 3 *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.

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

- 5 *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 **Roles/Sites** at the top.

Advanced Tunnel Options

As needed, use the options **Tunnel** template to configure and push these options.

Tunnel Advanced Options	
General	
IPsec Anti-replay Window	1024
UDP Destination Port	4163
UDP Flows	256
Min BW (Kbps)	40
Packet	
Coalescing Enabled	true
Coalescing Wait (ms)	0
Reorder Wait (ms)	100
FEC	disable
FEC Ratio	1:10
Tunnel Health	
Retry Count	30
DSCP	
FastFail Thresholds	
Fastfail Enabled	disable
Latency (ms)	0
Loss (%)	0
Jitter (ms)	0
Fastfail Wait-time Base Offset (ms)	150
Fastfail RTT Multiplication Factor	2

Definitions (alphabetically)

- **Admin Status** indicates whether the tunnel has been set to admin **Up** or **Down**.
- **Coalescing Enabled** allows the appliance to coalesce smaller packets into larger packets.
- **Coalescing Wait** is the number of milliseconds that the appliance should hold packets while attempting to coalesce smaller packets into larger ones.
- **DSCP** determines which DSCP marking the keep-alive messages should use.

- **Fastfail Thresholds** – When multiple tunnels are carrying data between two appliances, this feature determines how quickly to disqualify a tunnel from carrying data.

The Fastfail connectivity detection algorithm for the wait time from receipt of last packet before declaring a **brownout** is:

$$T_{wait} = Base + N * RTT_{avg}$$

where **Base** is a value in milliseconds, and **N** is the multiplier of the average Round Trip Time over the past minute.

For example, if:

$$\begin{aligned} Base &= 200mS \\ N &= 2 \end{aligned}$$

Then,

$$RTT_{avg} = 50mS$$

The appliance declares a tunnel to be in **brownout** if it doesn't see a reply packet from the remote end within 300mS of receiving the most recent packet.

In the Tunnel Advanced Options, **Base** is expressed as **Fastfail Wait-time Base Offset (ms)**, and **N** is expressed as **Fastfail RTT Multiplication Factor**.

- **Fastfail Enabled** – This option is triggered when a tunnel's keepalive signal doesn't receive a reply. The options are **disable**, **enable**, and **continuous**. If the disqualified tunnel subsequently receives a keepalive reply, its recovery is instantaneous.
 - If set to **disable**, keepalives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost.
 - If set to **enable**, keepalives are sent every second, and a missed reply increases the rate at which keepalives are sent from 1 per second to 10 per second. Failover occurs after 1 second.
 - When set to **continuous**, keepalives are continuously sent at 10 per second. Therefore, failover occurs after one tenth of a second.
- Thresholds for **Latency**, **Loss**, or **Jitter** are checked once every second.
 - Receiving 3 successive measurements in a row that exceed the threshold puts the tunnel into a brownout situation and flows will attempt to fail over to another tunnel within the next 100mS.
 - Receiving 3 successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.
- **FEC** (Forward Error Correction) can be set to **enable**, **disable**, and **auto**.
- **FEC Ratio** is an option when FEC is set to **auto**, that specifies the maximum ratio. The options are 1:2, 1:5, 1:10, or 1:20.
- **IPSec Anti-replay window** provides protection against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet. The decryptor keeps track of which packets it has seen on the basis of these numbers. The default window size is 64 packets.
- **Local IP** is the IP address for the local appliance.
- **Max BW (Kbps)** is the maximum bandwidth for this tunnel, in kilobits per second. This must be **less than or equal to** the upstream bandwidth of your WAN connection.
- **Min BW (Kbps)** is the minimum bandwidth for this tunnel, in kilobits per second.
- **Mode** indicates whether the tunnel protocol is **udp**, **gre**, or **ipsec**.

- **MTU (bytes)** (Maximum Transmission Unit) is the largest possible unit of data that can be sent on a given physical medium. For example, the MTU of Ethernet is 1500 bytes. Silver Peak provides support for MTUs up to 9000 bytes. **Auto** allows the tunnel MTU to be discovered automatically, and it overrides the MTU setting.
- **Remote IP** is the IP address for the remote appliance.
- **Reorder Wait (ms)** - Maximum time the appliance holds an out-of-order packet when attempting to reorder. The 100ms default value should be adequate for most situations. FEC may introduce out-of-order packets if the reorder wait time is not set high enough.
- **Retry Count** is the number of failed keep-alive messages that are allowed before the appliance brings the tunnel down.
- **Status** indications are as follows:
 - **Down** = The tunnel is down. This can be because the tunnel administrative setting is down, or the tunnel can't communicate with the appliance at the other end. Possible causes are:
 - Lack of end-to-end connectivity / routability (test with *iperf*)
 - Intermediate firewall is dropping the packets (open the firewall)
 - Intermediate QoS policy (**be** packets are being starved. Change control packet DSCP marking)
 - Mismatched tunnel mode (**udp / gre / ipsec**)
 - IPsec is misconfigured: (1) enabled on one side (see *show int tunnel configured*), or (2) mismatched pre-shared key
 - **Down - In progress** = The tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel.
 - **Down - Misconfigured** = The two appliances are configured with the same System ID. (see *show system*)
 - **Up - Active** = The tunnel is up and active. Traffic destined for this tunnel will be forwarded to the remote appliance.
 - **Up - Active - Idle** = The tunnel is up and active but hasn't had recent activity in the past five minutes, and has slowed the rate of issuing keep-alive packets.
 - **Up - Reduced Functionality** = The tunnel is up and active, but the two endpoint appliances are running mismatched software releases that give no performance benefit.
 - **UNKNOWN** = The tunnel status is unknown. This can be because the appliance is unable to retrieve the current tunnel status. Try again later.
- **UDP destination port** is used in UDP mode. Accept the default value unless the port is blocked by a firewall.
- **UDP flows** is the number of flows over which to distribute tunnel data. Accept the default.
- **Uptime** is how long since the tunnel came up.

Tunnel Groups Tab

Configuration > Tunnel Groups

A **Tunnel Group** consists of a set of appliances, paired with a configuration that defines how to build tunnels among them.

Use this page to create Tunnel Groups.

The screenshot shows the 'Tunnel Groups' configuration page. At the top, there are tabs for 'Topology' and 'Tunnel Groups x'. Below the tabs, the page title is 'Tunnel Groups' with a help icon and a link to 'Manage Appliances and Tunnel Groups'. There is a 'Settings' button. A 'Tunnel Group Name' field with a dropdown arrow and '+Add Rename Delete' links is present. The 'Topology' section has two radio buttons: 'Mesh' (selected) and 'Hub & Spoke'. A 'Connect to Hubs +Add' link is next to a large empty box. The 'Interfaces' section has a sub-section 'Interfaces To Connect' with two radio buttons: 'Connect All Available Interfaces' (selected) and 'Only Connect These Labels'. Below this are three checkboxes: 'MPLS', 'Internet', and 'LTE'. The 'Cross Connect' section has a checkbox for 'Cross Connect Interfaces' with a note: '(For example, allow MPLS <--> Internet)'. At the bottom, there are 'Save' and 'Cancel' buttons.

- Orchestrator automatically builds these tunnels in the background.
- Tunnel groups are self-healing. If a change is made to an IP address (as with DHCP) or to a Label, those changes propagate appropriately through the tunnel groups.
- To assign tunnel properties, use the **Tunnels** section of the **Configuration > Templates** tab.
- To **add** and **remove** appliances from Tunnel Groups, click **Manage Appliances and Tunnel Groups**.
- To **view** a list of tunnels, refer to the **Configuration > Tunnels** tab.
- To pause Orchestrator's tunnel management while you troubleshoot, click **Settings** and deselect **Enable**.

Topology

- You can choose either a **Mesh** or a **Hub & Spoke** topology.
- If choosing **Hub & Spoke**, choose the hubs you need from the **Select Hubs** area. If one you need isn't displayed, click **+Add**, as needed.

Interfaces

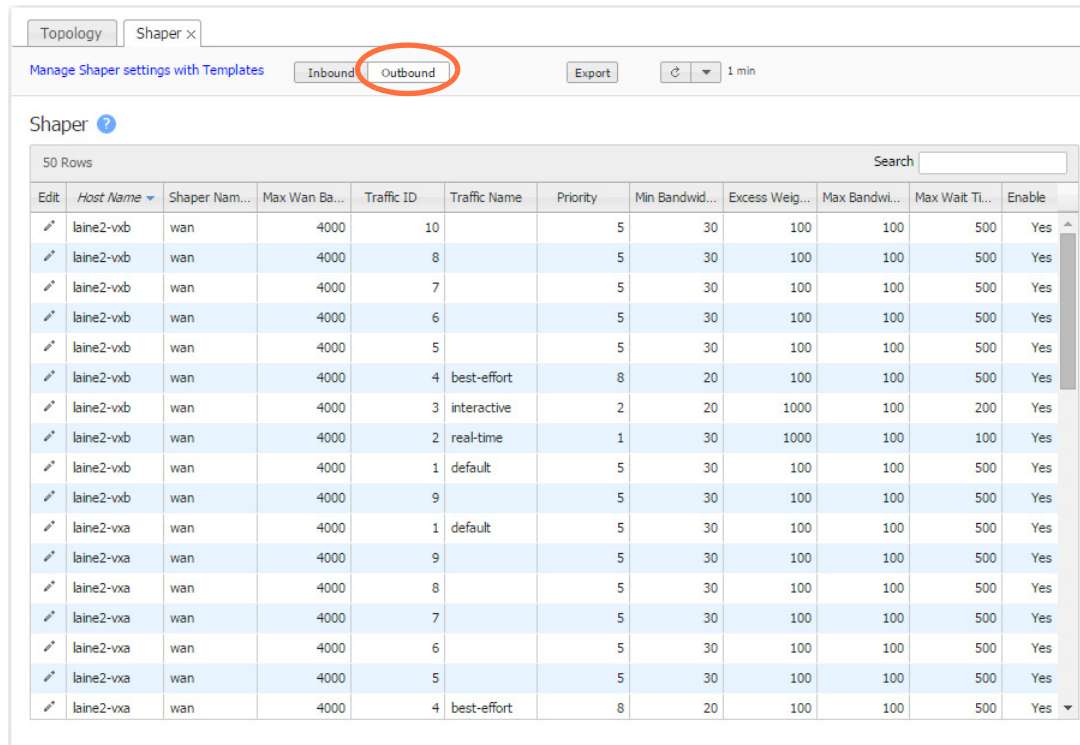
- **Connect all Available Interfaces** refers to WAN ports only. If an appliance is in Server mode, its WAN port is the **mgmt0** interface.
- **Only Connect These Labels** is an option when the appliance is at Release 8.0 or later, and you have used the Orchestrator to assign labels to interfaces. Generally, WAN interfaces are named according to the service or service provider.

Shaper Tab

Configuration > Shaper

This report provides a view of the Shaper settings.

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

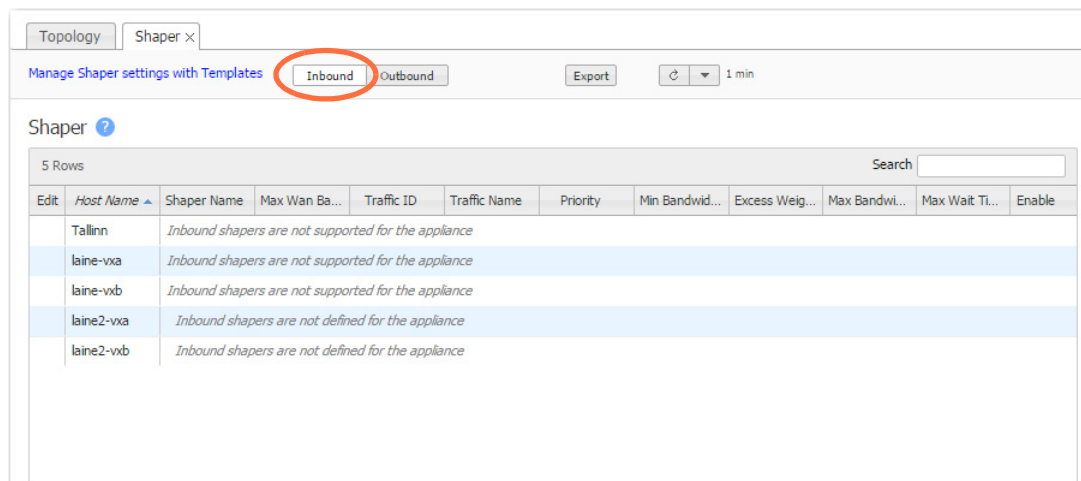


Manage Shaper settings with Templates Inbound Outbound Export 1 min

Shaper ?

50 Rows Search

Edit	Host Name	Shaper Name	Max Wan Ba...	Traffic ID	Traffic Name	Priority	Min Bandwid...	Excess Weig...	Max Bandwi...	Max Wait Tl...	Enable
	laine2-vxb	wan	4000	10		5	30	100	100	500	Yes
	laine2-vxb	wan	4000	8		5	30	100	100	500	Yes
	laine2-vxb	wan	4000	7		5	30	100	100	500	Yes
	laine2-vxb	wan	4000	6		5	30	100	100	500	Yes
	laine2-vxb	wan	4000	5		5	30	100	100	500	Yes
	laine2-vxb	wan	4000	4	best-effort	8	20	100	100	500	Yes
	laine2-vxb	wan	4000	3	interactive	2	20	1000	100	200	Yes
	laine2-vxb	wan	4000	2	real-time	1	30	1000	100	100	Yes
	laine2-vxb	wan	4000	1	default	5	30	100	100	500	Yes
	laine2-vxb	wan	4000	9		5	30	100	100	500	Yes
	laine2-vxa	wan	4000	1	default	5	30	100	100	500	Yes
	laine2-vxa	wan	4000	9		5	30	100	100	500	Yes
	laine2-vxa	wan	4000	8		5	30	100	100	500	Yes
	laine2-vxa	wan	4000	7		5	30	100	100	500	Yes
	laine2-vxa	wan	4000	6		5	30	100	100	500	Yes
	laine2-vxa	wan	4000	5		5	30	100	100	500	Yes
	laine2-vxa	wan	4000	4	best-effort	8	20	100	100	500	Yes



Manage Shaper settings with Templates Inbound Outbound Export 1 min

Shaper ?

5 Rows Search

Edit	Host Name	Shaper Name	Max Wan Ba...	Traffic ID	Traffic Name	Priority	Min Bandwid...	Excess Weig...	Max Bandwi...	Max Wait Tl...	Enable
	Tallinn	<i>Inbound shapers are not supported for the appliance</i>									
	laine-vxa	<i>Inbound shapers are not supported for the appliance</i>									
	laine-vxb	<i>Inbound shapers are not supported for the appliance</i>									
	laine2-vxa	<i>Inbound shapers are not defined for the appliance</i>									
	laine2-vxb	<i>Inbound shapers are not defined for the appliance</i>									

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

Definitions

- **Traffic Name:** Name assigned to a traffic class, either prescriptively or by the user.
- **Priority:** Determines the order in which to allocate each class's minimum bandwidth - 1 is first, 10 is last.
- **Min Bandwidth:** Refers to the percentage of bandwidth guaranteed to each traffic class, allocated by priority. However, if the sum of the percentages is greater than 100%, then lower-priority traffic classes might not receive their guaranteed bandwidth if it's all consumed by higher-priority traffic.

If you set **Min Bandwidth** to a value greater than **Max Bandwidth**, then **Max** overrides **Min**.

- **Excess Weighting:** If there is bandwidth left over after satisfying the minimum bandwidth percentages, then the excess is distributed among the traffic classes, in proportion to the weightings specified in the **Excess Weighting** column. Values range from 1 to 10,000.
- **Max Bandwidth:** You can limit the maximum bandwidth that a traffic class uses by specifying a percentage in the **Max Bandwidth** column. The bandwidth usage for the traffic class will never exceed this value.
- **Max Wait Time:** Any packets waiting longer than the specified **Max Wait Time** are dropped.

Subnets Tab

Configuration > Shaper

To add, edit, or delete a subnet, you must *select an individual appliance* from the navigation panel.

Edit	Appliance Name	Metric	Subnet/Mask	Is Local	Advertise to Peers	Exclude	Type	SaaS Application Name	Learned from Peer	Comment
✓	Tallinn	No Subnets defined for this appliance.								
✓	laine-vxa	No Subnets defined for this appliance.								
	laine-vxa	50	10.1.153.0/24	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Auto (added by system)			
	laine-vxa	50	10.1.154.0/24	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		laine-vxb	
✓	laine-vxb	No Subnets defined for this appliance.								
	laine-vxb	50	10.1.153.0/24	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		laine-vxa	
	laine-vxb	50	10.1.154.0/24	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Auto (added by system)			
✓	laine2-vxa	No Subnets defined for this appliance.								
	laine2-vxa	50	10.3.183.0/24	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Auto (added by system)			
	laine2-vxa	50	10.3.184.0/24	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		laine2-vxb	
✓	laine2-vxb	No Subnets defined for this appliance.								
	laine2-vxb	50	10.3.183.0/24	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		laine2-vxa	
	laine2-vxb	50	10.3.184.0/24	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Auto (added by system)			

All = (manually) Configured + Learned

No subnets were manually configured

What is subnet sharing?

Subnet sharing is one of the three strategies that Silver Peak uses to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. Auto-optimization strategies reduce the need to create explicit route map entries to optimize traffic. The other two strategies are **TCP-based** auto-opt and **IP-based** auto-opt.



Note Enabled by default, the global settings for all three reside on the **Templates** tab, under **System**.

How is subnet sharing implemented?

Each appliance builds a subnet table from entries added automatically by the system or manually by a user. When two appliances are connected by a tunnel, they exchange this information ("learn" it) and use it to route traffic to each other.

When would you need to use a Route Policy template?

Subnet sharing takes care of optimizing IP traffic.

Use and apply a Route Policy template for flows that are to be:

- sent pass-through (shaped or unshaped)
- dropped

- configured for a specific high-availability deployment
- routed based on application, ports, VLAN, DSCP, or ACL (Access Control List)

Subnet table columns

- **Subnet/Mask:** Actual subnet to be shared or learned
- **Metric:** Metric of the subnet. Value must be between 0 and 100. When a peer has more than one tunnel with a matching subnet (for example, in a high availability deployment), it chooses the tunnel with the greater numerical value.
- **Is Local:** Specifies if the subnet is local to this site.

The appliance sets this parameter for **automatically** added subnets because those subnets are directly attached to an appliance interface, and therefore are most likely local to the appliance.

Also, you can select the parameter when **manually** adding a subnet:

- Select this option for a manually added subnet if all the IP addresses in the subnet are known to be local.
- Deselect this option if the subnet is so large (for example, 0.0.0.0/0) that it may include IP addresses that are not local to this appliance. If a subnet is too wide, and it's marked **local**, then the stats will count any pass-through packets with an IP address within that range as WAN-to-LAN.
- **Exclude:** Use this option to prevent optimization of more specific subnets from a wider advertised subnet range.
- **Advertise to Peers:** Selected by default, it shares the subnet information with peers. Peers then learn it.

To add a subnet to the table without divulging it to peers, yet, deselect this option.

- **Type of subnet:**
 - **Auto (added by system)** = automatically added subnets of interfaces on this appliance
 - **Auto (added by saas optimization)** = automatically added subnets from SaaS services
 - **Added by user** = manually added/configured subnets for this appliance
 - **Learned from peer** = subnets added as a result of exchanging information with peer appliances
- **SaaS Application Name:** Specifies the SaaS application. For example, **Outlook**, **Office 365**, or **Salesforce**.
- **Learned from Peer:** Which peer appliance advertised (and shared) this subnet information

SSL Certificates Tab

Configuration > SSL Certificates

Silver Peak provides deduplication for Secure Socket Layer (SSL) encrypted WAN traffic by supporting the use of SSL certificates and keys.

The screenshot shows the 'SSL Certificates' management interface. At the top, there are tabs for 'Topology' and 'SSL Certificates x'. Below the tabs, there is a search bar and a 'Show' dropdown set to '25'. A table lists certificates with columns for 'Appliance Name', 'Issuer', 'Issued To', 'Certificate', and 'Expiration Date'. One entry is visible: 'DM-VX-B' issued to 'dm' by 'dm', with an expiration date of 'Jan 1 23:19:39 2015 GMT'. A 'View c#' link is next to the certificate name. An arrow points from this link to a 'View Certificate Content' dialog box. The dialog box displays the following certificate details:

```

Certificate:
Data:
Version: 3 (0x2)
Serial Number:
cd:a8:77:1b:f5:8d:14:ac
Signature Algorithm: sha1withRSAEncryption
Issuer: C=us, ST=ca, L=san jose, O=silverpeak, OU=eng, CN=dm/emailAddress=dm@dm.com
Validity
Not Before: Dec 2 23:19:39 2014 GMT
Not After : Jan 1 23:19:39 2015 GMT
Subject: C=us, ST=ca, L=san jose, O=silverpeak, OU=eng, CN=dm/emailAddress=dm@dm.com
Subject Public Key Info:
Public Key Algorithm: rsaEncryption
RSA Public Key: (2048 bit)
Modulus (2048 bit):
00:d7:ea:5b:15:6a:c1:43:67:8c:29:c8:01:2c:b8:
e1:eb:a6:8d:f2:d9:78:18:fd:bb:46:9b:38:b3:fc:
d0:2c:dd:85:83:f7:a6:02:6f:55:23:1a:db:a1:36:
98:4c:6d:18:51:22:f2:05:7d:29:94:12:dc:54:b2:
80:f5:61:7b:60:8c:57:58:bc:da:0c:d0:18:09:d3:
c8:c2:ca:be:64:b7:cf:a6:15:73:27:b5:91:29:8c:
8e:ce:2e:8d:42:fe:ff:05:d7:69:cf:73:ea:f7:d6:
23:fb:98:4f:8f:70:8e:51:98:78:4f:ca:36:a5:eb:
4e:01:6a:6d:97:bf:ad:a6:52:76:95:b8:9f:2e:71:
75:e7:b0:69:40:0b:d3:c8:bc:24:62:98:54:7d:d8:
2d:44:94:00:92:6a:e8:51:4b:6c:58:b1:c5:7b:05:
d0:88:89:f1:c4:fa:da:43:07:2c:bc:ee:19:2d:8b:
b7:88:4c:ad:62:35:d5:9a:39:eb:1f:9e:3c:85:78:
58:a9:e9:e5:7e:fd:30:33:74:39:1e:00:cb:19:45:
12:55:68:cd:fa:86:8a:1b:07:81:20:c2:6e:59:f9:
d7:22:53:f9:4d:7c:49:c9:e0:81:9e:fd:f3:83:c5:
23:99:cb:fd:2b:ad:8d:d2:26:da:13:74:06:31:e8:
27:0f
Exponent: 65537 (0x10001)
X509v3 extensions:
X509v3 Subject Key Identifier:
6E:BC:F0:A1:FE:AF
  
```

This report summarizes the SSL certificates installed on appliances **for decrypting non-SaaS traffic**.

- Silver Peak decrypts SSL data using the configured certificates and keys, optimizes the data, and transmits data over an IPsec tunnel. The peer Silver Peak appliance uses configured SSL certificates to re-encrypt data before transmitting.
- Peers that exchange and optimize SSL traffic must use the same certificate and key.
- Silver Peak supports:
 - X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX, generally for Microsoft servers), and RSA key 1024-bit and 2048-bit certificate formats.
 - SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.
- Silver Peak appliances support:
 - **Protocol versions:** SSLv3, SSLv3.3, TLS1.0, TLS1.1, TLS1.2
 - **Key exchanges:** RSA, DHE, ECDHE
 - **Authentication:** RSA
 - **Cipher algorithms:** RC4, 3DES, AES128, AES256, AES128-GCM, AES256-GCM
 - **Message Digests:** MD5, SHA, SHA256, SHA284

- For the SSL certificates to function, the following must also be true:
 - The tunnels are in **IPSec** mode for both directions of traffic.
 - In the *Optimization Policy*, **TCP acceleration** and **SSL acceleration** are enabled.



Tip For a historical matrix of SSL/TLS versions and ciphers for VXOA releases, click [here](#).

SSL CA Certificates Tab

Configuration > SSL CA Certificates

This tab lists any installed Certificate Authorities (CA) that the browser uses to validate up the chain to the root CA.

The screenshot displays the 'SSL CA Certificates' tab. At the top, there is a breadcrumb 'Configuration > SSL CA Certificates' and a search bar. Below the search bar, there is a 'Manage SSL CA Certificates with Templates' button and an 'Export' button. The main content area shows a table with the following data:

Edit	Appliance Name	Issuer	Issued To	Certificate	Expiration Date
	chateau	Silver Peak SSL Proxy	Silver Peak SSL Proxy	View	Dec 31 23:59:59 2034 GMT

At the bottom of the table, it says 'Showing 1 to 1 of 1 entries' and there are navigation buttons for 'First', 'Previous', '1', 'Next', and 'Last'.

If the enterprise certificate that you used for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, then you must add those CA certificates. If the browser can't validate up the chain to the root CA, it will warn you that it can't trust the certificate.

- Silver Peak supports:
 - X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
 - SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.



Tip For a historical matrix of SSL/TLS versions and ciphers for VXOA releases, click [here](#).

SSL for SaaS Tab

Configuration > SSL for SaaS

This report lists the appliances' signed substitute certificates.

The screenshot shows the 'SSL for SaaS' configuration page. At the top, there are tabs for 'Topology' and 'SSL for SaaS'. Below the tabs, there is a 'Manage SSL for SaaS with Templates' button, an 'Export' button, and a refresh icon with a '3 mins' timer. The main content area is titled 'SSL for SaaS' and includes a search bar and a 'Show 25' dropdown. A table lists two appliances: 'castle' and 'chateau'. The 'chateau' row is highlighted, and a 'View' link is visible. A modal window titled 'View Certificate Content' is open, displaying the following certificate details:

```

Certificate:
Data:
  Version: 3 (0x2)
  Serial Number: 14340313823633164355 (0xc7030771c1474843)
  Signature Algorithm: sha256WithRSAEncryption
  Issuer: CN=Silver Peak SSL Proxy, C=US, ST=California, L=Santa Clara, O=Silver Peak Systems
  Validity
    Not Before: Jan 1 00:00:00 2015 GMT
    Not After : Dec 31 23:59:59 2034 GMT
  Subject: CN=Silver Peak SSL Proxy, C=US, ST=California, L=Santa Clara, O=Silver Peak Systems
  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    Public-Key: (2048 bit)
  Modulus:
    00:b0:cf:eb:0d:d8:1b:51:5a:67:82:48:8e:d8:e4:
    2e:f2:b0:89:e3:17:4e:1f:4d:cb:c2:02:36:2f:dd:
    6f:a1:c1:c9:2a:d3:10:c3:26:da:b3:5e:73:ba:5d:
    ce:ce:d1:d7:13:9d:82:80:9a:58:cb:db:79:a6:0f:
    f9:0e:2e:43:e1:fc:60:65:4b:4b:06:fa:83:97:b2:
    fe:3f:35:64:23:38:18:72:db:73:8f:7b:98:58:9e:
    cc:e8:20:5c:2b:de:5f:e1:ae:6a:20:ae:a2:a5:4b:
    8f:81:47:15:0f:43:8f:1f:4d:7c:db:94:8b:3d:d2:
    93:51:9f:8a:43:45:9a:87:20:23:b8:f9:40:5f:bf:
    bf:52:a2:98:78:8b:4e:60:ee:fb:06:52:53:38:56:
    75:2e:3b:72:22:04:fb:46:1b:06:83:ce:ad:14:
    af:29:11:19:92:a8:12:00:a9:39:64:bd:1e:ea:2b:
    e1:ae:41:72:01:03:42:7c:21:17:5c:b4:eb:20:64:
  
```

To fully compress SSL traffic for a SaaS service, the appliance must decrypt it and then re-encrypt it.

To do so, the appliance generates a substitute certificate that must then be signed by a Certificate Authority (CA). There are two possible signers:

- For a *Built-In CA Certificate*, the signing authority is Silver Peak.
 - The appliance generates it locally, and each certificate is unique. This is an ideal option for Proof of Concept (POC) and when compliance is not a big concern.
 - To avoid browser warnings, follow up by importing the certificate into the browser from the client-side appliance.
- For a *Custom CA Certificate*, the signing authority is the Enterprise CA.
 - If you already have a subordinate CA certificate (for example, an SSL proxy), you can upload it to the Orchestrator and push it out to the appliances. If you need a copy of it later, just download it from here.
 - If this substitute certificate is subordinate to a root CA certificate, then also install the higher-level **SSL CA certificates** (into the **SSL CA Certificates** template) so that the browser can validate up the chain to the root CA.
 - If you **don't** already have a subordinate CA certificate, you can access any appliance's **Configuration > SaaS Optimization** page and generate a Certificate Signing Request (CSR).

- Silver Peak appliances support:
 - **Protocol versions:** SSLv3, SSLv3.3, TLS1.0, TLS1.1, TLS1.2
 - **Key exchanges:** RSA, DHE, ECDHE
 - **Authentication:** RSA
 - **Cipher algorithms:** RC4, 3DES, AES128, AES256, AES128-GCM, AES256-GCM
 - **Message Digests:** MD5, SHA, SHA256, SHA284



Tip For a historical matrix of SSL/TLS versions and ciphers for VXOA releases, click [here](#).

VRRP Tab

Configuration > VRRP

This tab summarizes the configuration and state for appliances deployed with **Virtual Router Redundancy Protocol (VRRP)**.

Edit	Appliance Name	Group	Interface	State	Admin	Virtual IP	Advertisement	Priority	Preemption	Master IP	Virtual MAC Address	State Uptime	Master State Transitions	IP Address
	Chicago													
	Dallas													
	Denver-EC													
	Los-Angeles													
	Seattle-EC													

In an out-of-path deployment, one method for redirecting traffic to the Silver Peak appliance is to configure 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**.

DEFINITIONS (alphabetically)

- **Admin** = The options are up (enable) and **down** (disable).
- **Advertisement Timer** = default is **1 second**.
- **Group ID** is a value assigned to the two peers. Depending on the deployment, the group can consist of an appliance and a router (or L3 switch), or two appliances. The valid range is **1 - 255**.
- **Interface** refers to the interface that VRRP is using for peering.
- **IP Address Owner** = A Silver Peak appliance cannot use one of its own IP addresses as the VRRP IP, so this will always be **No**.
- **Master IP** = Current VRRP Master's Interface or local IP address.
- **Master State Transitions** = Number of times the VRRP instance went from Master to Backup and vice versa. A high number of transitions indicates a problematic VRRP configuration or environment. If this is the case, check the configuration of all local appliances and routers, and review the log files.
- **Preemption**. Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
- **Priority**. The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.
- **State Uptime** = Time elapsed since the VRRP instance entered the state it's in.

- **State** = There are three options for the VRRP instance:
 - **Backup** = Instance is in VRRP backup state.
 - **Init** = Instance is initializing, it's disabled, or the interface is down.
 - **Master** = Instance is the current VRRP master.
- **Virtual IP**. The IP address of the VRRP instance. VRRP instances may run between two or more appliances, or an appliance and a router.
- **Virtual MAC address** = MAC Address that the VRRP instance is using. On an NX Appliance, this is in 00-00-5E-00-01-**{VRID}** format. On virtual appliances, the VRRP instance uses the interface's assigned MAC Address (for example, the MAC address that the hypervisor assigned to **wan0**).

WCCP Tab

Configuration > VRRP

Use this page to **view**, **edit**, and **delete** WCCP Service Groups.

Edit	Appliance Name	Group ID	Oper Status	Admin	Router IP	Protocol	Interface	Compatibility	Forwarding Met...	Advanced Settings
✓	Tallinn		No data available							
✓	laine-vxa		Not applicable as this appliance is in Bridge mode.							
✓	laine-vxb		Not applicable as this appliance is in Bridge mode.							
✓	laine2-vxa		Not applicable as this appliance is in Bridge mode.							
✓	laine2-vxb		Not applicable as this appliance is in Bridge mode.							

Web Cache Communications Protocol (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 optimizes traffic flows that the Route Policy tunnelizes. The appliance forwards all other traffic as pass-through or pass-through-unshaped, as per the Route Policy.

- For the Service Groups to be active, you must select **Enable WCCP**. Otherwise, the service groups are configured, but not in service.
- 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.
- If the appliance uses *auto-optimization*, then WCCP redirection must also be applied on the uplinks of the router or L3 switch to the core/WAN.



Refer to the *Silver Peak Network Deployment Guide* for examples, best practices, and deployment tips.

Definitions (alphabetically)

- **Admin** values are up and down. The default is up.
- **Advanced Settings**. You can only configure these options directly on the appliance. For more information, and best practices, refer to the *Silver Peak Network Deployment Guide*.
- **Compatibility** Mode. Select the option appropriate for your router. If a WCCP group is peering with a router running **Nexus OS**, then the appliance must adjust its WCCP protocol packets to be compatible. By default, the appliance is **IOS-compatible**.
- **Forwarding Method**, also known as the *Redirect Method*. 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 Route Encapsulation (GRE)** and **L2 redirection**.
 - **either** allows the appliance and the router to negotiate the best option. You should always select **either**. During protocol negotiation, if the router offers both GRE and L2 as redirection methods, the appliance will automatically select L2.

- **GRE** (Layer 3 Generic Routing Encapsulation) allows packets to reach the appliance even if there are other routers in the path between the forwarding router and the appliance. At high traffic loads, this option may cause high CPU utilization on some Cisco platforms.
- **L2** (Layer-2) redirection takes advantage of internal switching hardware that either partially or fully implements the WCCP traffic interception and redirection functions at Layer 2. Layer-2 redirection requires that the appliance and router be on the same subnet. It is also recommended that the appliance is given a separate subnet to avoid pass-through traffic from being redirected back to the appliance and causing a redirection/Layer-3 loop.
- **Group ID** refers to the Service Group ID.
- **Interface**. The default value is **wan0**.
- **Oper Status**. Common states: **INIT**, **Active - Designated**, **Active**
 - **INIT**. Initializing or down
 - **ACTIVE**. This indicates that the protocol is established and the router has assigned hash/mask buckets to this appliance.
 - **BACKUP**. This indicates that the protocol is established but the router has not assigned any hash/mask buckets to this appliance. This may be caused by using a Weight of **0**.
 - **Designated**. This state (in addition to Active/Backup) indicates that the appliance is the designated web-cache for the group. The designator communicates with the router(s) to assign hash/mask assignments. When there is more than one appliance in a group, the appliance with the lowest IP becomes the designator for that group.
- **Protocol**. Although many more protocols are supported, generally **TCP** and **UDP** are the focus. For troubleshooting, you may consider adding a group for **ICMP** as well.
- **Router IP** is the IP address of the WCCP router. For Layer 2 redirection, use the physical IP address of the interface that is directly connected to the appliance. For Layer 3 redirection, consider using a loopback IP. It is not recommended to use VRRP or HSRP IPs as router IPs.

Service Group Advanced Settings

- **Assignment Detail**
 - This field can be used to customize hash or mask values. If you have only one appliance or if you are using route-map or subnet sharing to tunnelize, use the default **LAN-ingress** setting.
 - **WAN-ingress** and **LAN-ingress** are not applicable if there is only one active appliance.
 - **WAN-ingress** and **LAN-ingress** are also not applicable if you are using route-map or subnet sharing to tunnelize.
 - If there is more than one active appliance and you're using *TCP-IP auto-optimization*:
 - Use **LAN-ingress** for WCCP groups that are used to redirect outbound traffic.
 - Use **WAN-ingress** for WCCP groups that are used to redirect inbound traffic.

This ensures that a connection will go through the same appliance in both inbound and outbound directions and avoid asymmetry.
 - **custom** provides granular control of the distribution of flows. Contact Silver Peak Technical Support for assistance.
- **Assignment Method** determines how redirected packets are distributed between the devices in a Service Group, effectively providing load balancing among the devices. The options are:

- **either**, which lets the appliance and router negotiate the best method for assignment. This is preferred. If the router offers both *hash* and *mask* methods, then the appliance will select the *mask* assignment method.
 - **hash**, for hash table assignment
 - **mask**, for mask/value sets assignment
- **Force L2 Return** is generally not selected. Normally, all Layer-3 redirected traffic that isn't optimized (that is, it's pass-through) is returned back to the WCCP router as GRE (L3 return). Processing returned GRE traffic may create additional CPU overhead on the WCCP router. **Force L2 Return** may be used to override default behavior and route pass-through traffic back to the appliance's next-hop router, which may or may not be the WCCP router. Use caution, as this may create a Layer 3 loop, if L2 returned traffic gets redirected back to the appliance by the WCCP router.
 - **Password**. This field is *optional*.
 - **Priority**. The lowest priority is **0**, and the default value is **128**. Only change this setting from the default if an interface has multiple WCCP service groups defined for the same protocol (for example, TCP) and you wish to specify which service group to use.
 - **Weight**. The default value is **100**. You may use this to influence WCCP hash/mask assignments for individual appliances when more than one appliance is in a cluster. For Active/Backup appliance configuration, use a Weight of **0** on the backup appliance.

The *Hash* and *Mask* areas are only accessible when you select **custom** in the **Assignment Detail** field.

Route Policies Tab

Configuration > Route Policies

The **Route Policies** report displays the route policy entries that exist on the appliance(s).

This includes the appliance-based defaults, entries applied manually (via the WebUI or CLI), and entries that result from applying an Orchestrator Route Policies template.

Route Policies														Search	
21 Rows															
Edit	Appliance...	Map	Prio...	Match Criteria						Set Actions				Comment	
				ACL	Protoc...	Source IP/Subnet	Dest IP/Subnet	Application	Source...	DSCP	Interface	Destination	Path		Tunnel Down Action
	Los-Ange...	map1 (a...	20000	ip	any	any	any	any	0:0	any	VoIP	Voice		drop	Voice overlay
	Los-Ange...	map1 (a...	20001	ip	any	any	any	any	0:0	any	VoIP	Guest_Wifi		drop	Guest_Wifi overlay
	Los-Ange...	map1 (a...	65535	ip	any	any	any	any	0:0	any	any	auto optimized	default	pass-through	
	Dallas	map1 (a...	20000	ip	any	any	any	any	0:0	any	VoIP	Voice		drop	Voice overlay
	Dallas	map1 (a...	20001	ip	any	any	any	any	0:0	any	VoIP	Backup		drop	Backup overlay
	Dallas	map1 (a...	20002	ip	any	any	any	any	0:0	any	VoIP	Guest_Wifi		drop	Guest_Wifi overlay
	Dallas	map1 (a...	20003	ip	any	any	any	any	0:0	any	VoIP	Email		drop	Email overlay
	Dallas	map1 (a...	65535	ip	any	any	any	any	0:0	any	any	auto optimized	default	pass-through	
	Chicago	map1 (a...	20000	ip	any	any	any	any	0:0	any	VoIP	Voice		drop	Voice overlay
	Chicago	map1 (a...	20001	ip	any	any	any	any	0:0	any	VoIP	Backup		drop	Backup overlay
	Chicago	map1 (a...	20002	ip	any	any	any	any	0:0	any	VoIP	Guest_Wifi		drop	Guest_Wifi overlay
	Chicago	map1 (a...	20003	ip	any	any	any	any	0:0	any	VoIP	Email		drop	Email overlay
	Chicago	map1 (a...	65535	ip	any	any	any	any	0:0	any	any	auto optimized	default	pass-through	
	Seattle-EC	map1 (a...	20000	ip	any	any	any	any	0:0	any	VoIP	Voice		drop	Voice overlay
	Seattle-EC	map1 (a...	20001	ip	any	any	any	any	0:0	any	VoIP	Guest_Wifi		drop	Guest_Wifi overlay

Each appliance's default behavior is to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. **Auto-optimization** strategies reduce the need to create explicit route map entries for optimization. The three strategies that Silver Peak uses are **TCP-based** auto-opt, **IP-based** auto-opt, and **subnet sharing**. By default, all three are enabled on the **Templates** tab, under **System**.

The Route Policy, then, only requires entries for flows that are to be:

- sent pass-through (shaped or unshaped)
- dropped
- configured for a specific high-availability deployment
- routed based on application, VLAN, DSCP, or ACL (Access Control List)

You may also want to create a Route Policy entry when multiple tunnels exist to the remote *peer*, and you want the appliance to dynamically select the best path based on one of these criteria:

- load balancing
- lowest loss
- lowest latency
- specified tunnel

Manage these instances on the **Templates** tab, or click the **Edit** icon to manage Route policies directly for a particular appliance.

Priority

- You can create rules with any priority between 1 and 65534.
 - If you are using Orchestrator templates to add route map entries, the Orchestrator will delete all entries from **1000 – 9999**, inclusive, before applying its policies.
 - You can create rules from **1 – 999**, which have higher priority than Orchestrator template rules.
 - Similarly, you can create rules from **10000 – 65534** which have lower priority than Orchestrator template rules.
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- An IP address can specify a subnet - for example: 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

QoS Policies Tab

Configuration > QoS Policies

The QoS Policy determines how flows are queued and marked.

The **QoS Policies** tab displays the QoS policy entries that exist on the appliances.

This includes the appliance-based defaults, entries applied manually (via the WebUI or CLI), and entries that result from applying a Orchestrator QoS Policy template.

Both the Shaper and the QoS Policy deal with traffic classes. How are they related?

>> **The Shaper defines and the QoS Policy assigns.** <<

Use the **Templates** tab to create and manage QoS policies for multiple appliances, or click the **Edit** icon to manage QoS Policies directly for a particular appliance.

Edit	Appliance Name	Map	Priority	Match Criteria							Set Actions			Comment		
				ACL	Protocol	Source IP/Subnet	Dest IP/Subnet	Application	Source/Dest	DSCP	Interface	Traffic Class	LAN QoS		WAN QoS	
✓	Chicago	map1 (active)	20000		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Voice overlay
✓	Chicago	map1 (active)	20001		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Backup overlay
✓	Chicago	map1 (active)	20002		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Guest_Wifi overlay
✓	Chicago	map1 (active)	20003		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Email overlay
✓	Chicago	map1 (active)	65535		ip	any	any	any	any	0:0	any	any	1 - default	trust-lan	trust-lan	
✓	Dallas	map1 (active)	20000		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Voice overlay
✓	Dallas	map1 (active)	20001		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Backup overlay
✓	Dallas	map1 (active)	20002		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Guest_Wifi overlay
✓	Dallas	map1 (active)	20003		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Email overlay
✓	Dallas	map1 (active)	65535		ip	any	any	any	any	0:0	any	any	1 - default	trust-lan	trust-lan	
✓	Denver-EC	map1 (active)	20000		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Voice overlay
✓	Denver-EC	map1 (active)	20001		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Backup overlay
✓	Denver-EC	map1 (active)	20002		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Guest_Wifi overlay
✓	Denver-EC	map1 (active)	20003		ip	any	any	any	any	0:0	any	VoIP	1 - default	trust-lan	trust-lan	Email overlay
✓	Denver-EC	map1 (active)	65535		ip	any	any	any	any	0:0	any	any	1 - default	trust-lan	trust-lan	

The QoS Policy's SET actions determine two things:

- to what traffic class a shaped flow -- 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

Priority

- You can create rules with any priority between 1 and 65534.
 - If you are using Orchestrator templates to add route map entries, the Orchestrator will delete all entries from **1000 – 9999**, inclusive, before applying its policies.
 - You can create rules from **1 – 999**, which have higher priority than Orchestrator template rules.
 - Similarly, you can create rules from **10000 – 65534** which have lower priority than Orchestrator template rules.
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- An IP address can specify a subnet - for example: 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Optimization Policies Tab

Configuration > Optimization Policies

The **Optimization Policies** report displays a polled, read-only view of the Optimization policy entries that exist on the appliance(s). This includes the appliance-based defaults, entries applied manually (via the WebUI or CLI), and entries that result from applying an Orchestrator Optimization Policy template.

Use the **Templates** tab to create and manage Optimization policies.

		Match Criteria										Set Actions						
Edit	Appliance N...	Map	Priority	ACL	Protocol	Source IP/S...	Dest IP/Sub...	Applicatio...	Source:D...	DSCP	Interface	Network...	IP Heade...	Payload C...	TCP Accel	TCP Accel...	Protocol...	Comment
✓	Chicago	map1 (acti...	20000		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Voice overlay
✓	Chicago	map1 (acti...	20001		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Backup overlay
✓	Chicago	map1 (acti...	20002		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Guest_WiFi overlay
✓	Chicago	map1 (acti...	20003		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Email overlay
✓	Chicago	map1 (acti...	65535		ip	any	any	any	0:0	any	any	balanced	Yes	Yes	Yes	☐	none	
✓	Dallas	map1 (acti...	20000		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Voice overlay
✓	Dallas	map1 (acti...	20001		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Backup overlay
✓	Dallas	map1 (acti...	20002		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Guest_WiFi overlay
✓	Dallas	map1 (acti...	20003		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Email overlay
✓	Dallas	map1 (acti...	65535		ip	any	any	any	0:0	any	any	balanced	Yes	Yes	Yes	☐	none	
✓	Denver-EC	map1 (acti...	20000		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Voice overlay
✓	Denver-EC	map1 (acti...	20001		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Backup overlay
✓	Denver-EC	map1 (acti...	20002		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Guest_WiFi overlay
✓	Denver-EC	map1 (acti...	20003		ip	any	any	any	0:0	any	VoIP	balanced	Yes	Yes	Yes	☐	none	Email overlay
✓	Denver-EC	map1 (acti...	65535		ip	any	any	any	0:0	any	any	balanced	Yes	Yes	Yes	☐	none	
✓	Los-Angeles	map1 (acti...	10		ip	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	balanced	Yes	Yes	Yes	☐	none	

Set Actions Definitions

- **Network Memory** addresses limited bandwidth. This technology uses advanced fingerprinting algorithms to examine all incoming and outgoing WAN traffic. Network Memory localizes information and transmits only modifications between locations.
 - **Maximize Reduction** optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP, where bandwidth savings are the primary concern.
 - **Minimize Latency** ensures that Network Memory processing adds no latency. This may come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It's also appropriate when the primary objective is to fully utilize the WAN pipe to increase the LAN-side throughput, as opposed to conserving WAN bandwidth.
 - **Balanced** is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types.
 - **Disabled** turns off Network Memory.
- **IP Header Compression** is the process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end. It's possible to compress the protocol headers due to the redundancy in header fields of the same packet, as well as in consecutive packets of a packet stream.
- **Payload Compression** uses algorithms to identify relatively short byte sequences that are repeated frequently. These are then replaced with shorter segments of code to reduce the size of transmitted data. Simple algorithms can find repeated bytes within a single packet; more sophisticated algorithms can find duplication across packets and even across flows.

- **TCP Acceleration** uses techniques such as selective acknowledgements, window scaling, and message segment size adjustment to mitigate poor performance on high-latency links.
- **Protocol Acceleration** provides explicit configuration for optimizing CIFS, SSL, SRDF, Citrix, and iSCSI protocols. In a network environment, it's possible that not every appliance has the same optimization configurations enabled. Therefore, the site that initiates the flow (the client) determines the state of the protocol-specific optimization.

Priority

- You can create rules with any priority between 1 and 65534.
 - If you are using Orchestrator templates to add route map entries, Orchestrator will delete all entries from **1000 – 9999**, inclusive, before applying its policies.
 - You can create rules from **1 – 999**, which have higher priority than Orchestrator template rules.
 - Similarly, you can create rules from **10000 – 65534** which have lower priority than Orchestrator template rules.
- Adding a rule increments the last Priority by 10. This leaves room for you to insert a rule in between rules without having to renumber subsequent priorities. Likewise, you can just edit the number.

Source or Destination

- An IP address can specify a subnet - for example: 10.10.10.0/24 (IPv4) or fe80::204:23ff:fed8:4ba2/64 (IPv6).
- To allow **any IP address**, use 0.0.0.0/0 (IPv4) or ::/0 (IPv6).
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use **0**.

Access Lists Tab

Configuration > Access Lists

This tab lists the configured **Access Control List (ACL)** rules.

8 Rows		Search										
Edit	Appliance /Na...	ACLs	Priority	Match Criteria						Set Actions		Comment
				Protocol	Source IP/Subnet	Dest IP/Subnet	Application	Source:Dest ...	DSCP	Interface	Permit	
✓	Tallinn			No Access Lists defined for this appliance.								
✓	laine-va	bettermaybe	10	ip	10.10.10.0/24	0.0.0.0/0	bit_torrent	0:0	aF33	any	deny	
✓	laine-va	bettermaybe	20	ip	10.20.10.0/24	0.0.0.0/0	encrypted	0:0	cs1	any	permit	
✓	laine-va	office-lab	10	l2tp	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any,24	permit	
✓	laine-va	office-lab	20	ip	0.0.0.0/0	0.0.0.0/0	edonkey	0:0	any	any	deny	
✓	laine-vb			No Access Lists defined for this appliance.								
✓	laine2-va			No Access Lists defined for this appliance.								
✓	laine2-vb			No Access Lists defined for this appliance.								

An **ACL** is a reusable **MATCH** criteria for filtering flows, and is associated with an action, **permit** or **deny**: An ACL can be a **MATCH** condition in more than one policy --- Route, QoS, or Optimization.

- An Access Control List (ACL) consists of one or more ordered access control rules.
- An ACL only becomes active when it's used in a policy.
- **Deny** prevents further processing of the flow by *that ACL, specifically*. The appliance continues to the next entry in the policy.
- **Permit** allows the matching traffic flow to proceed on to the policy entry's associated **SET** action(s).

User Defined Applications Tab

Configuration > User Defined Applications

This tab lists **user-defined applications** (UDA).

Edit	Appliance Name	Priority	Application	Protocol	Source IP/Subnet	Dest IP/Subnet	Source Port/Range	Dest Port/Range	DSCP	Interface
	Tallinn	No User Defined Applications for this appliance.								
	laine-vxa	10	newfield	tcp	10.10.10.0/24	0.0.0.0/0	12-140	0	any	any
	laine-vxa	20	tomplins	ip	0.0.0.0/0	0.0.0.0/0	0	0	af32	any,607
	laine-vxb	No User Defined Applications for this appliance.								
	laine2-vxa	No User Defined Applications for this appliance.								
	laine2-vxb	No User Defined Applications for this appliance.								

UDAs are specific to the appliance on which they're defined.

Where can you use them?

- Route Policy
- QoS Policy
- Optimization Policy
- Access Lists (ACL)
- Application Groups

Behavior

- For reporting symmetry, you must define the same application(s) on peer appliances. Otherwise, the application may be a UDA on one appliance, and yet be categorized as an **unassigned application** on another, paired appliance.
- In the context of flow and application statistics reports, user-defined applications are always surveyed before built-in applications.
- **Ports are unique.** If a port or a range includes a built-in port, then the custom application is the one that owns it.
- If two distinctly named user-defined applications have a port number in common, then report results will be skewed, depending on the priority assigned to the custom applications. A port is only counted once.
- If a UDA is in use, deleting it deletes **all** the dependent entries. A warning message appears before deletion.
- Multiple UDAs can have the same name. Whenever that name is referenced, the software sequentially matches against each UDA definition having that name. So, dependent entries are only deleted when you delete the **last** definition of that UDA.

Application Groups Tab

Configuration > Application Groups

Application groups associate applications into a common group that you can use as a MATCH criteria. The applications can be built-in, user-defined, or a combination of both.

Edit	Appliance Name	Group Name	Applications
✎	Tallinn	citrix	citrix-bcast, citrix-cgp, citrix-ica, citrix-ima
✎	Tallinn	encrypted	ddm_ssl, https, imap4s, ipsec, nntp, pop3s, smtps, ssh, telnets
✎	Tallinn	interactive	pcanywhere, pcoip, ssh, telnet, telnets, vnc, xwindows
✎	Tallinn	real-time	cisco_skinny, h_323, rtcp, rtsp, t_120
✎	Tallinn	replication	app_assure_replication, app_assure_svr_backup, aspera, avamar, bluearc, celerra, centera, c...
✎	laine-vxa	citrix	citrix-bcast, citrix-cgp, citrix-ica, citrix-ima
✎	laine-vxa	encrypted	ddm_ssl, https, imap4s, ipsec, nntp, pop3s, smtps, ssh, telnets
✎	laine-vxa	interactive	pcanywhere, pcoip, ssh, telnet, telnets, vnc, xwindows
✎	laine-vxa	real-time	cisco_skinny, h_323, rtcp, rtsp, t_120
✎	laine-vxa	replication	app_assure_replication, app_assure_svr_backup, aspera, avamar, bluearc, celerra, centera, c...
✎	laine-vxb	citrix	citrix-bcast, citrix-cgp, citrix-ica, citrix-ima
✎	laine-vxb	encrypted	ddm_ssl, https, imap4s, ipsec, nntp, pop3s, smtps, ssh, telnets
✎	laine-vxb	interactive	pcanywhere, pcoip, ssh, telnet, telnets, vnc, xwindows
✎	laine-vxb	real-time	cisco_skinny, h_323, rtcp, rtsp, t_120
✎	laine-vxb	replication	app_assure_replication, app_assure_svr_backup, aspera, avamar, bluearc, celerra, centera, c...
✎	laine2-vxa	citrix	citrix-bcast, citrix-cgp, citrix-ica, citrix-ima

- The **Group Name** cannot be empty or have more than 64 characters.
- Group names are not case-sensitive.
- A group can be empty or contain up to 128 applications.
- An application group cannot contain an application group.
- For reporting symmetry, you must define the same application groups on peer appliances. Otherwise, the application group may be named on one appliance, and yet be categorized as an **unassigned application** on another, paired appliance.

NAT Policies Tab

Configuration > NAT Policies

This report has two views to show the NAT policies configured on appliances:

- The **Basic** view shows whether NAT is enabled on all **Inbound** and **Outbound**.

NAT Policies		NAT All Inbound			NAT All Outbound		
Appliance Name	Enable	NAT IP	Fallback	Enable	NAT IP	Fallback	
laine-vxa	<input type="checkbox"/>	auto	<input type="checkbox"/>	<input type="checkbox"/>	auto	<input type="checkbox"/>	
laine-vxb	<input type="checkbox"/>	auto	<input type="checkbox"/>	<input type="checkbox"/>	auto	<input type="checkbox"/>	
laine2-vxa	<input type="checkbox"/>	auto	<input type="checkbox"/>	<input type="checkbox"/>	auto	<input type="checkbox"/>	
laine2-vxb	<input type="checkbox"/>	auto	<input type="checkbox"/>	<input type="checkbox"/>	auto	<input type="checkbox"/>	


- The **Advanced** view displays all the NAT map rules.

NAT Policies		Match Criteria										Set Actions			Comment
Appliance	Map	Prio	ACL	Protoc	Source IP/Subnet	Dest IP/Subnet	Application	Source	DSCP	Interface	NAT Type	NAT Direct	NAT IP	Fallback	Comment
laine-vxa	map1 (acti...	65535	ip	any	any	any	any	0:0	any	any	no-nat	none	auto	<input type="checkbox"/>	
laine-vxb	map1 (acti...	65535	ip	any	any	any	any	0:0	any	any	no-nat	none	auto	<input type="checkbox"/>	
laine2-vxa	map1 (acti...	65535	ip	any	any	any	any	0:0	any	any	no-nat	none	auto	<input type="checkbox"/>	
laine2-vxb	map1 (acti...	65535	ip	any	any	any	any	0:0	any	any	no-nat	none	auto	<input type="checkbox"/>	

Two use cases illustrate the need for NAT:


- 1 **Inbound NAT.** The appliance automatically creates a source NAT (Network Address Translation) map when retrieving subnet information from the Silver Peak Cloud portal. This ensures that traffic destined to SaaS servers has a return path to the appliance from which that traffic originated.

NAT with a SaaS Service



- 2 **Outbound NAT.** The appliance and server are in the cloud, and the server accesses the internet. As in the example below, a Citrix thin client accesses its cloud-based server, and the server accesses the internet.

NAT with the Internet



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 addressing requirements.

- Enabling **NAT all** applies NAT policies to pass-through traffic as well as optimized traffic, ensuring that black-holing doesn't occur. **NAT all** on outbound only applies 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. You can do this by deploying the appliance in Router mode in-path with two (or four) interfaces.

Advanced Settings

The appliance can perform **source network address translation** (Source NAT or SNAT) on inbound or outbound traffic.

There are two types of NAT policies:

- **Dynamic** – created automatically by the system for inbound NAT when the **SaaS Optimization** feature is enabled and SaaS service(s) are selected for optimization. The appliance polls the *Silver Peak Unity Cloud Intelligence* service for a directory of SaaS services, and NAT policies are created for each of the subnets associated with selected SaaS service(s), ensuring that traffic destined for servers in use by those SaaS services has a return path to the appliance.
- **Manual** – created by the administrator for specific IP addresses / ranges or subnets. When assigning priority numbers to individual policies within a NAT map, first view **dynamic policies** to ensure that the manual numbering scheme doesn't interfere with dynamic policy numbering (that is, the manually assigned priority numbers cannot be in the range: 4000-5000). The default (**no-NAT**) policy is numbered 65535.

The NAT policy map has the following criteria and **Set Actions**:

- **Source or Destination**
 - 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.
- **NAT Type**
 - **no-nat** is the *default*. No IP addresses are changed.
 - **source-nat** changes the source address and the source port in the IP header of a packet.
- **NAT Direction**
 - **inbound** NAT is on the LAN interface.
 - **outbound** NAT is on the WAN interface.
 - **none** -- the only option if the NAT Type is **no-nat**.
- **NAT IP**
 - **auto** -- Select if you want to NAT **all** traffic. The appliance then picks the first available NAT IP/Port.
 - **tunnel** -- Select if you only want to NAT **tunnel** traffic. Applicable only for inbound NAT, as outbound doesn't support NAT on tunnel traffic.
 - **[IP address]** -- Select if you want to make NAT use this IP address during address translation.
- **Fallback** -- If the IP address is full, the appliance uses the next available IP address.

When you select a specific IP, then ensure that the routing is in place for NAT-ted return traffic.

SaaS Optimization Tab

Configuration > SaaS Optimization

When SaaS optimization is enabled, this report provides a view of the information retrieved from the *Silver Peak Unity Cloud Intelligence Service*.

Configuration Tab

To directly access an appliance and configure the SaaS applications/services you want to optimize, select the desired row and click Edit. Once you've accessed the appliance:

The screenshot displays the 'SaaS Optimization Configuration' page. At the top, there are tabs for 'Configuration' and 'Monitoring', and an 'Export' button. Below the tabs, the page title is 'SaaS Optimization Configuration'. A search bar is located on the right. The main content is a table with 101 rows. The table has the following columns: 'Appliance Name', 'Application Name', 'Optimize', 'Advertise', 'RTT Threshold', and 'Domains'. The 'Optimize' and 'Advertise' columns contain checkboxes. The 'RTT Threshold' column shows '10 ms' for all entries. The 'Domains' column lists various domain names for each application.

Appliance Name	Application Name	Optimize	Advertise	RTT Threshold	Domains
Chicago	No SaaS Optimization defined for this appliance.				
Dallas	No SaaS Optimization defined for this appliance.				
Denver-EC	Adobe	<input checked="" type="checkbox"/>	<input type="checkbox"/>	10 ms	adobe.com
Denver-EC	AirWatch	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.air-watch.com
Denver-EC	Box	<input checked="" type="checkbox"/>	<input type="checkbox"/>	10 ms	*.app.box.com, *.box.com, *.box.net, *.boxcdn.net, *.boxcloud.com
Denver-EC	CCConne	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	mycccportal.com, *.mycccportal.com
Denver-EC	ConstantContact	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	constantcontact.com
Denver-EC	CornerstoneOnDemand	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	cornerstoneondemand.com
Denver-EC	Dropbox	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	dropbox.com, *.dropbox.com
Denver-EC	Eloqua	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	eloqua.com, eloquatrainingcenter.com
Denver-EC	GoToAssist	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com
Denver-EC	GoToMeeting	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotomeeting.com
Denver-EC	GoToTraining	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com
Denver-EC	GoToWebinar	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotoweinar.com, gotoassist.com
Denver-EC	Intuit	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	intuit.com
Denver-EC	Jobvite	<input checked="" type="checkbox"/>	<input type="checkbox"/>	10 ms	careers.jobvite.com, www.jobvite.com, hire.jobvite.com

- **Enable SaaS optimization** enables the appliance to contact Silver Peak's *Unity Cloud Intelligence Service* and download information about SaaS services. This option is located on the appliance's **Configuration > SaaS Optimization** page.
- Initially, you may want to set a higher **RTT Threshold** value so that you can see a broader scope of reachable data centers/servers for any given SaaS application/service. As a best practice, production **RTT Threshold** values should not exceed 50 ms.
- You can use the **RTT Threshold** and **Location** columns on the appliance's **Monitoring > SaaS Optimization** page to help you determine if you should reposition the SaaS-enabled Silver Peak appliance closer to the SaaS data center.

Monitoring Tab

Topology SaaS Optimization x

Manage SaaS Optimization with Templates Configuration Monitoring Export ↻

SaaS Optimization Monitoring ?

75 Rows Search

Edit	Appliance Name ▲	Application N...	Subnet	Server IP	Advertised	RTT	RTT Thresho...	Ping Method	Ping Port	Location
✎	Chicago	<i>No SaaS Optimization defined for this appliance.</i>								
✎	Dallas	<i>No SaaS Optimization defined for this appliance.</i>								
✎	Denver-EC	Adobe	173.240.105...	173.240.105.1	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Adobe	173.240.103...	173.240.103...	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Adobe	173.240.108...	173.240.108...	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Adobe	173.240.110...	173.240.110...	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Adobe	173.240.105...	173.240.105...	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Adobe	173.240.110...	173.240.110...	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Adobe	173.240.102...	173.240.102...	No	Unreachable	10 ms			Indianapolis, United States
✎	Denver-EC	Box	74.112.184...	74.112.184.0	No	Unreachable	10 ms			Los Altos, United States
✎	Denver-EC	Box	64.79.128.4...	64.79.128.49	No	Unreachable	10 ms			Las Vegas, United States
✎	Denver-EC	Box	208.184.35...	208.184.35.1	No	Unreachable	10 ms			Redwood City, United States
✎	Denver-EC	Jobvite	52.7.65.30/32	52.7.65.30	No	Unreachable	10 ms			Ashburn, United States
✎	Denver-EC	Jobvite	54.210.23.2...	54.210.23.251	No	Unreachable	10 ms			Ashburn, United States
✎	Denver-EC	Salesforce	136.146.212...	136.146.212.2	No	Unreachable	10 ms			San Francisco, United States

Threshold Crossing Alerts Tab

Configuration > Threshold Crossing Alerts

Threshold Crossing Alerts (TCAs) are pre-emptive, user-configurable alarms triggered when specific thresholds are crossed.

Threshold Crossing Alerts

Edit	Appliance Name	Name	Rising				Falling			
			Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
	Tallinn	File-system utilization	90%	85%	1	<input checked="" type="checkbox"/>	75%	75%	1	<input type="checkbox"/>
	Tallinn	LAN-side receive throughput	1000000 kbps	1000000 kbps	1	<input checked="" type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>
	Tallinn	Total number of flows	57600 flows	54400 flows	1	<input checked="" type="checkbox"/>	0 flows	0 flows	1	<input type="checkbox"/>
	Tallinn	Total number of optimized flows	256000 flows	256000 flows	1	<input type="checkbox"/>	0 flows	0 flows	1	<input type="checkbox"/>
	Tallinn	Tunnel OOP post-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
	Tallinn	Tunnel OOP pre-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
	Tallinn	Tunnel latency	1000 ms	850 ms	1	<input checked="" type="checkbox"/>	0 ms	0 ms	1	<input type="checkbox"/>
	Tallinn	Tunnel loss post-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
	Tallinn	Tunnel loss pre-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
	Tallinn	Tunnel reduction	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
	Tallinn	Tunnel utilization	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
	Tallinn	WAN-side transmit throughput	1000000 kbps	1000000 kbps	1	<input type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>
	laime-xxa	File-system utilization	90%	85%	1	<input checked="" type="checkbox"/>	75%	75%	1	<input type="checkbox"/>
	laime-xxa	LAN-side receive throughput	1000000 kbps	1000000 kbps	1	<input type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>
	laime-xxa	Total number of flows	90%	85%	1	<input checked="" type="checkbox"/>	0%	0%	1	<input type="checkbox"/>


Threshold Crossing Alerts

Edit	Name	Tunnel Name	TCA Name	Rising			Falling			
				Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
		auto_tun_1...	latency	15 ms	NaN	<input type="checkbox"/>	undefined ms	undefined ms	NaN	<input type="checkbox"/>

They alarm on both rising and falling threshold crossing events (i.e., floor and ceiling levels). For both levels, one value raises the alarm, while another value clears it.

- When you configure appliance and tunnel TCAs with an Orchestrator template, all alerts apply globally, so all of an appliance's tunnels have the same alerts.
- To create a tunnel-specific alert, go to **Configuration > Tunnels**, select the tunnel, click the Edit icon to access the tunnel directly, and then click the icon in the **Alert Options** column. Make your changes and click **OK**.
- To view globally applied system and tunnel alerts, click **System**.
- To view alerts that are specific to an individual tunnel, click **Tunnel**.

Times to Trigger - A value of 1 triggers an alarm on the first threshold crossing instance.



Rules:


- High raise threshold is greater than high clear threshold
- Low raise threshold is less than low clear threshold

ON by default:

- **Appliance Capacity** - triggers when an appliance reaches 95% of its total flow capacity. It is not configurable and can only be cleared by an operator.
- **File-system utilization** - percent of non-Network Memory disk space filled by the appliance. This TCA cannot be disabled.
- **Tunnel latency** - measured in milliseconds, the maximum latency of a one-second sample within a 60-second span

OFF by default:

- **LAN-side receive throughput** - based on a one-minute average, the LAN-side receive **TOTAL** for all interfaces
- **WAN-side transmit throughput** - based on a one-minute average, the WAN-side transmit **TOTAL** for all interfaces
- **TCAs based on an end-of-minute count:**
 - Total number of flows
 - Total number of optimized flows
- **TCAs based on a one-minute average:**
 - Tunnel loss post-FEC
 - Tunnel loss post-FEC
 - Tunnel OOP post-POC
 - Tunnel OOP post-POC
 - Tunnel reduction
 - Tunnel utilization (based on percent of configured maximum [system] bandwidth)

 **Note** Enabled by default, there is also an **Appliance Capacity** TCA that triggers when an appliance reaches 95% of its total flow capacity. It doesn't automatically clear, but can be cleared by an operator. It is also not configurable.

This table lists the **defaults** of each type of threshold crossing alert:

Table 4-1 Defaults values for Threshold Crossing Alerts

	TCA Name	Default [ON, OFF]	Default Values		allow rising	allow falling
			[Rising Raise, Rising Clear, Falling Raise, Falling Clear]			
Appliance Level	WAN-side transmit throughput	OFF	1 Gbps; 1 Gbps; 0; 0		4	4
	LAN-side receive throughput	OFF	1 Gbps; 1 Gbps; 0; 0		4	4
	Total number of optimized flows	OFF	256,000, 256,000; 0; 0		4	4
	Total number of flows	OFF	256,000, 256,000; 0; 0		4	4
	File-system-utilization	ON^a	95%; 85%; 0%; 0%		4	
Tunnel Level	Tunnel latency	ON	1000; 850; 0; 0		4	
	Tunnel loss pre-FEC	OFF	100%; 100%; 0%; 0%		4	
	Tunnel loss post-FEC	OFF	100%; 100%; 0%; 0%		4	
	Tunnel OOP pre-POC	OFF	100%; 100%; 0%; 0%		4	
	Tunnel OOP post-POC	OFF	100%; 100%; 0%; 0%		4	
	Tunnel utilization	OFF	95%; 90%; 0%; 0%		4	4
	Tunnel reduction	OFF	100%; 100%; 0%; 0%			4

a. Cannot be disabled.



Appliance Administration Tabs

This chapter describes the reports that display appliance administration parameters.

In This Chapter

- **Date/Time Tab** See page 138.
- **Domain Name Servers (DNS) Tab** See page 139.
- **SNMP Tab** See page 140.
- **NetFlow Tab** See page 141.
- **Logging Tab** See page 142.
- **Appliance User Accounts Tab** See page 144.
- **Auth/RADIUS/TACACS+ Tab** See page 145.
- **Banners Tab** See page 147.

Date/Time Tab

Administration > [General] Date/Time

This tab highlights significant time discrepancies among the devices recording statistics.

Relative to the appliance's
configured time

Edit	Appliance Name	Time Zone	NTP Enabled	NTP servers	Appliance Date/Time	Orchestrator Delta	Browser Delta
	Tallinn	UTC	<input type="checkbox"/>		2016/02/19 02:51:27	-0 hrs : 0 mins : 40 secs	-0 hrs : 0 mins : 40 secs
	laine-vxa	UTC	<input type="checkbox"/>		2016/02/19 02:22:55	-0 hrs : 29 mins : 12 secs	-0 hrs : 29 mins : 12 secs
	laine-vxb	UTC	<input type="checkbox"/>		2016/02/19 02:23:31	-0 hrs : 28 mins : 36 secs	-0 hrs : 28 mins : 36 secs
	laine2-vxa	UTC	<input checked="" type="checkbox"/>	172.20.20.37 (Version 3)	2016/02/19 02:52:07	-0 hrs : 0 mins : 0 secs	+0 hrs : 0 mins : 0 secs
	laine2-vxb	UTC	<input checked="" type="checkbox"/>	172.20.20.37 (Version 3)	2016/02/19 02:52:07	-0 hrs : 0 mins : 0 secs	+0 hrs : 0 mins : 0 secs

Appliance times should be within 1min of Orchestrator time AND client (browser) time - NTP is recommended.

If the **date and time** of an appliance, the Orchestrator server, and your browser aren't all synchronized, then charts (and stats) will inevitably have different timestamps for the same data, depending on which device you use to view the reports.

Recommendation: For consistent results, configure the appliance, the Orchestrator server, and your PC to use an NTP (Network Time Protocol) server.

Domain Name Servers (DNS) Tab

Administration > [General] DNS

This tab lists the Domain Name Servers that the appliances reference.

Topology | DNS x

Manage DNS with Templates | Export | ↻

DNS ?

5 Rows | Search

Edit	Appliance Na...	Primary DNS IP addr	Secondary DNS IP addr	Tertiary DNS IP addr	Domain Names
↗	Chicago	No DNS settings defined for this appliance.			
↗	Dallas	No DNS settings defined for this appliance.			
↗	Denver-EC	No DNS settings defined for this appliance.			
↗	Los-Angeles	1.1.1.1			
↗	Seattle-EC	No DNS settings defined for this appliance.			

A **Domain Name Server** (DNS) uses a table to map domain names to IP addresses. So, you can reference locations by a domain name, such as *mycompany.com*, instead of using the IP address.

Each appliance can support up to three name servers.

SNMP Tab

Administration > [General] SNMP

This tab summarizes what SNMP capabilities are enabled and which hosts can receive SNMP traps.

		SNMP			Trap Receivers		
Edit	Appliance Name	Enable SNMP	Enable SNMP Traps	Enable V3 User	Trap Receiver 1	Trap Receiver 2	Trap Receiver 3
	Tallinn	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			
	laine-vxa	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			
	laine-vxb	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			
	laine2-vxa	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			
	laine2-vxb	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			

- The Silver Peak appliance supports the Management Information Base (MIB) II, as described in RFC 1213, for cold start traps and warm start traps, as well as Silver Peak proprietary MIBs.
- The appliance issues an SNMP trap during reset--that is, when loading a new image, recovering from a crash, or rebooting.
- The appliance sends a trap every time an alarm is raised or cleared. Traps contain additional information about the alarm, including severity, sequence number, a text-based description of the alarm, and the time the alarm was created.

Term	Definition
Enable SNMP	Allows the SNMP application to poll this Silver Peak appliance. (For SNMP v1 and SNMP v2c)
Enable SNMP Traps	Allows the SNMP agent (in the appliance) to send traps to the receiver(s). (For SNMP v1 and SNMP v2c)
Enable V3 User	For additional security when the SNMP application polls the appliance, you can use SNMP v3, instead of using v1 or v2c. This provides a way to authenticate without using clear text.
Trap Receiver	IP address of a host configured to receive SNMP traps.

NetFlow Tab

Administration > [General] NetFlow

This tab summarizes how the appliances are configured to export statistical data to NetFlow collectors.

5 Rows		Search						
Edit	Appliance Name	Flow Exporting ...	Active Flow Time...	Traffic Type	Collector1		Collector2	
					IP Address	Port	IP Address	Port
	Tallinn	<input type="checkbox"/>	1	Outbound WAN				
	laine-vxa	<input checked="" type="checkbox"/>	1	Outbound WAN				
	laine-vxb	<input type="checkbox"/>	1	Outbound WAN				
	laine2-vxa	<input type="checkbox"/>	1	Outbound WAN				
	laine2-vxb	<input type="checkbox"/>	1	Outbound WAN				

- The appliance exports flows against two virtual interfaces -- **sp_lan** and **sp_wan** -- that accumulate the total of LAN-side and WAN-side traffic, regardless of physical interface.
- These interfaces appear in SNMP and are therefore "discoverable" by NetFlow collectors.
- **Flow Exporting Enabled** allows the appliance to export the data to collectors (and makes the configuration fields accessible).
- The **Collector's IP Address** is the IP address of the device to which you're exporting the NetFlow statistics. The default Collector Port is **2055**.
- In **Traffic Type**, you can select as many of the traffic types as you wish. The default is **Outbound WAN**.

Logging Tab

Administration > [General] Logging

This tab summarizes the configured logging parameters:

- **Log Configuration** refers to local logging.
- **Log Facilities Configuration** refers to remote logging.

		Log Configuration			Log Facilities Configuration		
Edit	Appliance Name	Minimum Severity	Log File Size Threshold	Number of Logs to Ke...	System	Audit	Flow
✓	Tallinn	Notice	50	30	local1	local0	local2
✓	laine-vxa	Notice	50	30	local1	local0	local2
✓	laine-vxb	Notice	50	30	local1	local0	local2
✓	laine2-vxa	Notice	50	30	local1	local0	local2
✓	laine2-vxb	Notice	50	30	local1	local0	local2

Minimum Severity Levels

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

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.
INFORMATIONAL	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.
- 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**.

Remote Logging

- You can 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 may not accept as low a severity level as you are forwarding to it.
- Each message/event type (**System / Audit / Flow**) is assigned to a syslog facility level (**local0** to **local7**).

Appliance User Accounts Tab

Administration > [User Management] Users

This tab provides data about the **user accounts** on each appliance.

The screenshot shows the 'User Accounts' tab in the Silver Peak Unity Orchestrator. The interface includes a search bar, an 'Export' button, and a table with 11 rows. The table columns are 'Edit', 'Appliance Name', 'User Name', 'Capability', and 'Enabled'. The rows are as follows:

Edit	Appliance Name	User Name	Capability	Enabled
	Tallinn	admin	admin	<input checked="" type="checkbox"/>
	Tallinn	monitor	monitor	<input checked="" type="checkbox"/>
	laine-vxa	admin	admin	<input checked="" type="checkbox"/>
	laine-vxa	monitor	monitor	<input checked="" type="checkbox"/>
	laine-vxa	myself	admin	<input checked="" type="checkbox"/>
	laine-vxb	admin	admin	<input checked="" type="checkbox"/>

- The Silver Peak appliance's **built-in user database** supports user names, groups, and passwords.
- Each appliance has two default users, **admin** and **monitor**, who cannot be deleted.
- Each **User Name** belongs to one of two user groups -- **admin** or **monitor**.
 - The **monitor** group supports reading and monitoring of all data, in addition to performing all actions. This is equivalent to the Command Line Interface's (CLI) enable mode privileges.
 - The **admin** group supports full privileges, along with permission to add, modify, and delete. This is equivalent to the CLI's **configuration** mode privileges.
- Named user accounts can be added via Appliance Manager or the Command Line Interface (CLI).
- The table lists all users known to the appliances, whether or not their accounts are enabled.

Auth/RADIUS/TACACS+ Tab

Administration > [User Management] Auth/RADIUS/TACACS+

This tab displays the configured settings for authentication and authorization.

If the appliance relies on either a RADIUS or TACACS+ server for those services, then those settings are also reported.

All settings are initially applied via the **Auth/RADIUS/TACACS+** configuration **template**.

Authentication and Authorization

The screenshot shows the configuration page for Authentication and Authorization. It includes a search bar and a table with 5 rows. The table columns are: Edit, Appliance Name, Authentication Order (First, Second, Third), and Authorization Information (Map Order, Default User).

Edit	Appliance Name	Authentication Order			Authorization Information	
		First	Second	Third	Map Order	Default User
	Tallinn	local			remote-first	admin
	laine-vxa	local			remote-first	admin
	laine-vxb	local			remote-first	admin
	laine2-vxa	local			remote-first	admin
	laine2-vxb	local			remote-first	admin

- **Authentication** is the process of validating that the end user, or a device, is who they claim to be.
- **Authorization** is the action of determining what a user is allowed to do. Generally, authentication precedes authorization.
- When it's possible to validate against more than one database (local, RADIUS server, TACACS+ server), **Authentication Order** specifies which method to try in what sequence.
- **Map order**. The default—and recommended—value is **remote-first**.
- **Default user**. The default—and recommended—value is **admin**.

RADIUS and TACACS+

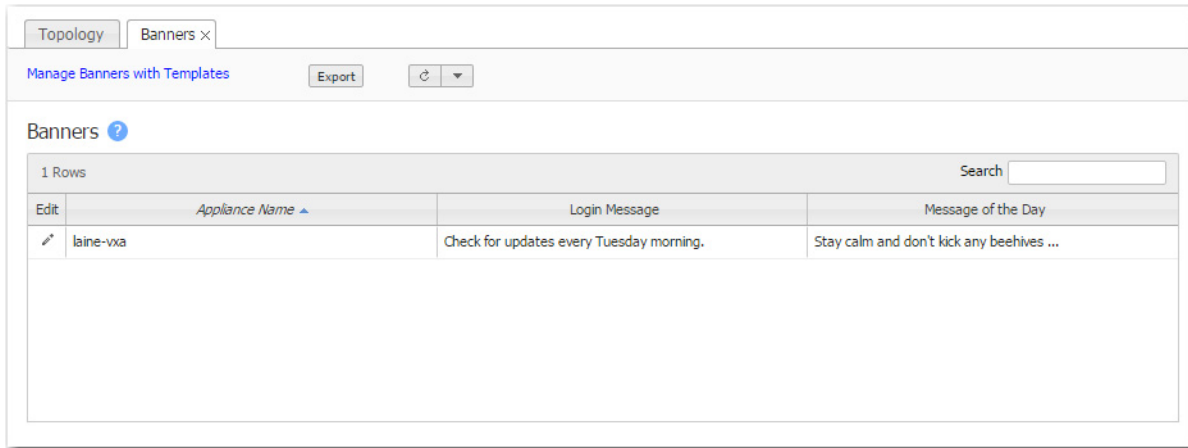
The screenshot shows the 'Auth/RADIUS/TACACS+' configuration page. At the top, there are tabs for 'Topology' and 'Auth/RADIUS/TACACS+' with a close button. Below the tabs, there is a navigation bar with 'Manage Authentication with Templates', 'Authentication', 'RADIUS and TACACS', and 'Export' buttons. The main content area is titled 'RADIUS/TACACS+ Servers' and features a search bar. Below the search bar is a table with the following columns: Edit, Appliance Name, Server Type, Order, Server IP, Auth Port, Auth Type, Timeout, Retries, and Enabled. The table is currently empty, and the text 'No Data Available' is displayed in the center.

- **Server Type.** RADIUS or TACACS+
- **Auth Port.** For RADIUS, the default value is **1812**. For TACACS+, the default value is **49**.
- **Auth Type.** [TACACS+] The options are **pap** or **ascii**.
- **Timeout.** If a logged-in user is inactive for an interval that exceeds the inactivity time-out, the appliance logs them out and returns them to the login page. You can change that value, as well as the maximum number of sessions, in the **Session Management template**.
- **Retries.** The number of retries allowed before lockout.
- **Enabled.** Whether or not the server is enabled.

Banners Tab

Administration > [User Management] Banners

This tab lists the banner messages on each appliance.



The screenshot displays the 'Banners' tab in a management interface. At the top, there are tabs for 'Topology' and 'Banners x'. Below the tabs, there is a section titled 'Manage Banners with Templates' with an 'Export' button and a refresh icon. The main content area is titled 'Banners' and shows '1 Rows'. A search bar is located on the right side of the table. The table has three columns: 'Appliance Name', 'Login Message', and 'Message of the Day'. The data row shows the appliance name 'laine-vxa', the login message 'Check for updates every Tuesday morning.', and the message of the day 'Stay calm and don't kick any beehives ...'.

Edit	Appliance Name ▲	Login Message	Message of the Day
✎	laine-vxa	Check for updates every Tuesday morning.	Stay calm and don't kick any beehives ...

- The **Login Message** appears before the login prompt.
- The **Message of the Day** appears after a successful login.

Alarms

Monitoring > Alarms

This chapter describes alarm categories and definitions. It also describes how to view and handle alarm notifications.



Threshold crossing alerts are related to alarms. They are preemptive, user-configurable thresholds that declare a Major alarm when crossed. For more information about their configuration and use, see *“Threshold Crossing Alerts Template” on page 74* and *“Threshold Crossing Alerts Tab” on page 134*.

In This Chapter

- **Understanding Alarms** See page 150.
- **Viewing Alarms** See page 161.
- **Specifying Alarm Recipients** See page 163.

Understanding Alarms

The Orchestrator and the appliances have separate alarm summaries and alarm tables.

Silver Peak appliances "push" all their alarms to the Orchestrator database, which then updates the Appliance Alarms table.

Each entry represents one current condition that may require human intervention. Because alarms are *conditions*, they may come and go without management involvement.

Whereas merely acknowledging most alarms does **not** clear them, some alarm conditions are set up to be self-clearing when you acknowledge them. For example, if you remove a hard disk drive, it generates an alarm; once you've replaced it and it has finished rebuilding itself, the alarm clears.

Categories of Alarms

Alarms have one of four severity levels: **Critical**, **Major**, **Minor**, and **Warning**. Only **Critical** and **Major** alarms are service-affecting.



- **Critical** alarms require immediate attention, and reflect conditions that affect an appliance or the loss of a broad category of service.
- **Major** alarms reflect conditions which should be addressed in the next 24 hours. An example would be an unexpected traffic class error.
- **Minor** alarms can be addressed at your convenience. An example of a minor alarm would be a user not having changed their account's default password, or a degraded disk.
- **Warnings** inform you of conditions that may become problems over time. For example, a software version mismatch.

Types of Appliance Alarms

The appliance can raise alarms based on issues with tunnels, software, equipment, and Threshold Crossing Alerts (TCAs). The latter are visible on the appliance but managed by the Orchestrator.

Although Appliance Manager (the WebUI) doesn't display **Alarm Type ID (Hex)** codes, the data is available for applications that can do their own filtering, such as SNMP.

Table 6-1 Silver Peak Appliance Alarms

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Tunnel	00010009	CRITICAL	An unexpected GRE packet was detected from tunnel peer. RESOLUTION: Check for tunnel encapsulation mismatch.
	00010003	CRITICAL	Tunnel keepalive version mismatch RESOLUTION: Tunnel peers are running incompatible software versions. <ul style="list-style-type: none"> • Normal during a software upgrade. • Run the same or compatible software releases among the tunnel peers.
	00010008	CRITICAL	Tunnel local IP address not owned by this appliance. RESOLUTION: Delete the tunnel and re-create it with a valid IP address.
	00010001	CRITICAL	Tunnel state is Down RESOLUTION: Cannot reach tunnel peer. <ul style="list-style-type: none"> • Check tunnel configuration [Admin state, Source IP/Dest IP, IPsec] • Check network connectivity.
	00010007	MAJOR	Duplicate license detected in peer (only applies to virtual appliance) RESOLUTION: Install unique license on all virtual appliances. To check and/or change license: <ul style="list-style-type: none"> • In WebUI: Administration > License & Registration • In Orchestrator: Administration > Licenses
	0001000a	MAJOR	Software version mismatch between peers results in reduced functionality. RESOLUTION: Tunnel peers are not running the same release of software. This results in reduced functionality. Run the same or compatible software releases among the tunnel peers.
	00010000	MAJOR	Tunnel remote ID is misconfigured RESOLUTION: System ID is not unique. <ul style="list-style-type: none"> • Virtual Appliance: Was the same license key used? • Physical Appliance: Change System ID in the rare case of a duplicate ID (CLI command: <code>system id < ></code>)
	00010005	MINOR	Tunnel software version mismatch RESOLUTION: Tunnel are not running the same release of software. They will function, but with reduced functionality. <ul style="list-style-type: none"> • Normal during an upgrade. • Run the same software version to eliminate the alarm and fully optimize.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Software	0004001c	CRITICAL	EC license not granted RESOLUTION: Please obtain additional EC (EdgeConnect) licenses.
	0004000c	CRITICAL	Invalid virtual appliance license. RESOLUTION: Enter a new license key on the <System Page> to proceed.
	00040016	CRITICAL	Software capability license has expired. RESOLUTION: You must have HTTPS connectivity to internet to renew the licensing token.
	00040003	CRITICAL	The licensing for this virtual appliance has expired. [For VX series only] ^a RESOLUTION: Enter a new license.
	00040004	CRITICAL	There is no license installed on this virtual appliance. [For VX series only] ^a RESOLUTION: Enter a valid license.
	00040005	MAJOR	A disk self-test has been run on the appliance. RESOLUTION: Reboot the appliance. Traffic will not be optimized until this is performed.
	00040013	MAJOR	A peer name has been specified in the route-map configuration matching no existing remote peer RESOLUTION: Correct route-map entry or build tunnel.
	00040019	MAJOR	Application deleted on portal RESOLUTION: Contact Customer Service.
	0004000d	MAJOR	Dual wan-next-hop topology is no longer supported. RESOLUTION: Reconfigure appliance as single bridge with one next-hop, or dual bridge with two IP addresses and two next-hops.
	00040010	MAJOR	Major inconsistency among tunnel traffic class settings found during upgrade. RESOLUTION: Review the WAN shaper traffic class settings.
	0004001b	MAJOR	Portal registration data incorrect RESOLUTION: Please provide valid portal account registration information.
	00040002	MAJOR	Significant change in time of day has occurred, and might compromise statistics. Please contact TAC. RESOLUTION: Appliance statistics could be missing for a substantial period of time. Contact Customer Service.
	00040015	MAJOR	Software capability license needs to be renewed before it expires. RESOLUTION: Software will automatically renew the licensing token as long as it has HTTPS connectivity to the internet.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Software (cont.)	00040001	MAJOR	System is low on resources RESOLUTION: Contact Customer Service.
	00040011	MAJOR	Tunnel IP header disable setting was discarded during upgrade. RESOLUTION: Review the optimization map header compression settings.
	0004000a	MAJOR	Virtual appliance license expires on mm/dd/yyyy. [15-day warning] RESOLUTION: Enter a new license key on the <System Page> to avoid loss of optimization or potential traffic disruption.
	0004001a	MINOR	Performance limited by maximum Boost bandwidth RESOLUTION: Recommend subscribing to more Boost bandwidth.
	00040012	WARNING	A very large range has been configured for a local subnet. RESOLUTION: Please confirm that you intended to configure such a large local subnet.
	00040014	WARNING	Interface shaper max bandwidth exceeds system max bandwidth RESOLUTION: Review the interface shaper max bandwidth settings. Please make sure it doesn't exceed system max bandwidth.
	0004000f	WARNING	Minor inconsistency among tunnel traffic class settings found during upgrade. RESOLUTION: Review the WAN shaper traffic class settings.
	0004000e	WARNING	Setting default system next-hop to VLAN next-hop no longer necessary. RESOLUTION: Use the VLAN IP address as tunnel source endpoints instead of bvi0.
	00040017	WARNING	Silver Peak portal is unreachable. RESOLUTION: Appliance cannot connect to Silver Peak portal using HTTPS. This connectivity is needed for internet applications classification.
	00040018	WARNING	Silver Peak portal is unreachable. RESOLUTION: Appliance cannot connect to Silver Peak portal using HTTPS Websockets.
	00040009	WARNING	The NTP server is unreachable. RESOLUTION: Check the appliance's NTP server IP and version configuration: <ul style="list-style-type: none"> • Can the appliance reach the NTP server? • Is UDP port 123 open between the appliance's mgmt0 IP and the NTP server?
	00040007	WARNING	The SSL certificate is not yet valid. RESOLUTION: The SSL certificate has a future start date. It will correct itself when the future date becomes current. Otherwise, install a certificate that is current.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Software (cont.)	00040008	WARNING	The SSL certificate has expired. RESOLUTION: Reinstall a valid SSL certificate that is current.
	00040006	WARNING	The SSL private key is invalid. RESOLUTION: The key is not an RSA standard key that meets the minimum requirement of 1024 bits. Regenerate a key that meets this minimum requirement.
	0004000b	WARNING	Virtual appliance license expires on mm/dd/yyyy. [45-day warning] RESOLUTION: Enter a new license key on the <System Page> to avoid loss of optimization or potential traffic disruption.
Equipment	0003002b	CRITICAL	Bridge creation failed RESOLUTION: Check log messages for more details on the failure.
	00030029	CRITICAL	Bridge loop is detected RESOLUTION: Make sure bridge ports are connected to different virtual switches and restart the appliance. Traffic will not be optimized until this is resolved.
	00030007	CRITICAL	Encryption card hardware failure RESOLUTION: Contact Customer Service.
	00030003	CRITICAL	Fan failure detected RESOLUTION: Contact Customer Service.
	00030024	CRITICAL	Insufficient configured memory size for this virtual appliance RESOLUTION: Assign more memory to the virtual machine, and restart the appliance. Traffic will not be optimized until this is resolved.
	00030025	CRITICAL	Insufficient configured processor count for this virtual appliance RESOLUTION: Assign more processors to the virtual machine, and restart the appliance. Traffic will not be optimized until this is resolved.
	00030026	CRITICAL	Insufficient configured disk storage for this virtual appliance RESOLUTION: Assign more storage to the virtual machine, and restart the appliance. Traffic will not be optimized until this is resolved.
	00030005	CRITICAL	LAN/WAN fail-to-wire card failure RESOLUTION: Contact Customer Service.
	0003002a	CRITICAL	Network interface is unassigned RESOLUTION: Assign the network interface to an existing MAC address, and then restart the appliance. Or, if the network interface isn't being used, then set its admin state to down.
	00030021	CRITICAL	NIC interface failure RESOLUTION: Contact Customer Service.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Equipment (cont.)	00030004	CRITICAL	System is in Bypass mode RESOLUTION: Normal with factory default configuration, during reboot, and if user has put the appliance in Bypass mode. Contact Customer Service if the condition persists.
	0003001d	MAJOR	Bonding members have different speed/duplex RESOLUTION: Check interface speed/duplex settings and negotiated values on wan0/wan1 and lan0/lan1 etherchannel groups.
	0003001c	MAJOR	[Flow redirection] cluster peer is down RESOLUTION: <ul style="list-style-type: none"> • Check flow redirection configuration on all applicable appliances. • Check L3/L4 connectivity between the peers. • Open TCP and UDP ports 4164 between the cluster peer IPs if they are blocked.
	00030017	MAJOR	Disk removed by operator RESOLUTION: Normal during disk replacement. Insert disk using UI/Orchestrator. Contact Customer Service if insertion fails.
	00030001	MAJOR	Disk is failed RESOLUTION: Contact Customer Service to replace disk.
	00030015	MAJOR	Disk is not in service RESOLUTION: <ul style="list-style-type: none"> • Check to see if the disk is properly seated. • Contact Customer service for further assistance.
	0003000b	MAJOR	Interface is half duplex RESOLUTION: Check speed/duplex settings on the router/switch port.
	0003000c	MAJOR	Interface speed is 10 Mbps RESOLUTION: <ul style="list-style-type: none"> • Check speed/duplex settings. • Use a 100/1000 Mbps port on the router/switch.
	00030027	MAJOR	Interfaces have different MTUs. [LAN0 and WAN0]. RESOLUTION: Check interface MTU settings on lan0/wan0(pairwise) on dual bridge mode and lan0/lan1/wan0/wan1... on single bridge mode.
	00030028	MAJOR	Interfaces have different MTUs. [LAN1 and WAN1]. RESOLUTION: Check interface MTU settings on lan1/wan1 or tlan1/twan1 interfaces.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Equipment (cont.)	00030022	MAJOR	LAN next-hop unreachable ^b RESOLUTION: Check appliance configuration: <ul style="list-style-type: none"> • LAN-side next-hop IP • Appliance IP / Mask • VLAN IP / Mask • VLAN ID
	0003001a	MAJOR	LAN/WAN interface has been shut down due to link propagation of paired interface RESOLUTION: Check cables and connectivity. For example, if lan0 is shut down, check why wan0 is down. Applicable only to in-line (bridge) mode.
	00030018	MAJOR	LAN/WAN interfaces have different admin states RESOLUTION: Check interface admin configuration for lan0/wan0 (and lan1/wan1). Applicable only to in-line mode.
	00030019	MAJOR	LAN/WAN interfaces have different link carrier states RESOLUTION: Check interface configured speed settings and current values (an0/wan0, lan1/wan1). Applicable only to in-line mode.
	0003000a	MAJOR	Management interface link down RESOLUTION: <ul style="list-style-type: none"> • Check cables. • Check interface admin status on the router.
	00030009	MAJOR	Network interface link down RESOLUTION: Is the system in Bypass mode? <ul style="list-style-type: none"> • Check cables. • Check interface admin status on the router.
	00030020	MAJOR	Power supply not connected, not powered, or failed RESOLUTION: <ul style="list-style-type: none"> • Connect to a power outlet. • Check power cable connectivity.
	0003002c	MAJOR	System optimization disabled RESOLUTION: Turn on system optimization.
	00030023	MAJOR	Unexpected system restart RESOLUTION: Power issues? Was the appliance shutdown ungracefully? Contact Customer Service if the shutdown was not planned.
	00030012	MAJOR	VRRP instance is down RESOLUTION: Check the interface. Is the link down?

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Equipment (cont.)	00030014	MAJOR	WAN next-hop router discovered on a LAN port (box is in backwards) RESOLUTION: <ul style="list-style-type: none"> • Check WAN next-hop IP address. • Check lan0 and wan0 cabling (in-line mode only). • If it cannot be resolved, call Customer Service.
	00030011	MAJOR	WAN next-hop unreachable ^b RESOLUTION: <ul style="list-style-type: none"> • Check cables on Silver Peak appliance and router. • Check IP/mask on Silver Peak appliance and router. Next-hop should be only a single IP hop away. • To troubleshoot, use: <code>show cdp neighbor,</code> <code>show arp,</code> and <code>ping -I <appliance IP> <next-hop IP>.</code>
	0003001e	MAJOR	WCCP adjacency(ies) down RESOLUTION: Cannot establish WCCP neighbor: <ul style="list-style-type: none"> • Check WCCP configuration on appliance and router. • Verify reachability. • Enable debugging on router: <code>debug ip wccp packet</code>
	0003001f	MAJOR	WCCP assignment table mismatch RESOLUTION: Check WCCP mask/hash assignment configuration on all Silver Peak appliances and ensure that they match.
	00030002	MINOR	Disk is degraded RESOLUTION: Wait for disk to recover. If it does not recover, contact Customer Service.
	00030016	MINOR	Disk is rebuilding RESOLUTION: Normal. If rebuilding is unsuccessful, contact Customer Service.
	0003001b	MINOR	Disk SMART threshold exceeded RESOLUTION: Contact Customer Service to replace disk.
	0003002d	MINOR	Non-optimal configured memory size for this virtual appliance RESOLUTION: Assign more memory to the virtual machine and restart the appliance. Traffic will be sub-optimal until this is resolved.
	0003002e	MINOR	Non-optimal configured processor count for this virtual appliance RESOLUTION: Assign more processors to the virtual machine and restart the appliance. Traffic will be sub-optimal until this is resolved.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Equipment (cont.)	0003002f	MINOR	Non-optimal configured disk storage for this virtual appliance RESOLUTION: Assign more storage to the virtual machine and restart the appliance. Traffic will be sub-optimal until this is resolved.
	00030008	WARNING	Network interface admin down RESOLUTION: Check Silver Peak interface configuration.
	00030013	WARNING	VRRP state changed from Master to Backup RESOLUTION: VRRP state has changed from Master to Backup. <ul style="list-style-type: none"> • Check VRRP Master for uptime. • Check VRRP Master for connectivity.
Threshold Crossing Alerts (TCAs)	00050001	WARNING	The average WAN-side transmit throughput of X Mbps over the last minute [exceeded, fell below] the threshold of Y Mbps RESOLUTION: User configured. Check bandwidth reports for tunnel bandwidth.
	00050002	WARNING	The average LAN-side receive throughput of X Mbps over the last minute [exceeded, fell below] the threshold of Y Mbps RESOLUTION: User configured. Check bandwidth reports.
	00050003	WARNING	The total number of X optimized flows at the end of the last minute [exceeded, fell below] the threshold of Y RESOLUTION: User configured. Check flow and real-time connection reports.
	00050004	WARNING	The total number of X flows at the end of the last minute [exceeded, fell below] the threshold of Y RESOLUTION: User configured. Check flow and real-time connection reports.
	00050005	WARNING	The file system utilization of X% at the end of the last minute [exceeded, fell below] the threshold of Y RESOLUTION: Contact Customer Service.
	00050006	WARNING	The peak latency of X during the last minute [exceeded, fell below] the threshold of Y RESOLUTION: User configured. <ul style="list-style-type: none"> • Check Latency Reports. If latency is too high, check routing between the appliances and QoS policy on upstream routers. • Check tunnel DSCP marking. If latency persists, contact ISP and Silver Peak support.

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Threshold Crossing Alerts (TCAs) (cont.)	00050007	WARNING	<p>The average pre-FEC loss of X% over the last minute [exceeded, fell below] the threshold of Y%</p> <p>RESOLUTION: User configured.</p> <ul style="list-style-type: none"> • Check Loss Reports. • Check for loss between Silver Peak appliances (interface counters on upstream routers). • Use network bandwidth measurement tools such as iperf to measure loss. • Contact ISP (Internet Service Provider).
	00050008	WARNING	<p>The average post-FEC loss of X% over the last minute [exceeded, fell below] the threshold of Y%</p> <p>RESOLUTION: User configured.</p> <ul style="list-style-type: none"> • Check Loss Reports. • Check for loss between Silver Peak appliances (interface counters on upstream routers). • Use network bandwidth measurement tools such as iperf to measure loss. • Enable/Adjust Silver Peak Forward Error Correction (FEC). • Contact ISP (Internet Service Provider).
	00050009	WARNING	<p>The average pre-POC out-of-order packets of X% over the last minute [exceeded, fell below] the threshold of Y%</p> <p>RESOLUTION: User configured.</p> <ul style="list-style-type: none"> • Check Out-of-Order Packets Reports. <p>Normal in a network with multiple paths and different QoS queues.</p> <p>Normal in a dual-homed router or 4-port in-line [bridge] configuration.</p> <ul style="list-style-type: none"> • Contact Customer Service if out-of-order packets are not 100% corrected.
	0005000a	WARNING	<p>The average post-POC out-of-order packets of X% over the last minute [exceeded, fell below] the threshold of Y%</p> <p>RESOLUTION: User configured.</p> <ul style="list-style-type: none"> • Check Out-of-Order Packets Reports. <p>Normal in a network with multiple paths and different QoS queues.</p> <p>Normal in a dual-homed router or 4-port in-line [bridge] configuration.</p> <ul style="list-style-type: none"> • Contact Customer Service if out-of-order packets are not 100% corrected.
	0005000b	WARNING	<p>The average tunnel utilization of X% over the last minute [exceeded, fell below] the threshold of Y%</p> <p>RESOLUTION: User configured.</p> <p>Check bandwidth reports for tunnel bandwidth utilization.</p>

Table 6-1 Silver Peak Appliance Alarms (Continued)

Subsystem	Alarm Type ID (Hex)	Alarm Severity	Alarm Text
Threshold Crossing Alerts (TCAs) (cont.)	0005000c	WARNING	The average tunnel reduction of X% over the last minute [exceeded, fell below] the threshold of Y% RESOLUTION: User configured. <ul style="list-style-type: none"> • Check bandwidth reports for deduplication. • Check if the traffic is pre-compressed or encrypted.
	0005000d	WARNING	The total number of flows <num-of-flows> is approaching the capacity of this appliance. Once the capacity is exceeded, new flows will be <dropped bypassed>. RESOLUTION: If this condition persists, a larger appliance will be necessary to fully optimize all flows.

- a. The VX appliances are a family of virtual appliances, comprised of the VX-n000 software, an appropriately paired hypervisor and server, and a valid software license.
- b. If there is either a **LAN Next-Hop Unreachable** or **WAN Next-Hop Unreachable** alarm, resolve the alarm(s) immediately by configuring the gateway(s) to respond to ICMP pings from the Silver Peak appliance IP Address.

Viewing Alarms

Monitoring > Alarms

Orchestrator and appliance alarms display in the same table.

Host Name	Alarm Time	Cleared Time	Severity	Source	Alarm Description	Recommended Action	Ack
Orchestrator	03-Mar-16 10:08		Warning		Disk partition /home is more than 53...		<input type="checkbox"/>
Orchestrator	03-Mar-16 09:53	03-Mar-16 11:08	Minor		Backup configuration not set		<input type="checkbox"/>
Portland-EC	03-Mar-16 09:49	03-Mar-16 09:55	Critical	system	Appliance cannot connect to Orchest...	Appliance cannot connect to Orchestrator using HTTPS. This connectivity is required fo...	<input type="checkbox"/>
Los-Angeles	03-Mar-16 09:14	03-Mar-16 09:21	Critical	system	Appliance cannot connect to Orchest...	Appliance cannot connect to Orchestrator using HTTPS. This connectivity is required fo...	<input type="checkbox"/>
Portland-EC	03-Mar-16 09:14	03-Mar-16 09:14	Critical	orchestrator	Orchestrator cannot reach this appli...		<input type="checkbox"/>
Minneapolis	03-Mar-16 09:14	03-Mar-16 09:14	Critical	orchestrator	Orchestrator cannot reach this appli...		<input type="checkbox"/>
Los-Angeles	03-Mar-16 09:14	03-Mar-16 09:14	Critical	orchestrator	Orchestrator cannot reach this appli...		<input type="checkbox"/>
Boston	03-Mar-16 09:14	03-Mar-16 09:14	Critical	orchestrator	Orchestrator cannot reach this appli...		<input type="checkbox"/>
Los-Angeles	03-Mar-16 09:13	03-Mar-16 09:15	Critical	to_Minneapolis_MPLS-...	Tunnel state is Down	Tunnel peer is unreachable. Check tunnel configuration. Verify Local & Remote IPs, Ad...	<input type="checkbox"/>
Salt-Lake-Ct...	03-Mar-16 09:13	03-Mar-16 09:14	Critical	orchestrator	Orchestrator cannot reach this appli...		<input type="checkbox"/>
Boston	03-Mar-16 09:13	03-Mar-16 09:15	Critical	to_Salt-Lake-City-EC_E...	Tunnel state is Down	Tunnel peer is unreachable. Check tunnel configuration. Verify Local & Remote IPs, Ad...	<input type="checkbox"/>
Minneapolis	03-Mar-16 09:13	03-Mar-16 09:15	Critical	to_Salt-Lake-City-EC_E...	Tunnel state is Down	Tunnel peer is unreachable. Check tunnel configuration. Verify Local & Remote IPs, Ad...	<input type="checkbox"/>
			Critical	to_Minneapolis_MPLS-...	Tunnel state is Down	Tunnel peer is unreachable. Check tunnel configuration. Verify Local & Remote IPs, Ad...	<input type="checkbox"/>

- The table has three filters:
 - **Active** - all uncleared alarms. Acknowledged alarms go to the bottom of this list.
 - **History** - filtered to show only cleared alarms
 - **All** - all uncleared and cleared alarms
- The Orchestrator keeps alarms for 90 days.
- Alarms have one of four severity levels: **Critical**, **Major**, **Minor**, or **Warning**. Only **Critical** and **Major** alarms are service-affecting.
 - **Critical** alarms require immediate attention, and reflect conditions that affect an appliance or the loss of a broad category of service.
 - **Major** alarms reflect conditions which should be addressed in the next 24 hours -- for example, an unexpected traffic class error.
 - **Minor** alarms can be addressed at your convenience -- for example, a degraded disk.
 - **Warnings** inform you of conditions that may become problems over time -- for example, the network interface is admin down.

Additional alarm indications

- A cumulative (Orchestrator + appliances) alarm summary always displays at the right side of the header. Clicking it opens the **Alarm** tab.

Alarms **31 Critical** **856 Major** **4 Minor** **2 Warning**

- On the **Topology** tab, appliances color-code how many of their severest alarms are open.

Specifying Alarm Recipients

The Orchestrator sends out alarm notifications as soon as they're received, to the email addresses specified.

A description area (not used as a filter)

Refers to the **appliance groups** in the navigation area

The screenshot shows the 'Alarm Recipients' configuration window. At the top left is an 'Add Recipient' button. Below it is a search bar. The main area is a table with the following structure:

Configuration Name ▲	Alarm Type	Severity				Email Addresses	Groups	
		Critical	Major	Minor	Warning			
main alarm email alert	Appliance ▼	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		Silver Peak Systems	×
	Orchestrator	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	dmerwin@silver-peak.com	Silver Peak Systems	×

At the bottom right of the window are buttons for 'Save', 'Reload', and 'Close'.



Monitoring Status and Performance

This chapter focuses on reports related to performance, traffic, and appliance status.

Also helpful in monitoring, *Alarms* and *Threshold Crossing Alerts* are addressed in other chapters.

In This Chapter

- **About Reports** See page 166.
- **Configuring and Distributing Custom Reports** See page 168.
- **Viewing Appliance Statistics** See page 170.
- **Viewing Application Statistics** See page 176.
- **Viewing Tunnel Statistics** See page 180.
- **Viewing Flows** See page 188.
- **Monitoring Status & Reporting** See page 206.

About Reports

This section discusses types of reports and understanding traffic direction.

Types of Reports

Reports and statistics help you bracket a problem, question, or analysis. The Orchestrator's collections of reports basically fall into two broad categories:

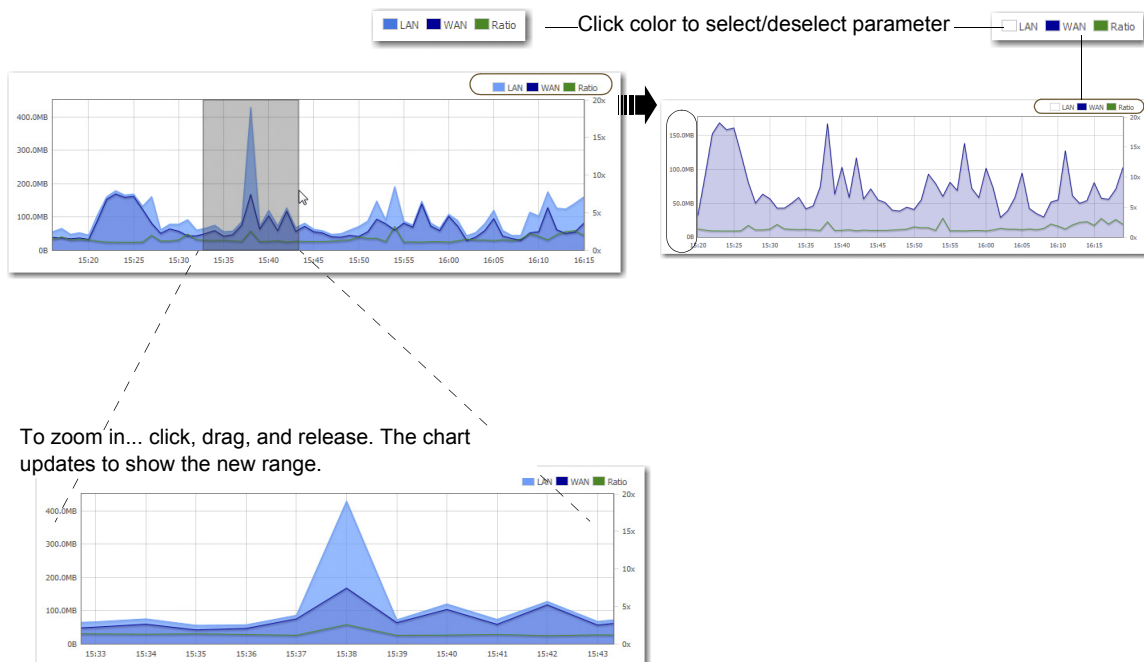
- Statistics related to **network performance** and **application performance**. These provide visibility into the network, enabling you to investigate problems, and address trends, and evaluate your WAN utilization.
- Reports related to **status** of the network and appliances. For example, alarms, threshold crossing alerts, reachability between the Orchestrator and appliances, scheduled jobs, etc.

Interpreting Charts

Some charts feature spark lines, as well as selectable (and modifiable) time ranges for collected data. Others show data in a bar chart format.

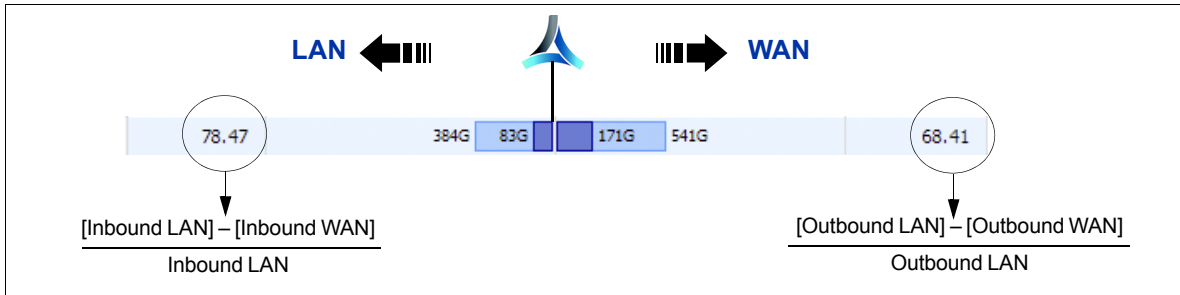
Line Charts

Line charts consist of filters, chart displays, and a modifiable time range area.



Bar Charts

For each direction of traffic — inbound and outbound — the overlapping bars are paired to show the full volume of traffic and the reduced, optimized size of the same traffic.



Configuring and Distributing Custom Reports

Monitoring > [Status & Reporting] Schedule & Run Reports

Use the **Schedule & Run Reports** tab to create, configure, run, schedule, and distribute reports.

Prepopulated from email settings in
Orchestrator Administration > Getting Started Wizard

Topology
Schedule & Run Reports ×

Schedule & Run Reports ? [View Reports](#)

for Global Report

Name
Global Report ▼

New Report Delete Report

Email Recipients [email image sizes](#)

eval-support@silver-peak.com, fakeemail@fake.com
(separate with commas or semicolons)

Appliances in Report Use Tree Selection

All Appliances

Data Granularity - Time Range

Daily - 14 days

Hourly - 24 hours

Minutely - 240 minutes

Scheduled or Single Report

Run Scheduled Report

Every day at 0:30 starting 03-Jun-14 0:30 PDT Edit

Run Single Report with Custom Time Range

2016-02-12 14:25 - 2016-02-19 14:25

Top Traffic Type

10 ▼ Optimized Traffic ▼

Application Charts

Application Reduction ↗

Application Pie Chart ↗

Application Trends ↗

Filter App

Tunnel Charts

Health Dashboard ↗

Flows Count ↗

Bandwidth Utilization ↗

Packets Count ↗

Loss ↗

Loss By Time ↗

Out-Of-Order Packets ↗

Out-Of-Order By Time ↗

Latency ↗

Latency By Time ↗

Tunnels Max Bandwidth ↗

Data Transfer & Reduction ↗

Tunnel Throughput By Time ↗

Filter Tunnel

Appliance Charts

Bandwidth Cost Savings ↗

Data Transfer & Reduction ↗

Bandwidth Utilization ↗

Appliance Throughput By Time ↗

Appliance Max Bandwidth ↗

Appliance Flow Count ↗

Appliance Packets Count ↗

DSCP Reduction ↗

Traffic Class Reduction ↗

Save Cancel Run Now

- On schedule or on demand, the Orchestrator can generate Daily, Hourly, and/or Minute Reports containing user-selected charts.
- Each report is a separate PDF file, and takes its filename based on the date, time, granularity, and name of the generated report.
- Along with the PDF report(s), the Orchestrator also generates a corresponding .zip file containing the raw data in .csv files. To open the .zip file, use either Winrar or 7-Zip.
- To access all reports residing on the Orchestrator server, click **View Reports**. The Orchestrator retains reports and zipped .csv files for 30 days.



Tip To specify the timezone for scheduled jobs and reports, go to **Orchestrator Administration > [General] Schedule Timezone**.

- The Orchestrator server also sends **reports via email**, using a Silver Peak SMTP server in Amazon Web Services.
 - To send a test email and/or to configure another SMTP server instead, click **SMTP server settings**.
 - If a test email doesn't arrive within minutes, check your firewall.
- **Global Report** - Once you enable it, this preconfigured subset of charts runs at 00:30 each day. This allows time to complete end-of-day processing. You can modify which charts to include and when/whether to run the report, but you cannot delete it.

Data Collection & Management

- The Orchestrator **polls** each of the appliances at **15-minute intervals**, based on the time that the Orchestrator was powered on. So, if the Orchestrator powered on at 14:26, it polls at 14:41, 14:56, 15:11, and 15:26, etc.
- A day begins at 00:00 and ends at 23:59:59.
- A Daily or Hourly report begins at the top of the hour. A Minute report begins at the last poll period.
- Report stats aggregate to 1 minute.
- Reach of reports: **Daily** = 14 days, **Hourly** = 24 hours, **Minute** = 4 hours
- In charts, the Orchestrator displays only the maximum peak in each prescribed time interval.
- Reports return the top ten filtered or unfiltered items.

Viewing Appliance Statistics

Charts feature spark lines, as well as selectable (and modifiable) time ranges.

The following charts exist for monitoring appliances:

- **Health Dashboard** See page 171.
- **Appliance Data Transfer & Reduction** See page 172.
- **Appliance Max Bandwidth** See page 172.
- **Appliance Bandwidth Utilization** See page 173.
- **Appliance Bandwidth Trends** See page 173.
- **Appliance Bandwidth Cost Savings** See page 174.
- **Appliance Flow Count** See page 175.
- **Appliance Packet Count** See page 175.

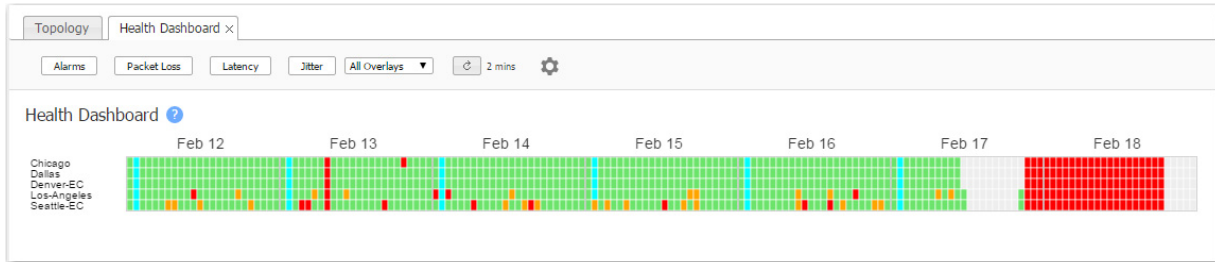
These items are also grouped with appliance statistics in the **Monitoring** menu:


- The **Flows** tables is discussed in more detail, separately, in *“Viewing Flows” on page 188.*
- **Alarms** are addressed in *Chapter 6, “Alarms.”*

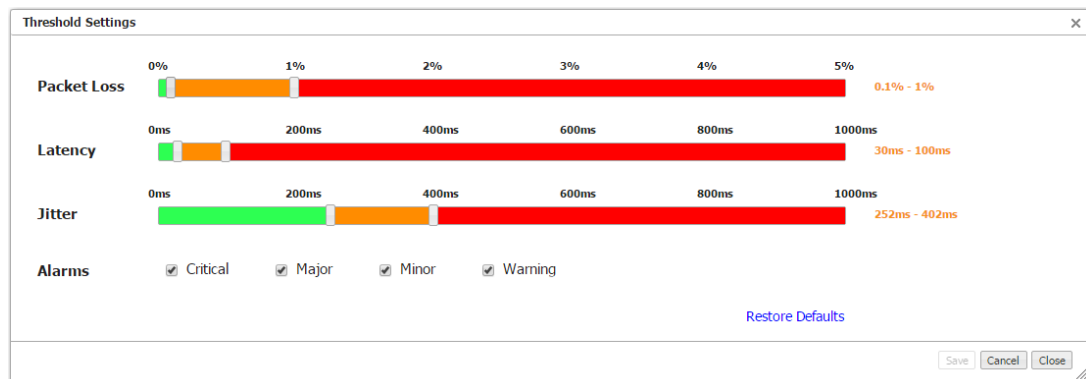
Health Dashboard

Monitoring > [Appliances] Health Dashboard

The **Health Dashboard** provides a high-level view of your network's health, based on the filter thresholds you configure.



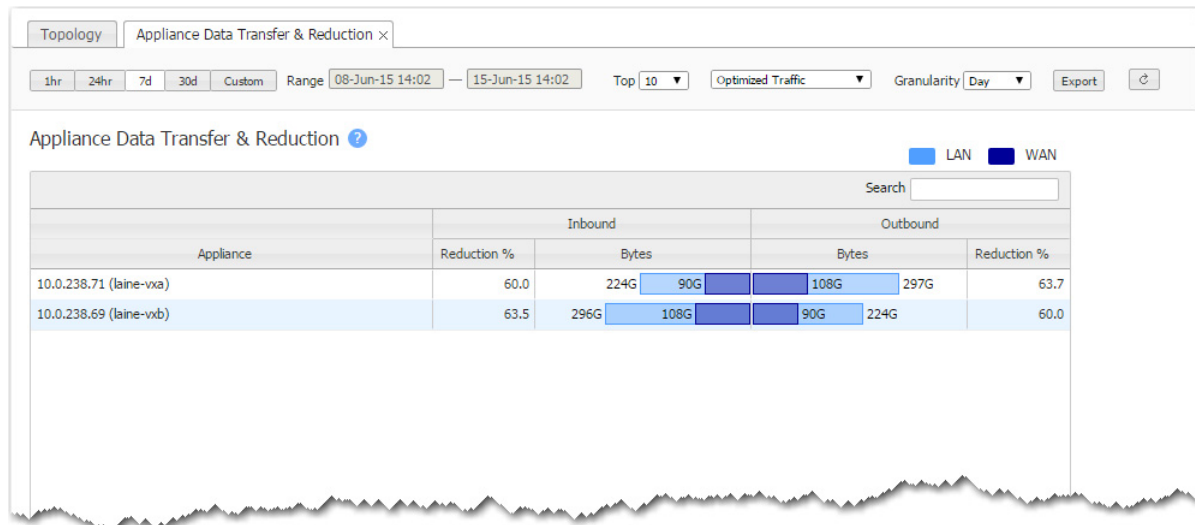
- Filters are available for *packet loss*, *latency*, and *jitter*. For each, you can configure two thresholds. You can also filter for various levels of *alarms*.
- Each block represents one hour and uses color coding to display the most severe event among the selected filters. Clicking a block displays a pop-up with specifics about that event, what value triggered it, and any additional threshold breach for that appliance during the same hour.
 - **Green** = normal operation
 - **Orange** = marginal
 - **Red** = needs immediate attention
 - **Aqua** = warning (an alarm level)
 - **Grey** = no data available
- Threshold settings apply globally. They are not retroactive; in other words, setting new thresholds does not redisplay historical data based on newly edited values.
- Deleting an appliance deletes its data.
- If you are using overlays...
 - You can view each overlay's health individually.
 - If you remove an individual overlay, its individual data is not recoverable. However, its historical data remains included in **All Overlays**.
- To access threshold configuration, click the gear icon ().



Appliance Data Transfer & Reduction

Monitoring > [Appliances] Data Transfer & Reduction

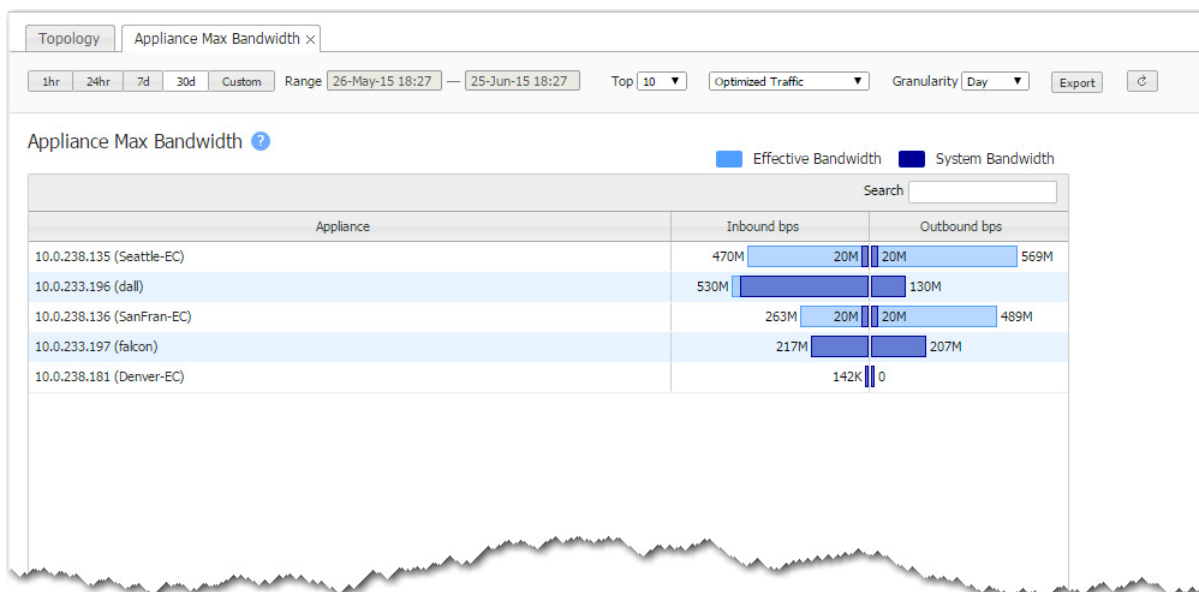
The **Appliance Data Transfer & Reduction** chart lists the top appliances based on the total volume of inbound and outbound traffic before reduction. It shows how many bytes the Silver Peak appliance saved when transferring data, aggregated over a selectable time period.



Appliance Max Bandwidth

Monitoring > [Appliances] Max Bandwidth

The **Appliance Max Bandwidth** chart lists the top appliances by the peak throughput (in either direction), within a selected time period. It compares the system bandwidth of the appliance to the effective bandwidth it's providing.

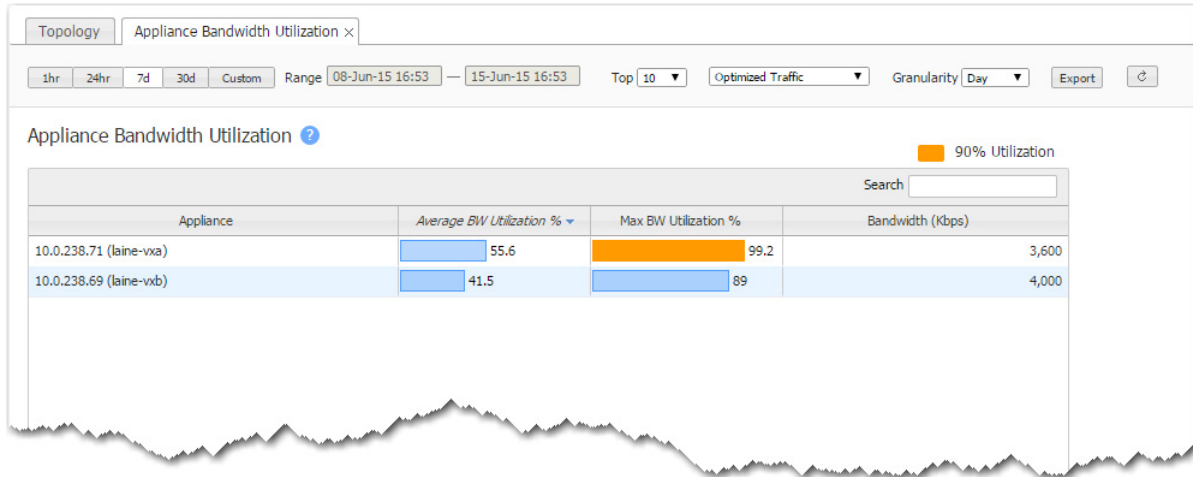


Appliance Bandwidth Utilization

Monitoring > [Appliances] Bandwidth Utilization

The **Appliance Bandwidth Utilization** chart lists the top appliances by the average percent of available bandwidth used. This helps you see if an appliance that is optimizing traffic is reaching its capacity.

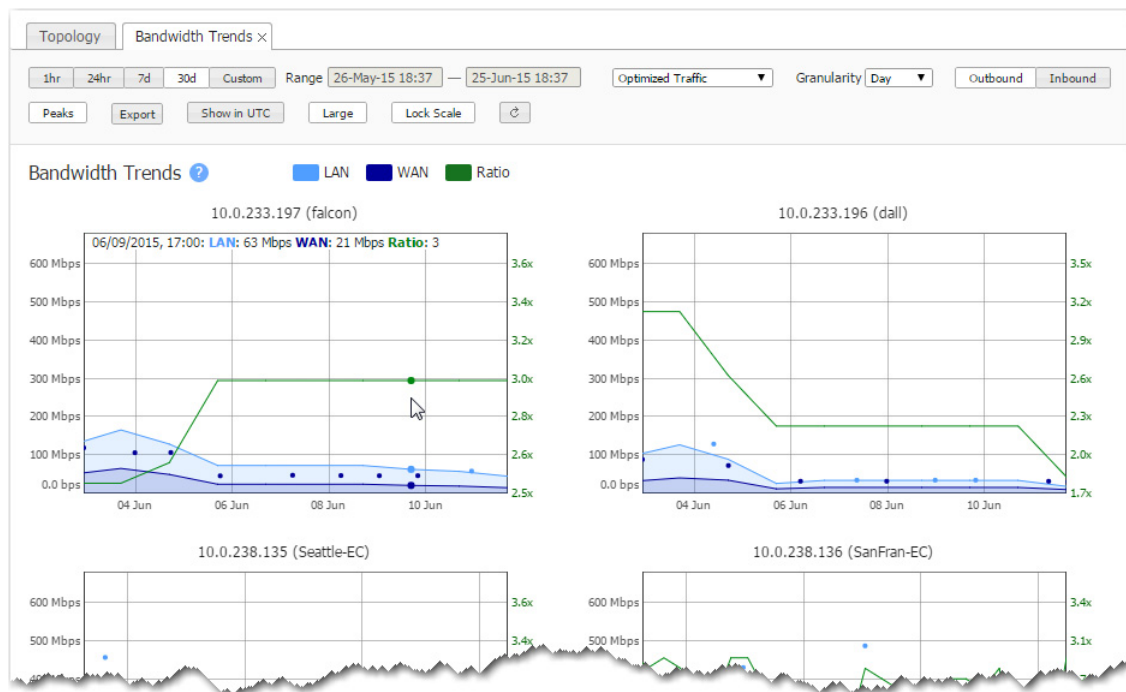
To see if your data link is nearing capacity, refer to the **Tunnel Bandwidth Utilization** chart.



Appliance Bandwidth Trends

Monitoring > [Appliances] Bandwidth Trends

The **Appliance Bandwidth Trends** chart shows bandwidth usage over time.



Appliance Bandwidth Cost Savings

Monitoring > [Appliances] Bandwidth Cost Savings

The **Bandwidth Cost Savings** chart shows how much money and time the Silver Peak appliances could have saved based on reduced bandwidth usage. The monthly figures are calculated by extrapolating the savings from the selected time range.

To view or edit the various data link costs, click **Configure Cost**.

The screenshot displays the 'Bandwidth Cost Savings' interface. At the top, there are filters for time range (1hr, 24hr, 7d, 30d, Custom) and a selected range from 15-Jun-15 16:08 to 15-Jun-15 17:08. The interface shows 'Monthly Cost Saved \$17,595' and 'Monthly Hours Saved 2,754'. A table lists two appliances with their respective cost and time savings. A 'Costs Configuration' dialog box is open, showing a list of bandwidth options and their corresponding monthly costs.

Appliance Name	Monthly Cost Saved	Monthly Hours Sav...	Bandwidth (Kbps)	Reduction %	Inbound Bytes	Outbound Bytes	Reduction %
10.0.238.71 (laine...)	\$9,151	1,432	3,600	50.6	2.1G	1.1G	50.3
10.0.238.69 (laine...)	\$8,444	1,322	4,000	50.5	2.2G	1.1G	50.5

Bandwidth	Cost Per Month	ISP	
1.6 Mbps	\$750.00		X
2.1 Mbps	\$1,000.00		X
3.2 Mbps	\$1,500.00		X
10.2 Mbps	\$4,600.00		X
20.5 Mbps	\$9,200.00		X
45.8 Mbps	\$23,000.00		X
52.7 Mbps	\$27,000.00		X
102.4 Mbps	\$51,000.00		X
158.7 Mbps	\$95,000.00		X
636.9 Mbps	\$382,000.00		X
1.0 Gbps	\$630,000.00		X
2.5 Gbps	\$1,500,000.00		X
10.1 Gbps	\$15,000,000.00		X

Calculations

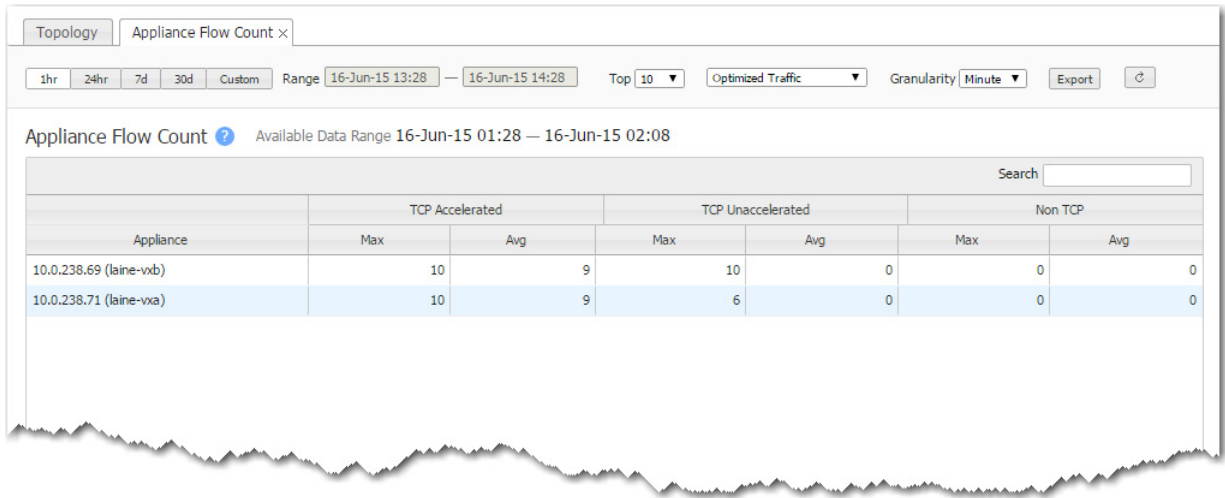
- For the monthly cost savings, it subtracts the maximum bytes you could send **without** the appliance from how many bytes the Silver Peak **actually** sent, and multiplies the difference by the data link cost.
- For the monthly time savings, it uses the data link speed to calculate how many more hours it would have taken to send those additional bytes **without** the Silver Peak appliance.

Appliance Flow Count

Monitoring > [Appliances] Flow Count

The **Appliance Flow Count** chart lists the top appliances according to which ones had the most flows within a selected time period.

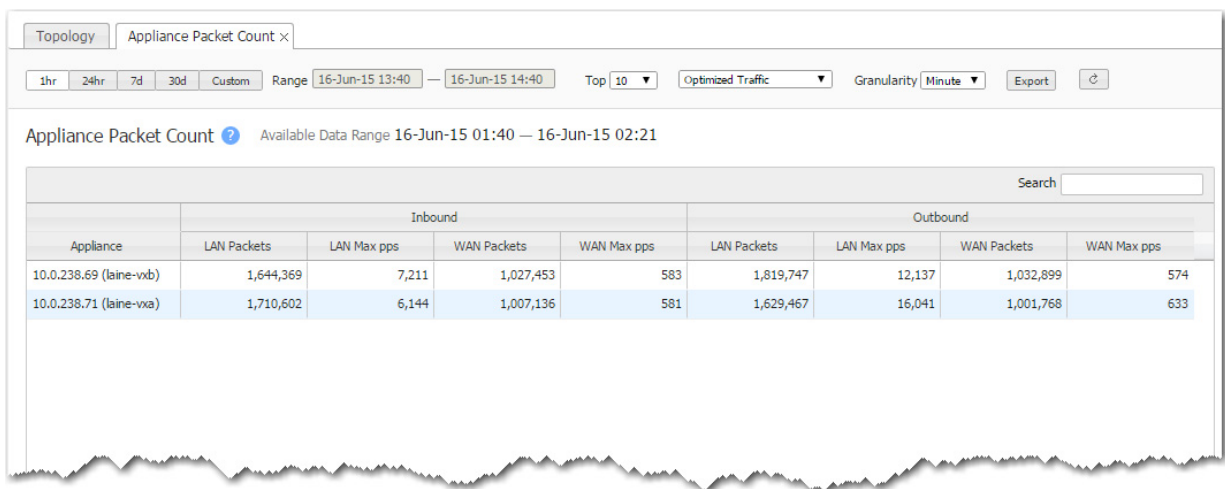
When you filter on **All Traffic**, the **Created** and **Deleted** columns display the number of new and ended flows for that same time period. The **Max** column value is from a one-minute window within the time range.



Appliance Packet Count

Monitoring > [Appliances] Packet Count

The **Appliance Packet Count** chart lists the top appliances according to the sum of the inbound and outbound LAN packets, showing how much traffic was sent.



Viewing Application Statistics

Charts feature spark lines, as well as selectable (and modifiable) time ranges for collected data.

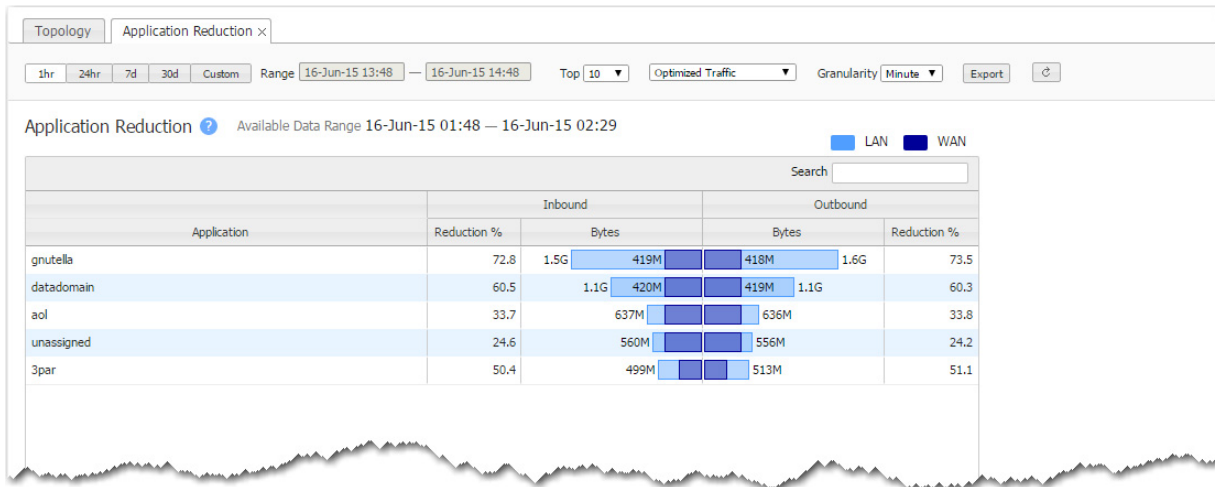
The following charts exist for monitoring applications:

- **Application Reduction** See page 176.
- **Application Pie Charts** See page 177.
- **Application Trends** See page 178.
- **DSCP Reduction** See page 179.
- **Traffic Class Reduction** See page 179.

Application Reduction

Monitoring > [Applications] Application Reduction

The **Application Reduction** chart shows which applications have sent the most bytes.

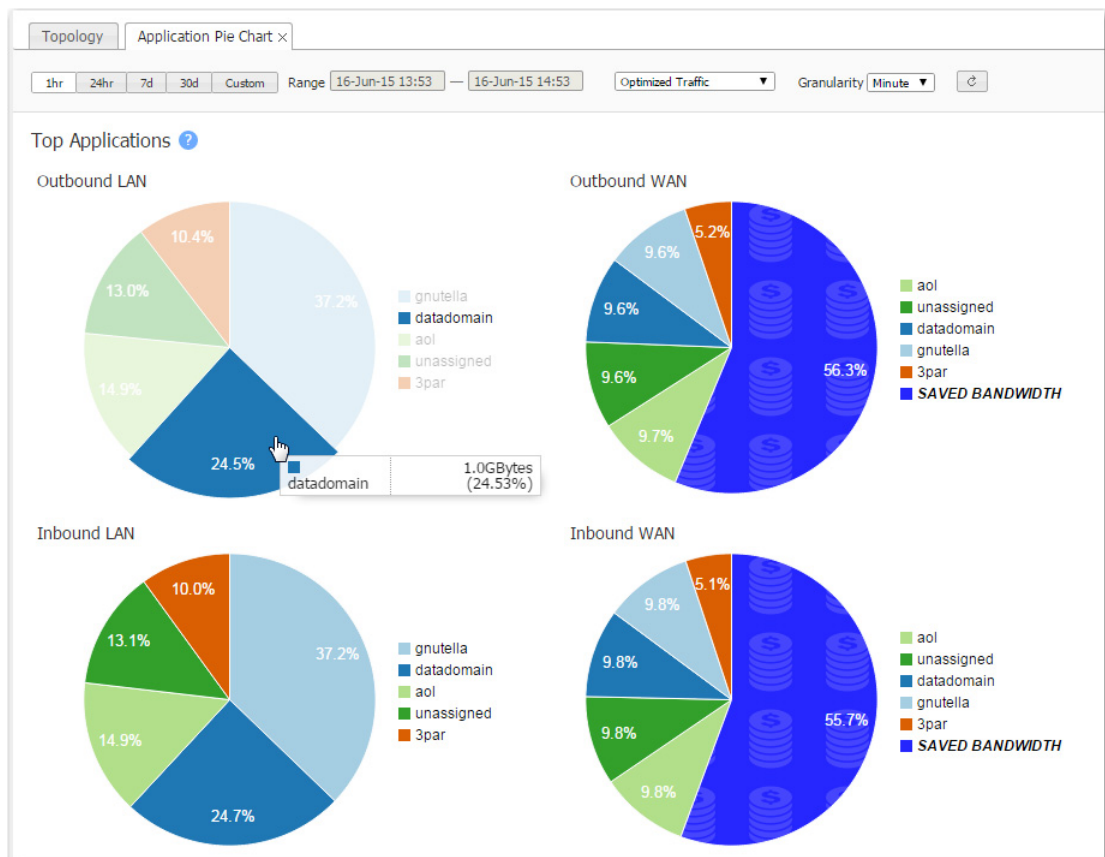


Application Pie Charts

Monitoring > [Applications] Application Pie Charts

The **Application Pie Charts** show what proportion of the bytes an application consumes on the LAN and on the WAN.

- Mousing over the charts and the legends reveals additional information.
- The WAN charts identify what percentage of the bandwidth the Silver Peak appliance saved by optimizing the traffic.



Application Trends

Monitoring > [Applications] Application Trends

The **Application Trends** chart answers the following questions:

- What proportion of traffic does each application account for over time?
- The top 10 applications account for what portion of the total traffic?



Note When it comes to flow and application statistics reports, user-defined applications are always checked before built-in applications.

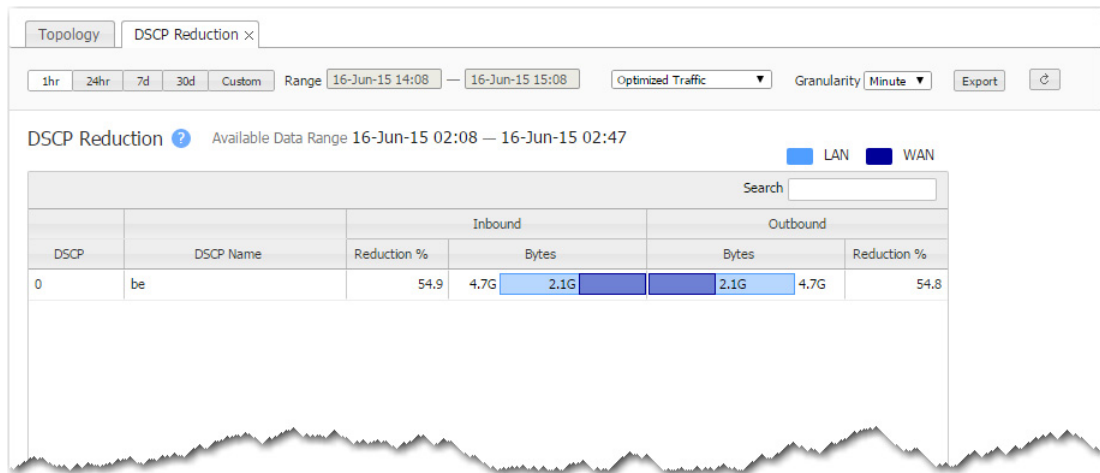
Ports are unique. If a port or a range includes a built-in port, then the custom application is the one that lays claim to it.

If two distinctly named user-defined applications have a port number in common, then report results will be skewed, depending on the priority assigned to the custom applications. A port is only counted once.

DSCP Reduction

Monitoring > [Applications] DSCP Reduction

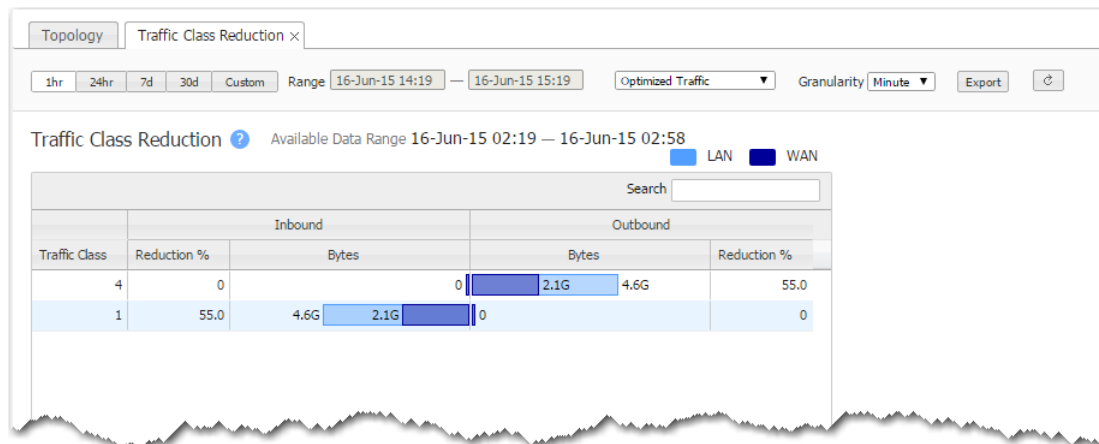
The **DSCP Reduction** chart shows which DSCP classes are sending the most data.



Traffic Class Reduction

Monitoring > [Applications] Traffic Class Reduction

The **Traffic Class Reduction** chart shows which applications have sent the most bytes.



Viewing Tunnel Statistics

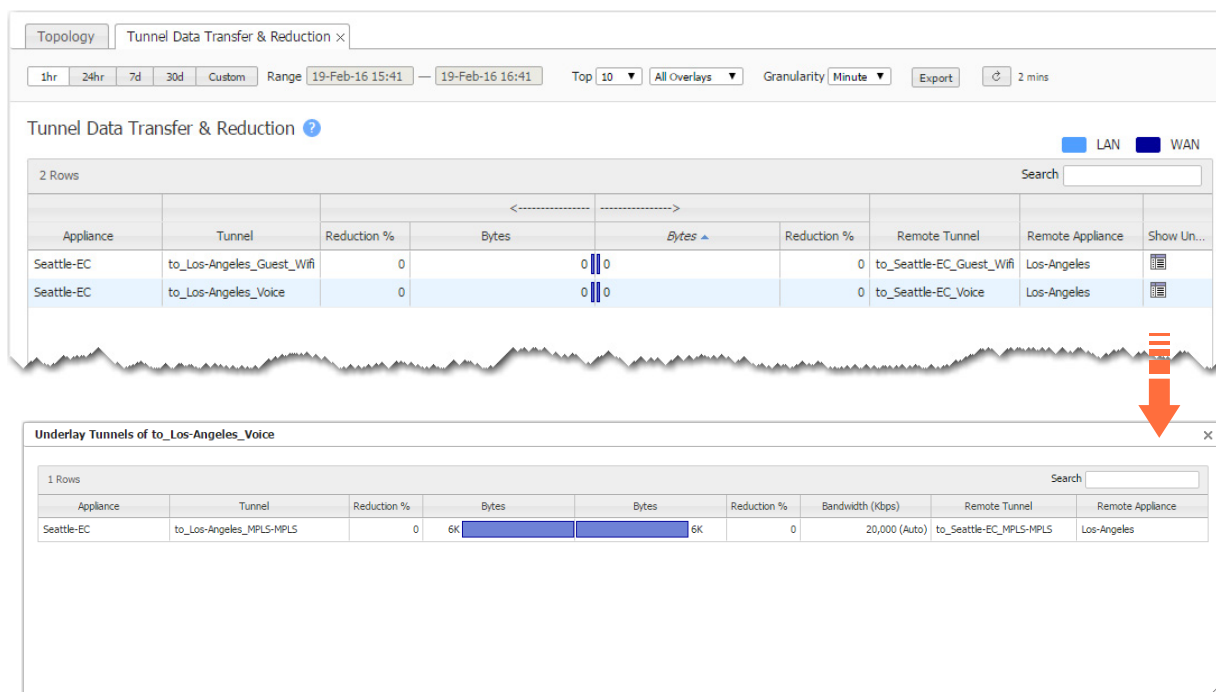
The following charts exist for monitoring tunnels:

- **Tunnel Data Transfer & Reduction** See page 180.
- **Tunnel Bandwidth Trends** See page 181.
- **Tunnel Max Bandwidth** See page 181.
- **Tunnel Bandwidth Utilization** See page 182.
- **Latency** See page 182.
- **Latency Trends** See page 183.
- **Loss** See page 183.
- **Loss Trends** See page 184.
- **Out of Order Packets** See page 184.
- **Out of Order Packets Trends** See page 185.
- **Tunnel Flow Count** See page 186.
- **Tunnel Packet Count** See page 186.
- **Tunnels Summary** See page 187.

Tunnel Data Transfer & Reduction

Monitoring > [Tunnels] Data Transfer & Reduction

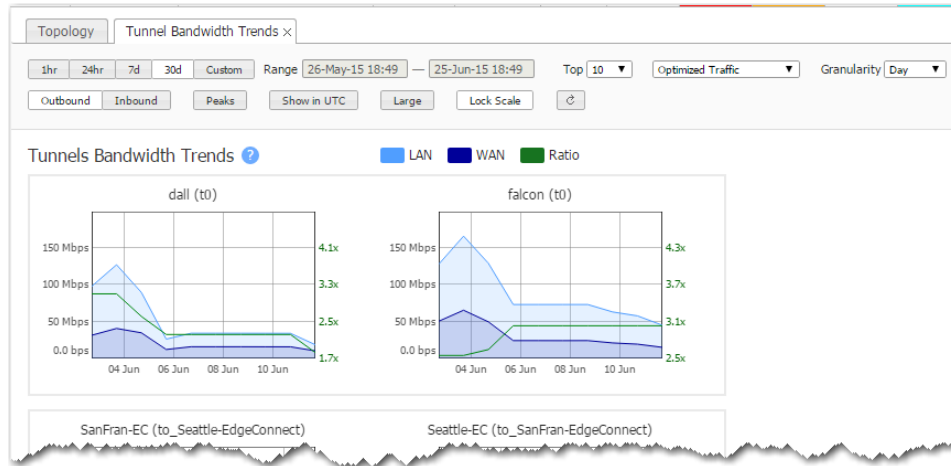
The **Tunnel Data Transfer & Reduction** chart shows which tunnels are sending the most bytes -- that is, the tunnels that are the most active.



Tunnel Bandwidth Trends

Monitoring > [Tunnels] Bandwidth Trends

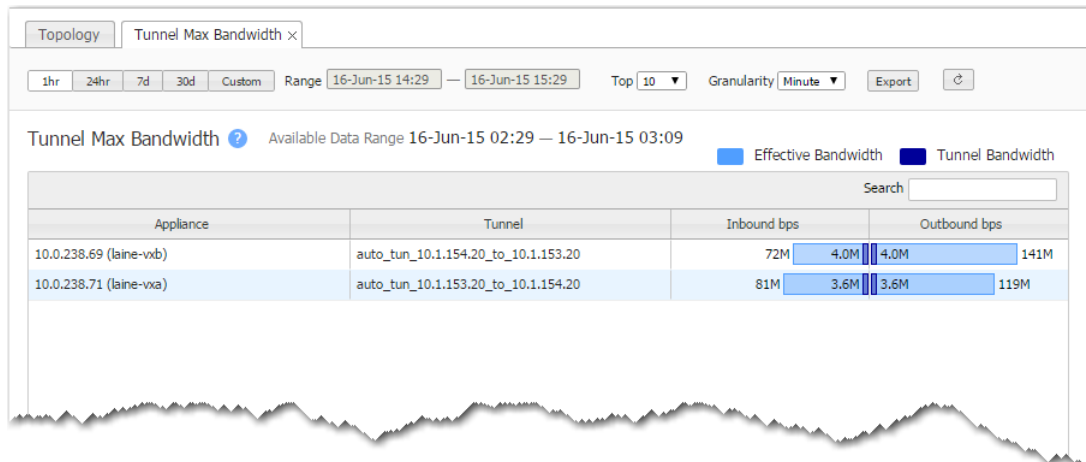
The **Tunnel Bandwidth Trends** chart shows tunnel bandwidth usage over time.



Tunnel Max Bandwidth

Monitoring > [Tunnels] Max Bandwidth

The **Tunnel Max Bandwidth** chart lists the top tunnels by the peak throughput (in either direction), within a selected time period. It shows how quickly data could have been sent through the tunnel.

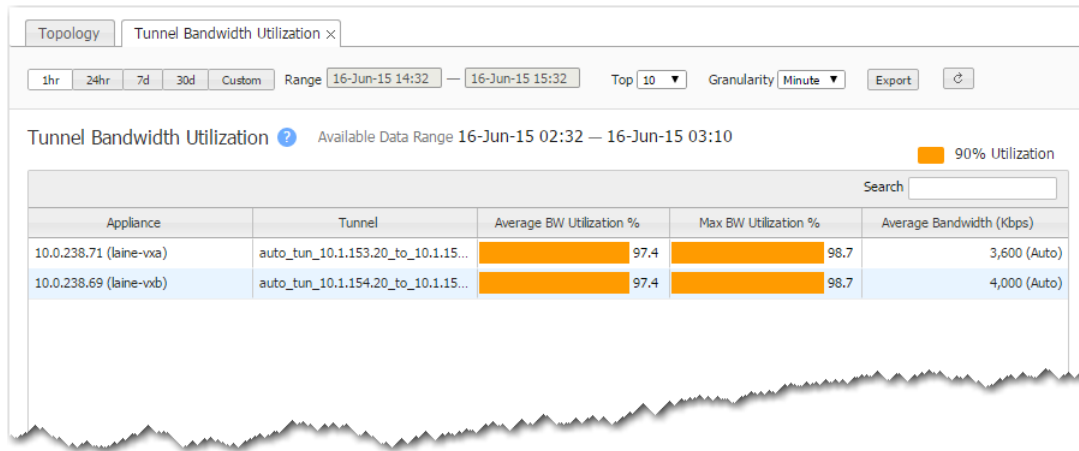


Tunnel Bandwidth Utilization

Monitoring > [Tunnels] Bandwidth Utilization

The **Tunnel Bandwidth Utilization** chart lists the top tunnels by the average percent of available bandwidth used. This helps you see if a data link is reaching its capacity.

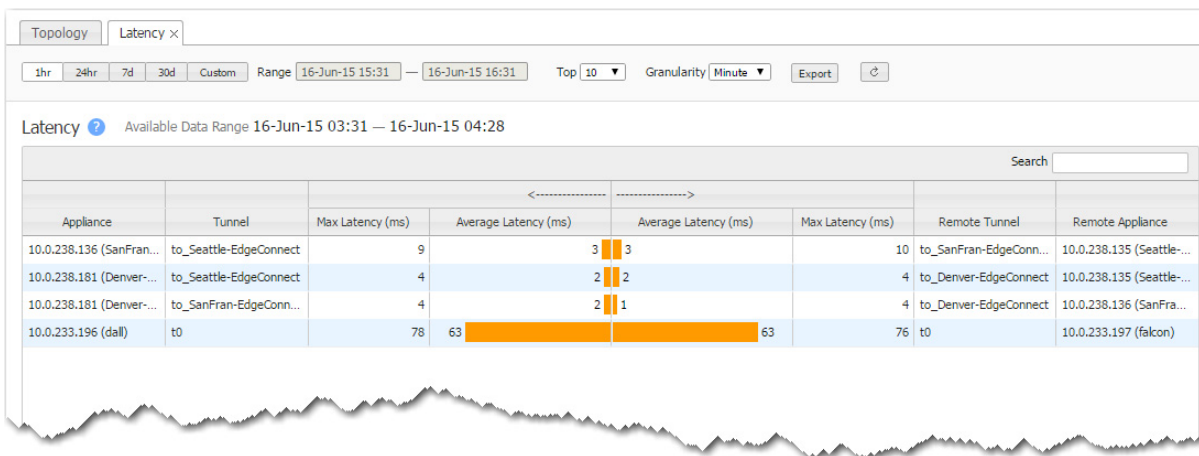
To see if your appliance is nearing capacity, refer to the **Appliance Bandwidth Utilization** chart.



Latency

Monitoring > [Tunnels] Latency

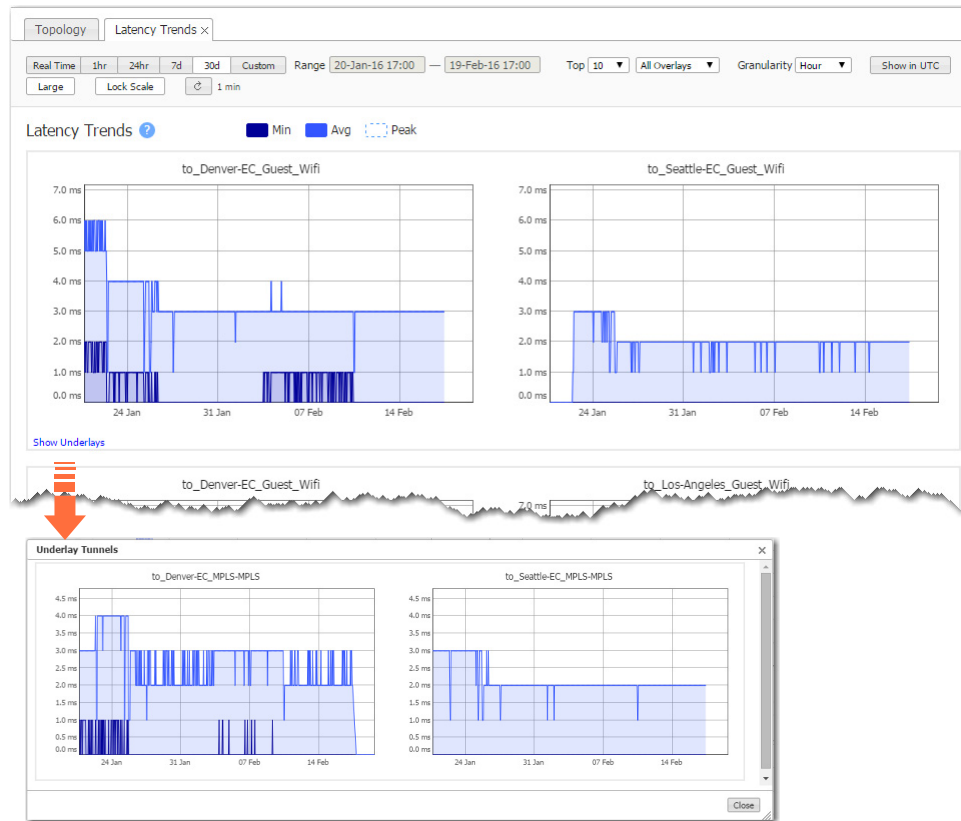
The **Latency** chart shows which tunnels have the most transmission delay, generally as a result of congestion.



Latency Trends

Monitoring > [Tunnels] Latency Trends

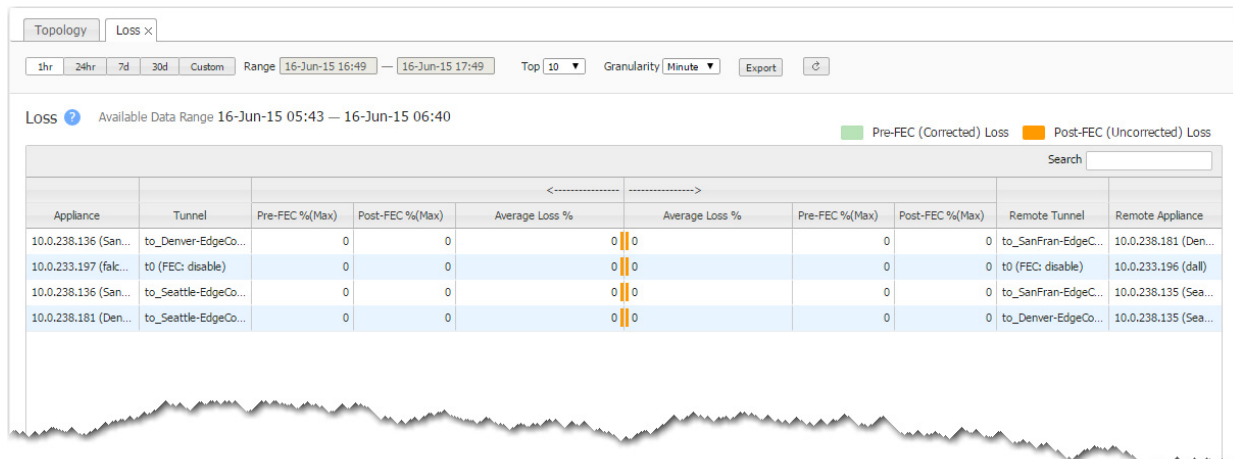
The **Latency Trends** chart shows tunnel latency over time.



Loss

Monitoring > [Tunnels] Loss

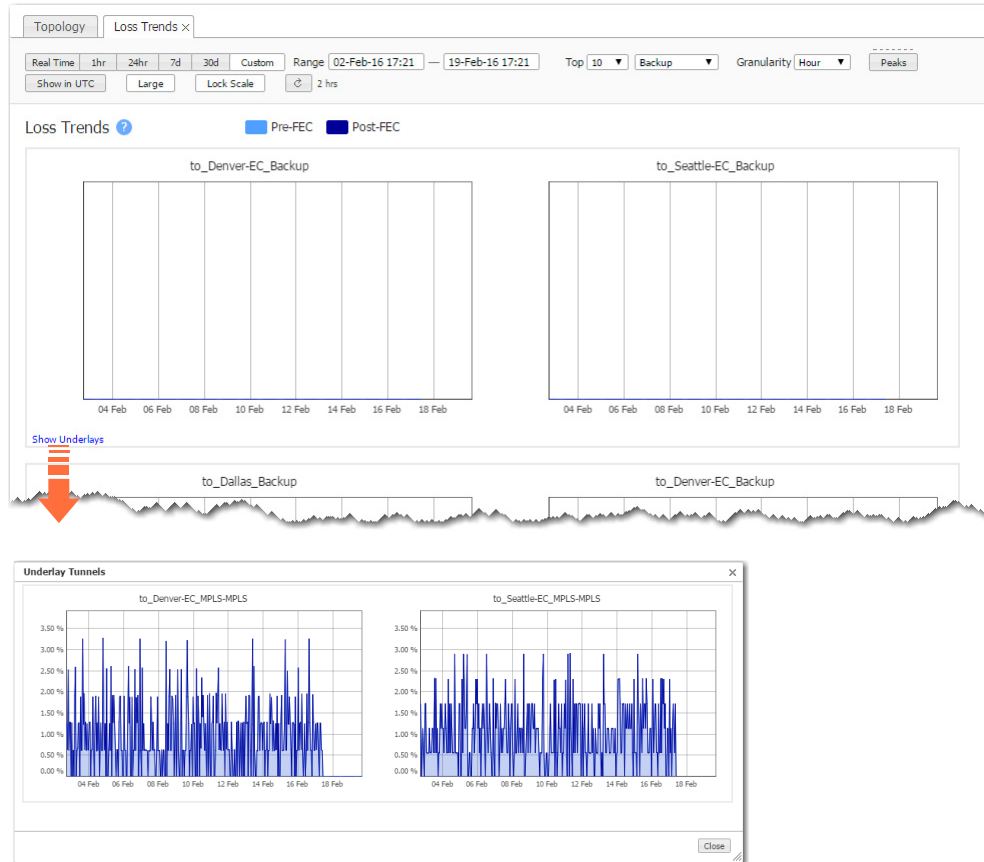
The **Loss** chart shows which tunnels have the most dropped packets.



Loss Trends

Monitoring > [Tunnels] Loss Trends

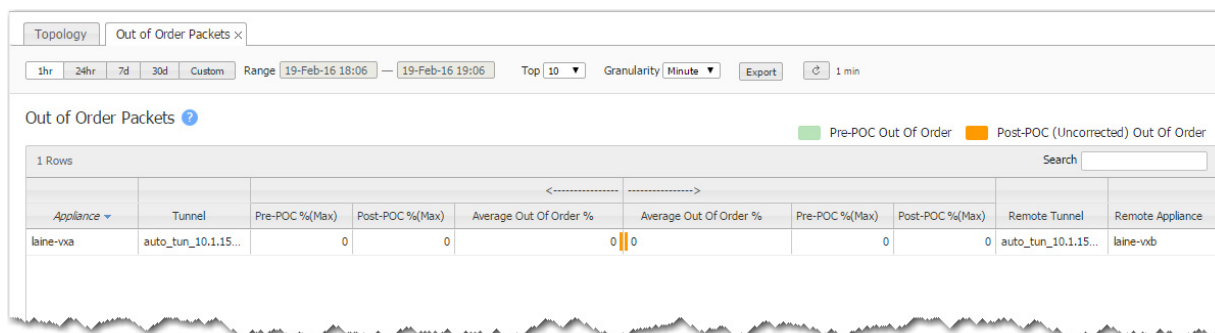
The **Loss Trends** chart shows tunnel packet loss over time, before and after Forward Error Correction (FEC).



Out of Order Packets

Monitoring > [Tunnels] Out of Order Packets

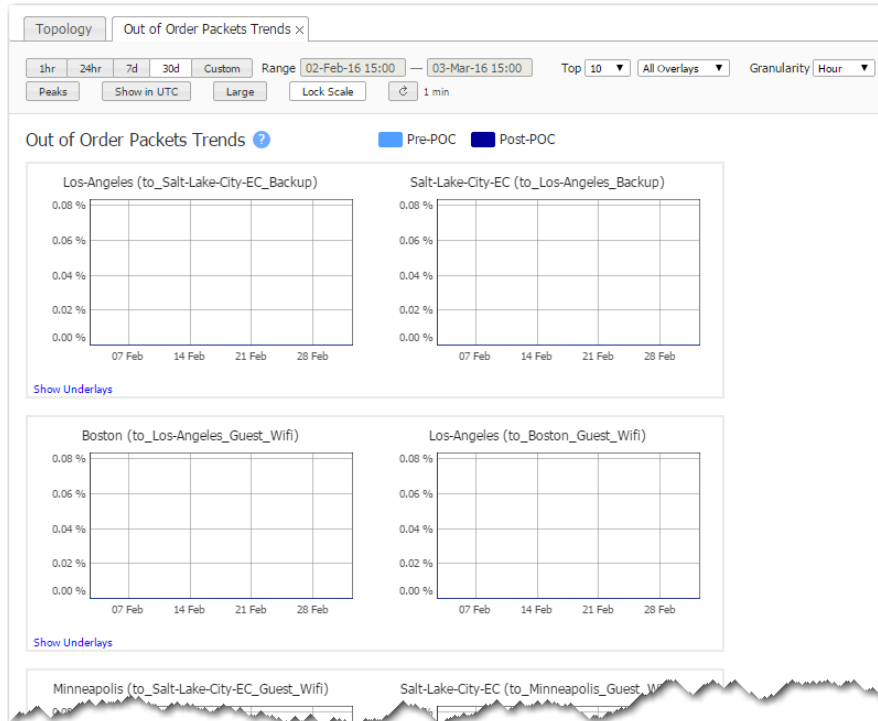
The **Max Out of Order Packets** chart shows which tunnels receive the most packets out of sequence relative to how they were sent.



Out of Order Packets Trends

Monitoring > [Tunnels] Out of Order Packets Trends

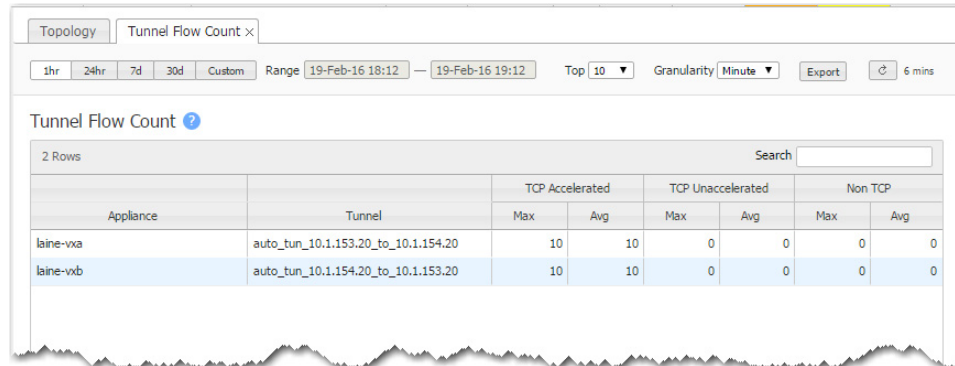
The **Out of Order Packets Trends** chart shows tunnel packets out of order over time, before and after Packet Order Correction (POC).



Tunnel Flow Count

Monitoring > [Tunnels] Flow Count

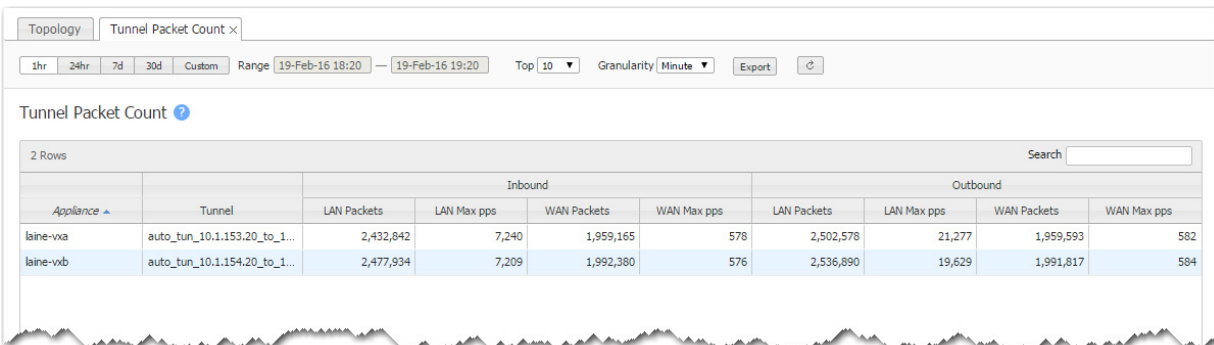
The **Tunnel Flow Count** chart lists the tunnels with the most flows, on average. It differentiates flows into TCP (accelerated and unaccelerated) and non-TCP, and also shows peak values.



Tunnel Packet Count

Monitoring > [Tunnels] Packet Count

The **Tunnel Packet Count** chart shows which tunnels sent the most packets.



Tunnels Summary

Monitoring > Tunnels Summary

The **Tunnels Summary** summarizes tunnel statistics — including reduction, throughput, latency, and packet loss.

Topology
Tunnels Summary ×

1hr 24hr 7d 30d Custom
Range 27-Jul-15 12:08 — 27-Jul-15 13:08
Top 10
Granularity Minute
Export
1 min

Tunnels Summary Available Data Range 27-Jul-15 12:07 — 27-Jul-15 13:07

Tunnel	Inbound					Outbound					Latency (ms)		Packets Loss %		Bandwidth
	LAN	WAN	Reduction %	LAN Through...	WAN Throug...	LAV	WAN	Reduction %	LAN Through...	WAN Throug...	Avg	Max	Avg	Max	
SanFran-EC: to_Seattle-E...	18 GB	7.4 GB	58.95	2.6 Gbps	1.0 Gbps	19 GB	7.5 GB	61.07	2.4 Gbps	990 Mbps	41.05	46.00	0	0	20,000 (Auto)
Seattle-EC: to_SanFran-E...	19 GB	7.6 GB	60.96	2.4 Gbps	985 Mbps	18 GB	7.4 GB	59.11	2.6 Gbps	1.0 Gbps	41.10	48.00	0	0	20,000 (Auto)
SanFran-EC: to_Denver-EC	0	0	0.00	0	0	0	0	0.00	0	0	20.03	23.00	0	0	20,000 (Auto)
Seattle-EC: to_Denver-EC	0	0	0.00	0	0	0	0	0.00	0	0	20.29	23.00	0	0	20,000 (Auto)
Denver-EC: to_SanFran-EC	0	0	0.00	0	0	0	0	0.00	0	0	19.49	23.00	0	0	20,000 (Auto)
Denver-EC: to_Seattle-EC	0	0	0.00	0	0	0	0	0.00	0	0	20.09	22.00	0	0	20,000 (Auto)

Viewing Flows

Monitoring > [Appliances] Flows

Flows are useful for troubleshooting and for detailed visibility into the network.

The **Flows** page retrieves a list of existing connections. The maximum visible number depends on which browser you use.

- The page displays a default set of columns, along with individual links to flow details and to any alerts.
- You can display additional columns from a customization list.

This section discusses the following topics:

- **How Flows Are Counted** See page 188.
- **How Flows are Organized** See page 189.
- **Customizing Which Columns Display** See page 191.
- **Flow Details** See page 192.
 - **Error Reasons for TCP Acceleration Failure** See page 196.
 - **Error Reasons for CIFS Acceleration Failure** See page 199.
 - **Error Reasons for SSL Acceleration Failure** See page 200.
 - **Error Reasons for Citrix Acceleration Failure** See page 203.
- **Resetting Flows to Improve Performance** See page 205.

How Flows Are Counted

When it comes to flow and application statistics reports, user-defined applications are always checked before built-in applications.

Ports are unique. If a port or a range includes a built-in port, then the custom application is the one that lays claim to it.

If two distinctly named user-defined applications have a port number in common, then report results will be skewed, depending on the priority assigned to the custom applications. A port is only counted once.

How Flows are Organized

Click to select the filter.
Active filters are highlighted.

Enter specific addresses and/or use zeroes (in the octet) as wildcards. The page lists flows that have either endpoint.


The screenshot shows the 'Flows' page in a network management system. At the top, there are filter fields for IP1, Port1, IP2, and Port2, all set to 0.0.0.0 and 0. Below these are dropdown menus for Application (All), Traffic (All), Protocol (All), and VLAN id. A 'Max Flows' field is set to 7000. A search icon and '40 min' are also visible.

Below the filters is a table of active flows. The table has columns for Host Name, Application, IP1, Port1, IP2, Port2, Detail, Flow Chart, Reduction %, Inbound Bytes, Outbound Bytes, Reduction %, Protocol, Outbound Tunnel, Start Time, and End Time. The first few rows show flows for applications like 'grutella' and 'datadomain'.

Clicking the icon in the **Details** column displays a detailed flow report.

The screenshot shows a 'Flow details' window for IP1: 10.1.153.10, Port1: 48410, IP2: 10.1.154.10, and Port2: 5191. The window is divided into several sections:

- Stats:** Outbound Ratio (1.33), Inbound Ratio (0.28), Outbound LAN bytes (2,371,764,679), Inbound LAN bytes (6,243,852), etc.
- Route:** Map Name (map1), Priority in Map (default), Configured Tx Action (pass-through), Rx Action (auto_tun_10.1.153.20_to_10.1.154.20), Tx Reason (Auto-opt), Application (aol), Protocol (top), Flow Direction (Outbound), Flow Redirected From (Auto-opt Status), Auto-opt Transit Node 1 (10.1.153.20), Auto-opt Transit Node 2 (10.1.154.20), Auto-opt Transit Node 3 (0.0.0.0), Auto-opt Transit Node 4 (0.0.0.0), LAN-side VLAN (None).
- Optimization:** Map Name (map1), Priority in Map (default), TCP Acceleration Configured (Yes), TCP Acceleration Status (Yes), TCP Acceleration Info, TCP Asymmetric (No), Proxy Remote Acceleration (No), CIFS Acceleration Configured (No), CIFS Acceleration Status (No), CIFS Acceleration Info, CIFS Server Side (No), CIFS SMB Signed (No), SRDF Acceleration Configured (No), SRDF Acceleration Status (No), SSL Acceleration Configured (No), SSL Acceleration Status (No), SSL Acceleration Reason, Citrix Acceleration Configured (No), Citrix Acceleration Status (No).



The following filters are available:

Parameter or Action	Definition
Flow Categories	The number after each option specifies how many flows fit the criteria <ul style="list-style-type: none"> • All – all flows • Optimized – optimized flows • Optimized* – these flows originally had a Status of Alert, and the user chose to no longer receive Alerts of the same type • Pass-through – includes shaped and unshaped traffic • Alert – notifies the user of any issue that might be inhibiting optimization, and offers a possible solution
Bytes Transferred	Choose from Total or Last 5 minutes .
Flow Timing	Choose from the following: <ul style="list-style-type: none"> • Active • Active + Ended Last 5 minutes • Started Last 5 minutes • Ended Last 5 minutes • Ended
Flows to Slow Devices	For debugging. A <i>slow device</i> is one that cannot tolerate having its connections accelerated. Generally, this occurs when the WAN side is congested, resulting in too much data on the LAN side. To protect the health of the appliance, you'll need to disable TCP acceleration in the Optimization Policy.
IP1 (2) / Port1 (2)	The IP address of an endpoint(s) that you want to use as a filter: <ul style="list-style-type: none"> • Entering a specific endpoint returns flows that have that endpoint. • Entering 0 in any IP address's octet position acts as a wild card for that position. 0 in the Port field is also a wild card. • The two IP address (and port) fields are independent of each other. In other words, you can filter on two separate endpoints.
Application	Select which standard or user-defined application (or application group) to use as a filter criteria. The default value is All .
Traffic	Select the type of traffic connections you want to retrieve: <ul style="list-style-type: none"> • All – all optimized and pass-through traffic. • Policy Drop – traffic with a Set Action of Drop in the Route Policy • Optimized Traffic – the sum of all optimized traffic. That is, all tunneled traffic. • Pass-through Shaped – all unoptimized, shaped traffic. • Pass-through Unshaped – all unoptimized, unshaped traffic. • [a named Tunnel] – that specific tunnel's optimized traffic.
Protocol	Select from the list. The default value is All .
VLAN Id	Enter only the integer value for the VLAN Id.
Internet Service	For sorting by domain, country, or city.
Max Flows	The upper limit depends on what browser you're using.
Reset Flows	Resetting the flow kills it and restarts it. It is service-affecting.
Reclassify Flows	Reclassifying the flow is not service-affecting. If a policy change makes a flow stale or inconsistent, then reclassifying makes a best-effort attempt to conform the flow to the change. If the flow can't be successfully "diverted" to this new policy, then an Alert asks if you want to Reset.

Customizing Which Columns Display

Following are some customization guidelines:

- The default set of columns includes the following:

Mgmt IP	Status	Protocol
Host Name	Detail	Outbound Tunnel
Application	Flow Chart	Start Time
IP1	Inbound Reduction %	End Time
PORT1	Inbound Bytes	
IP2	Outbound Bytes	
PORT2	Outbound Reduction %	

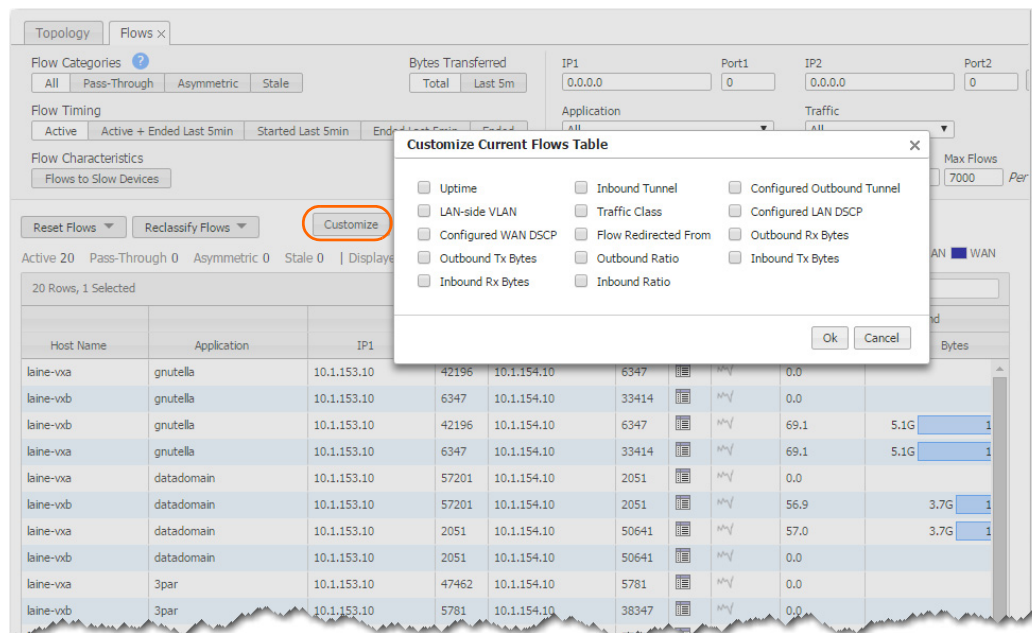
- You can customize by **adding** the following additional columns:

Uptime	Inbound Tunnel	Configured Outbound Tunnel
LAN-side VLAN	Traffic Class	Configured LAN DSCP
Configured WAN DSCP	Flow Redirected From	Outbound Rx Bytes
Outbound Tx Bytes	Outbound Ratio	Inbound Tx Bytes
Inbound Rx Bytes	Inbound Ratio	

- When you **Export** the data, **all default and possible custom columns are included** in the .csv file.

◆ To customize the screen display

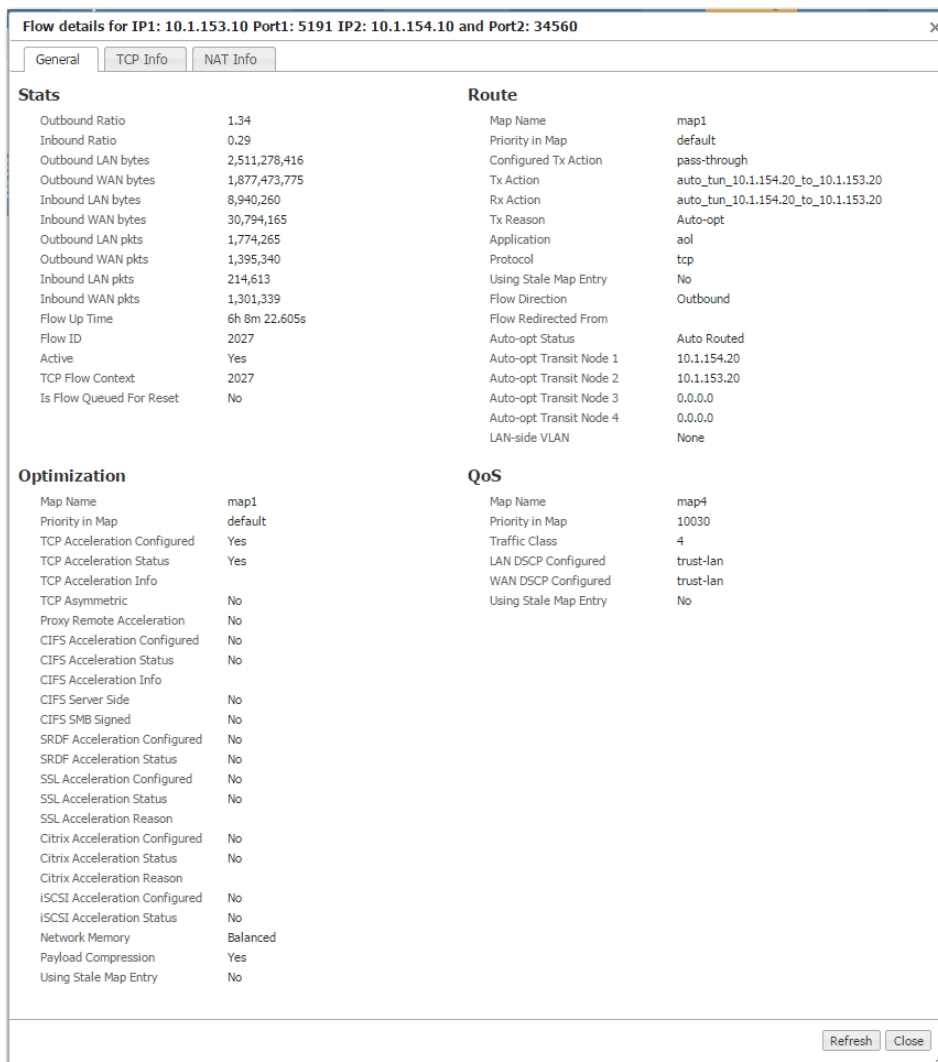
- To access the **Customize Current Flows Table**, click **Customize**.



- Select additional columns, and click **OK**. The columns append to the right side of the table.

Flow Details

Silver Peak Support uses the **Flow Detail** page for troubleshooting.



Most of the information on the **Flow Detail** page exceeds what is included in the **Current Flows** table.

Field	Definition
Stats Information	
Outbound Ratio	For the outbound traffic, a ratio of the Outbound LAN bytes divided by the Outbound WAN bytes. When this ratio is less than 1.0, it's attributable to a fixed overhead (for WAN transmission) being applied to traffic that either is not compressible or consists of few packets.
Inbound Ratio	For the inbound traffic, a ratio of the Inbound WAN bytes divided by the Inbound LAN bytes.
Outbound LAN bytes	Total number of bytes received from the LAN [outbound traffic]
Outbound WAN bytes	Total number of bytes sent to the WAN [outbound traffic]
Inbound LAN bytes	Total number of bytes sent to the LAN [inbound traffic]

Field	Definition (Continued)
Inbound WAN bytes	Total number of bytes received from the WAN [inbound traffic]
Outbound LAN pkts	Total number of packets received from the LAN [outbound traffic]
Outbound WAN pkts	Total number of packets sent to the WAN [outbound traffic]
Inbound LAN pkts	Total number of packets sent to the LAN [inbound traffic]
Inbound WAN pkts	Total number of packets received from the WAN [inbound traffic]
Flow Up Time	The length of time that there has been a connection between the endpoints.
Flow ID	A unique number that the appliance assigns to the flow.
Active	Whether the flow is Active [yes] or not [No].
TCP Flow Context	Silver Peak uses this for debugging purposes.
Is Flow Queued for Reset	Whether the flow is waiting to be reset (after user input) or not.
Route	
Map Name	The name of the Route Policy.
Priority in Map	The number of the entry in the Route Policy that the flow matches.
Configured Tx Action	The SET action configured in the Route Policy's Tunnel field.
Tx Action	How the traffic is actually being transmitted. Usually, this is a tunnel name.
Rx Action	By what path or method the appliance is receiving this flow's traffic.
Tx Reason	Any error associated with packet transmission to the WAN.
Application	Name of the application to which that flow's traffic belongs.
Protocol	The flow's protocol.
Using Stale Map Entry	Whether or not the flow is using a policy entry that has been edited or deleted since the flow began.
Flow Direction	Whether the flow is Inbound or Outbound .
Flow Redirected From	The IP address of the appliance that's redirecting this flow to this appliance.
Auto-opt Status	Whether it matched a specific Route Policy or was Auto Routed.
Auto-opt Transit Node (1, 2, 3, 4)	The IP addresses of the hops between this appliance and the other end of the connection.
LAN-side VLAN	Specifies the VLAN tag (1 – 4095) or None.
Optimization	
Map Name	The name of the Optimization Policy.
Priority in Map	The number of the entry in the Optimization Policy that the flow matches.
TCP Acceleration Configured	Whether or not TCP acceleration is configured in the Optimization Policy.
TCP Acceleration Status	Whether TCP is accelerated [Yes] or not [No].
TCP Acceleration Info	The reason that a TCP flow is not accelerated.. For a list of error codes, see <i>"Error Reasons for TCP Acceleration Failure" on page 196.</i>

Field	Definition (Continued)
TCP Asymmetric	When the answer is YES , the Silver Peak appliance is able to intercept connection establishment in only one direction. As a result, this flow is not accelerated. When this happens, it indicates that there is asymmetric routing in the network.
Proxy Remote Acceleration	Which side is accelerating the flow
CIFS Acceleration Configured	Whether or not CIFS acceleration is configured in the Optimization Policy [Yes/No]
CIFS Acceleration Status	Whether CIFS is accelerated [Yes] or not [No].
CIFS Acceleration Info	The reason that a CIFS flow is not accelerated. For a list of error codes, see "Error Reasons for CIFS Acceleration Failure" on page 199
CIFS Server Side	[Yes/No] If Yes , then this is the server side and the appliance is not accelerating (only the client side accelerates).
CIFS SMB Signed	Specifies whether or not the CIFS traffic is SMB-signed by the server: <ul style="list-style-type: none"> • Yes means it was signed. If that's the case, then the appliance was unable to accelerate any CIFS traffic. • No means it wasn't signed. If that's the case, then server requirements did not preclude CIFS acceleration. • Overridden means that SMB signing is ON and the appliance overrode it.
SRDF Acceleration Configured	Whether or not SRDF acceleration is configured in the Optimization Policy [Yes/No]
SRDF Acceleration Status	Whether SRDF is accelerated [Yes] or not [No].
SSL Acceleration Configured	Whether or not SSL acceleration is configured in the Optimization Policy [Yes/No]
SSL Acceleration Status	If a certificate has been appropriately installed via the Orchestrator, then SSL traffic can be deduplicated. Whether SSL is accelerated [Yes] or not [No].
SSL Acceleration Reason	The reason that an SSL flow is not accelerated. For a list of error codes, see "Error Reasons for SSL Acceleration Failure" on page 200
Citrix Acceleration Configured	Whether or not Citrix cgp (gateway) or ica protocol acceleration is configured in the Optimization Policy [Yes/No]
Citrix Acceleration Status	Whether Citrix is accelerated [Yes] or not [No].
Citrix Acceleration Reason	The reason that a Citrix flow is not accelerated. For a list of error codes, see "Error Reasons for Citrix Acceleration Failure" on page 203
iSCSI Acceleration Configured	Whether or not iSCSI acceleration is configured in the Optimization Policy [Yes/No]
iSCSI Acceleration Status	The reason that an iSCSI flow is not accelerated.

Field	Definition (Continued)
Network Memory	<p>There are four Network Memory settings:</p> <ul style="list-style-type: none"> • Maximize Reduction — optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP where bandwidth savings are the primary concern. • Minimize Latency — ensures that no latency is added by Network Memory processing. This may come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It is also appropriate if WAN bandwidth saving is not a primary objective, and instead it is desirable to fully utilize the WAN pipe to increase LAN-side throughput. • Balanced — This is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types. • Disabled — No Network Memory is performed.
Payload Compression	Whether or not payload compression is turned on.
Using Stale Map Entry	Whether or not the flow is using a Route Policy entry that has been edited or deleted since the flow began.
QoS Information	
Map Name	The name of the QoS Policy.
Priority in Map	The number of the entry in the QoS Policy that the flow matches.
Traffic Class	The number of the traffic class assigned by the QoS to the flow, based on the MATCH conditions satisfied.
LAN DSCP Configured	The LAN DSCP marking that the QoS policy assigned to the flow, based on the MATCH conditions satisfied.
WAN DSCP Configured	The WAN DSCP marking that the QoS policy assigned to the flow, based on the MATCH conditions satisfied.
Using Stale Map Entry	Whether or not the flow is using a policy entry that has been edited or deleted since the flow began.

Error Reasons for TCP Acceleration Failure

Following is a list of possible errors, along with a brief description and possible resolutions.

Error Reason	Description
asymmetric flow	Appliance did not receive a SYN-ACK. RESOLUTION: Most likely reason is asymmetric routing.
client advertised zero MSS	Flow is not accelerated because an endpoint did not send the TCP MSS option. RESOLUTION: Sometimes older operating systems (like Windows 95) do not send the TCP MSS option. You will have to upgrade the operating system software on the endpoints.
connection reset by peer	During setup, this TCP flow's endpoint(s) reset the connection. RESOLUTION: This is a transient condition. If it persists, take a tcpdump for this flow from both the client and server machines and contact Silver Peak Support.
connection to be deleted	Flow is not accelerated due to an internal error. RESOLUTION: Contact Silver Peak Support for further help.
disabled in Optimization Map	TCP Acceleration disabled in the Optimization Map. RESOLUTION: If you want this flow to be TCP accelerated, enable it in the optimization map.
disabled to allow debug	Flow is not accelerated because it has been disabled by tunbug debug console. RESOLUTION: Contact Silver Peak Support for further help.
first packet not a SYN	Appliance did not see the TCP SYN for this flow and therefore could not accelerate it. RESOLUTION: This could be due to various reasons: <ol style="list-style-type: none"> 1. The flow is already established before the appliance sees the first packet for the flow. If so, then resetting the flow will fix the problem. 2. WCCP or PBR is not set up correctly to redirect outbound traffic to the appliance. Check the WCCP or PBR configuration on the router. 3. You have routing issues, so the appliance is not seeing some of the traffic (for example, some packets come to the appliance while others go through another router). If so, you must review and fix your routing. 4. If you are in a cluster of Silver Peak appliances, you may have received a flow redirection timeout. If so, you must investigate why it takes so long for the Silver Peak appliance clusters to communicate with each other.
IP briefly blacklisted	Appliance did not receive a TCP SYN-ACK from remote end within 5 seconds and allowed the flow to proceed unaccelerated. Consequently, the destination IP address has been blacklisted for one minute. RESOLUTION: Wait for a minute and then reset the flow. If the problem reappears, the two most likely reasons are: 1) The remote server is slow in responding to TCP connection requests, or 2) a firewall is dropping packets containing Silver Peak TCP options. To check for either of these causes, perform a tcpdump on the server, with the filter set to these IP addresses: <ul style="list-style-type: none"> • If you don't see a TCP SYN from the client, it is due to firewall or routing issues. • If you notice that SYN-ACK was sent by the server after 5 seconds, it is due to a slow server.

Error Reason	Description (Continued)
keep alive failure	<p>Appliance did not receive a TCP SYN-ACK from the remote end within 5 seconds and allowed the flow to proceed unaccelerated.</p> <p>RESOLUTION: Wait for a minute and then reset the flow. If the problem reappears, the two most likely reasons are: 1) The remote server is slow in responding to TCP connection requests, or 2) a firewall is dropping packets containing Silver Peak TCP options.</p> <p>To check for either of these causes, perform a tcpdump on the server, with the filter set to these IP addresses:</p> <ul style="list-style-type: none"> • If you don't see a TCP SYN from the client, it is due to firewall or routing issues. • If you notice that SYN-ACK was sent by the server after 5 seconds, it is due to a slow server.
no remote appliance detected	<p>Appliance did not receive Silver Peak TCP option in the inbound direction.</p> <p>RESOLUTION: This could be due to various reasons:</p> <ol style="list-style-type: none"> 1. WCCP or PBR is not configured properly on the peer appliance. 2. Silver Peak routing policy not configured properly on the peer appliance. 3. Peer appliance is out of resources. 4. Routing is not configured properly on the router.
out of TCP memory	<p>Appliance is out of resources for accelerating TCP flows.</p> <p>RESOLUTION: Contact Silver Peak about upgrading to an appliance with higher flow capacity.</p>
remote appliance dropped out of accel	<p>Flow is not accelerated because Silver Peak flag is not set in TCP header or there was a mismatch in internal settings.</p> <p>RESOLUTION: Contact Silver Peak Support for further help.</p>
retransmission timeout	<p>Flow is not accelerated due to TCP protocol timeouts.</p> <p>RESOLUTION: This is a transient condition. You can reset the flow and then verify that it gets accelerated. If it does not, then take a tcpdump for this flow from both the client and server machines and contact Silver Peak Support.</p>
Route Map set to drop packets	<p>Flow is not accelerated because the route policy is set to drop packets.</p> <p>RESOLUTION: Fix the Set Action in the route policy entry.</p>
Route Map set to pass-through	<p>Flow is not accelerated because the route policy is set to send packets pass-through.</p> <p>RESOLUTION: Fix the Set Action in the route policy entry.</p>
software version mismatch	<p>Flow is not accelerated due to software version mismatch between two appliances.</p> <p>RESOLUTION: Upgrade software on one or both appliances to the same version of software.</p>
stale flow	<p>Flow is not accelerated due to an internal error. Before the previous flow could terminate cleanly, a new flow began with the same parameters.</p> <p>RESOLUTION: Contact Silver Peak Support for further help.</p>
SYN packet fragmented	<p>Flow is not accelerated for unknown reasons. Please contact Silver Peak Support for further help.</p> <p>RESOLUTION: Contact Silver Peak Support for further help. You may want to reset the connection to see if the problem resolves.</p>

Error Reason	Description (Continued)
system flow limit reached	Appliance has reached its limit for the total number of flows that can be accelerated. RESOLUTION: Contact Silver Peak about upgrading to an appliance with higher flow capacity.
tandem SP appliance involved	Appliance saw Silver Peak TCP option in the outbound direction. This implies that another Silver Peak appliance precedes this one and is responsible for accelerating this flow. RESOLUTION: Check the flow acceleration status on an upstream appliance.
TCP auto-optimization failed	Automatic optimization logic failed to accelerate this flow. These are handled for each auto-opt subcode below: <ul style="list-style-type: none"> <li data-bbox="760 594 1421 793">• TCP auto-optimization failed - NOSPS Auto-optimization failed because the peer appliance is not participating in automatic TCP acceleration. This can be due to various reasons: 1. Peer appliance is configured to not participate in optimization. 2. WCCP or PBR is not configured properly on the peer side. 3. Routing is not configured properly to send traffic to the peer appliance. <li data-bbox="760 804 1421 972">• TCP auto-optimization failed - NOTUNNEL Auto-optimization failed because there is no tunnel between this appliance and its peer, for two possible reasons: 1) Auto-tunnel is disabled. If so, manually create a tunnel. 2) Auto-tunnel is enabled, but needs time to finish creating the tunnel. If so, wait ~30 seconds for tunnel completion, and then reset this flow. <li data-bbox="760 982 1421 1066">• TCP auto-optimization failed - INVALID_OPT This is generally due to an internal error. Contact Silver Peak Support for further help. <li data-bbox="760 1077 1421 1140">• TCP auto-optimization failed - MISC Contact Silver Peak Support for further help. <li data-bbox="760 1150 1421 1245">• TCP auto-optimization failed - TUNNELDOWN Automatic optimization failed because the tunnel between this appliance and its peer is down.
TCP state mismatch	Flow is not accelerated due to an internal error. This flow will be automatically reset soon. RESOLUTION: This is a transient condition. You can wait for this flow to reset, or you can reset it manually now.
terminated by user	Flow has been reset by the user or automatically reset by the system. RESOLUTION: This is a transient condition. The flow is in the process of being reset.
tunnel down	Flow is not accelerated because the tunnel is down. RESOLUTION: Investigate why the tunnel is down.
unknown cause	Flow is not accelerated for unknown reasons. RESOLUTION: Contact Silver Peak Support for further help. You may want to reset the connection to see if the problem resolves.

Error Reasons for CIFS Acceleration Failure

When there is an acceleration failure, the appliance generates an **Alert** link that you can access on the **Flows** page. The **Alert** details the reason and the possible resolution.

Following is a list CIFS reason codes:

Reason Text	Description
CIFS optimization is disabled in the Optimization Policy	CIFS is disabled in the optmap.
SMB signing is required by the server	SMB signing is enforced by the server, and this requirement precludes optimization.
SMB version 2 is enforced by the client	SMB version 2 protocol is enforced by the client, and this requirement precludes optimization.
The flow limit for CIFS optimization has been exceeded	Maximum flow limit reach for CIFS optimized flows.
Sub-optimal read-write optimization - Non standard server	Sub-optimal read/write optimization due to non-standard server. For example, Windows XP cannot process more than 10 simultaneous outstanding requests.
Metadata optimization disabled - NTNOTIFY failure	Metadata optimization is disabled due to change notification failure.
Metadata optimization disabled - OPEN failure	Metadata optimization is disabled because proxy cannot open the root share. To resolve, check the root share permissions.
Metadata optimization disabled - Unsupported Dialect	Endpoints are using an unsupported CIFS dialect. To resolve, upgrade the CLIFS client/server.
Metadata optimization disabled - Unsupported Server	Unsupported CIFS server, like UNIX/Samba. To resolve, switch to standard servers like Windows/NetApp..
Metadata optimization disabled - Unsupported Client	Unsupported CIFS client, like UNIX/smbclient. To resolve, switch to standard clients like Windows/Mac.

Error Reasons for SSL Acceleration Failure

When there is an acceleration failure, the appliance generates an **Alert** link that you can access on the **Flows** page. The **Alert** details the reason and the possible resolution.

Silver Peak supports:

- X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
- SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.

Silver Peak appliances support the following:

- **Protocol versions:** SSLv3, TLS1.0, TLS1.1, TLS1.2
- **Cipher algorithms:** AES128, AES256, RC4, 3DES
- **Key exchanges:** RSA, DHE, ECDHE
- **Digests:** MD5, SHA1, SHA2

Following is a list of the reasons you may receive a failure message for SSL acceleration.

If the resolution calls for removing or reinstalling the certificate, refer to [“SSL Certificates Template” on page 69](#).

Error Reason	Description
error processing certificate	<p>Failure in processing certificate.</p> <p>RESOLUTION: Check the certificate. Possible problems include:</p> <ul style="list-style-type: none"> • There may be an issue with certificate format. • The certificate doesn't match the one that's installed on the server.
error processing client hello1	<p>Failed to create client hello, protocol error, invalid SSL packet, or internal error</p> <p>RESOLUTION: Check the SSL protocols on the client and the server. They must be compatible with what Silver Peak supports. If you find that they're incompatible, you must remove it and install the correct certificate.</p>
error processing client hello2	<p>Unsupported client SSL protocol version or options</p> <p>RESOLUTION: Check the SSL protocol on the client and the server. They must be compatible with what Silver Peak supports.</p>
error processing client hello3	<p>Invalid random number in SSLv2 client hello, protocol error, invalid SSL packet, or internal error</p> <p>RESOLUTION: Check the SSL protocol on the client and the server. They must be compatible with what Silver Peak supports.</p>
error processing SAN certificate	<p>Error while processing SAN certificate.</p> <p>RESOLUTION: Check the Subject Alternate Name fields in the SAN certificate. It may be an issue with SAN certificate format or with the certificate not matching the one that's installed on the server. If it's incorrect, you must remove it, and install the correct certificate.</p>
error processing server hello	<p>Error while processing server hello</p> <p>RESOLUTION: Contact Silver Peak Support for assistance.</p>

Error Reason	Description (Continued)
extension parse error	<p>TLS extension parse error, due to unknown TLS extensions</p> <p>RESOLUTION:</p> <ol style="list-style-type: none"> 1. Check the appliance syslog messages (that correspond to the client IP address) for SSL errors to determine which TLS extension is not supported. 2. Disable this (these) extensions in the client-side application's SSL settings. Typically, this application would be your browser.
invalid certificate	<p>SSL certificate is invalid or has expired.</p> <p>RESOLUTION: Remove the certificate, and reinstall the correct certificate.</p>
invalid client cipher	<p>Client negotiated unsupported cipher algorithm</p> <p>RESOLUTION: Check the client-side application's SSL cipher algorithm settings to verify that they're compatible with what Silver Peak supports.</p>
invalid client proto version	<p>Client negotiated unsupported SSL protocol version.</p> <p>RESOLUTION: Check the client-side application's SSL protocol settings to verify that they're compatible with what Silver Peak supports.</p>
invalid handshake condition	<p>Received invalid SSL packet or unsupported SSLv2 session resume request during handshake</p> <p>RESOLUTION: Contact Silver Peak Support for assistance.</p>
invalid key	<p>SSL private key is invalid</p> <p>RESOLUTION: Check that the private key file that was installed is correct and matches the server's private key.</p>
invalid server cipher	<p>Server negotiated unsupported cipher algorithm</p> <p>RESOLUTION: Check the SSL server's cipher algorithm settings.</p>
invalid server proto version	<p>Server negotiated unsupported SSL version</p> <p>RESOLUTION: Check the server-side application's SSL protocol settings to verify that they're compatible with what Silver Peak supports.</p>
memory flow control	<p>The appliance SSL memory is full and cannot accelerate additional flows.</p> <p>RESOLUTION: Contact Silver Peak support for assistance.</p>
miscellaneous error	<p>Generic proxy layer internal error</p> <p>RESOLUTION: Contact Silver Peak Support for assistance.</p>
missing active session	<p>Active session not found, cannot accelerate the SSL session. The appliance did not participate in the full handshake phase where the certificate information was exchanged between the client and the server.</p> <p>Or, the certificate was missing or did not match the server's certificate.</p> <p>RESOLUTION: If the certificate is missing, install the correct one. Otherwise, restart the client SSL application.</p>
missing certificate	<p>A matching SSL certificate was not found.</p> <p>RESOLUTION: Install the certificate on both appliances.</p>
missing key	<p>A matching SSL key was not found.</p> <p>RESOLUTION: Install the correct certificate and key.</p>
missing pending session	<p>Pending session not found, possible failure in client hello.</p> <p>RESOLUTION: Contact Silver Peak Support for assistance.</p>
missing resume session	<p>Do not have a session to resume in session cache. The session in Silver Peak's cache might have expired.</p> <p>RESOLUTION: To get full SSL acceleration, restart the application.</p>

Error Reason	Description (Continued)
missing SAN certificate	Did not find a matching SAN certificate. RESOLUTION: Install the missing SAN certificate.
no ipsec on tunnel	IPsec is not configured on the tunnel. RESOLUTION: Configure IPsec on the tunnel.
possibly no certs installed	Possibly no SSL certificate installed. RESOLUTION: If the Orchestrator shows no SSL certificate, install an appropriate one.
server-side advertised no dedup	Peer appliance SSL did not optimize the flow. RESOLUTION: On the other appliance, access the Current Flows report, and look at the reason code. (In some cases, the code is displayed only on one side).
ssl max flows limit	Exceeded maximum SSL optimized flows limit.
unsupported client cipher	Received unsupported cipher suite in SSLv2 client hello message. RESOLUTION: Check the client-side application's SSL cipher algorithm settings to verify that they're compatible with what Silver Peak supports. Check the client-side SSL protocol version settings. Silver Peak does not support SSLv2.
unsupported compress method	Unsupported SSL compression method negotiated. The SSL compression method should be disabled on both the client and the server. RESOLUTION: On both the client and the server, disable the SSL compression method.
unsupported extension	Unsupported TLS extension negotiated. RESOLUTION: 1. Check the appliance syslog messages (that correspond to the client IP address) for SSL errors to determine which TLS extension is not supported. 2. Disable this (these) extensions in the client-side application's SSL settings. Typically, this application would be your browser.
unsupported server cipher	Received unsupported cipher suite in SSLv2 server hello message. RESOLUTION: Check the server-side application's SSL cipher algorithm settings to verify that they're compatible with what Silver Peak supports. Check the server-side SSL protocol version settings. Silver Peak does not support SSLv2.
unsupported server protocol	Unsupported SSL protocol: SSLv2 server hello message not supported. RESOLUTION: Check the server-side application's SSL protocol settings to verify that they're compatible with what Silver Peak supports.

Error Reasons for Citrix Acceleration Failure

When there is an acceleration failure, the appliance generates an **Alert** link that you can access on the **Flows** page. The **Alert** details the reason and the possible resolution.

Reason Text	Description
Exceeded max flows	Flow will not be accelerated because max citrix flow limit has been reached. RESOLUTION: Check
Exceeded max CGP sessions	Flow will not be accelerated because max Citrix CGP session limit has been reached.
Missing CGP data	Flow will not be accelerated because session resume did not find the CGP session.
Connection Alloc failure	Connection element could not be allocated. RESOLUTION: Contact Silver Peak.
Pending Full Accel	Full acceleration criteria in ICA protocol negotiation not yet satisfied. RESOLUTION: Relaunch the Citrix session.
Encryption level not supported	Encryption level not Basic/Secure on the Citrix server or client. RESOLUTION: Check the encryption level setting on the Citrix server.
Packet Alloc failure	Packet allocation has failed. Packets will be forwarded. RESOLUTION: Contact Silver Peak.
Citrix Protocol not as expected	Some pre-defined pattern in Citrix ICA Protocol negotiation is not as expected. RESOLUTION: 1. Verify that non-Citrix traffic is not being sent over the Citrix ports. 2. Check the Citrix protocol versions used in the client and server and call Silver Peak.
Citrix optimization failed due to an unknown error	RESOLUTION: 1. Disabling Citrix Acceleration in the Optimization Policy is recommended. 2. Review the system logs to find the exact error code and contact Silver Peak.
Citrix rc5 padding error	Citrix ICA record did not align on 8-byte boundary or a padding error occurred. RESOLUTION: Contact Silver Peak.
Citrix rc5 Decrypt error	Citrix ICA record failed decryption. RESOLUTION: Contact Silver Peak.
Citrix rc5 Encrypt error	Citrix ICA record failed encryption. RESOLUTION: Contact Silver Peak.
Citrix rc5 Misc error	Citrix ICA RC5 unknown error. RESOLUTION: See system logs. Contact Silver Peak.
Citrix rc5 crypto buffer too short	Citrix RC5 crypto buffer passed in to encryption/decryption was too short. RESOLUTION: Contact Silver Peak.
Citrix rc5 crypto buffer too long	Citrix RC5 crypto buffer passed in to encryption/decryption was too long. RESOLUTION: Contact Silver Peak.

Reason Text	Description (Continued)
Citrix rc5 crypto invalid length	Citrix RC5 crypto invalid length was passed in for encryption/decryption. RESOLUTION: Contact Silver Peak.
Citrix rc5 crypto initialization failed	Citrix RC5 crypto engine failed initialization. RESOLUTION: Contact Silver Peak.

Resetting Flows to Improve Performance

In the list of **Alerts**, you can look for the flows that aren't being accelerated, but *could* be. Generally, this means flows that use TCP protocol and are not TCP-accelerated:

- This includes tunnelized TCP traffic that is **not** TCP-accelerated. TCP connections are not accelerated if they already exist when the tunnel comes up or when the appliance reboots.
- Pass-through connections are neither tunnelized nor accelerated if they already exist when a new tunnel is added and/or when an ACL is added or edited.

Unaccelerated TCP flows can be reset to allow them to reconnect at a later time. It is assumed that the connection end-points will re-establish the flows. When these flows are reconnected, the appliance recognizes them as new and accelerates them. Note that the time it takes to reset a flow may vary, depending on the traffic activity.



CAUTION **Resetting a flow interrupts service for that flow.** The appliance cannot restore the connection on its own; it relies on the end points to re-establish the flow. Use it only if service interruption can be tolerated for a given flow.



Tip For information about configuring the appliance to automatically reset TCP flows, see the Advanced TCP Options in *"TCP Acceleration Options"* on page 63.

Monitoring Status & Reporting

The following reports exist for monitoring status and reporting:

- **View Reports** See page 206.
- **Scheduled & Historical Jobs** See page 207.
- **Realtime Charts** See page 208.
- **Historical Charts** See page 208.
- **Reachability Tab** See page 209.

View Reports

Monitoring > [Status & Reporting] View Reports

This link opens a separate browser tab with links to the generated web reports.

Directory: /webreports/		
01.Feb.16-00.30.03-Daily-Global_Report-appliancesDataTransferAndReduction.jpeg	69859 bytes	Feb 1, 2016 12:30:46 AM
01.Feb.16-00.30.03-Daily-Global_Report-applicationReduction.jpeg	76057 bytes	Feb 1, 2016 12:30:09 AM
01.Feb.16-00.30.03-Daily-Global_Report-summary.jpeg	8914 bytes	Feb 1, 2016 12:30:54 AM
01.Feb.16-00.30.03-Daily-Global_Report-tunnelDataTransferAndReduction.jpeg	64944 bytes	Feb 1, 2016 12:30:39 AM
01.Feb.16-00.30.03-Daily-Global_Report-tunnelsLatency.jpeg	53829 bytes	Feb 1, 2016 12:30:31 AM
01.Feb.16-00.30.03-Daily-Global_Report-tunnelsLoss.jpeg	65079 bytes	Feb 1, 2016 12:30:17 AM
01.Feb.16-00.30.03-Daily-Global_Report-tunnelsMaxOutOfOrderPackets.jpeg	66557 bytes	Feb 1, 2016 12:30:24 AM
01.Feb.16-00.30.03-Hourly-Global_Report-appliancesDataTransferAndReduction.jpeg	68558 bytes	Feb 1, 2016 12:32:34 AM
01.Feb.16-00.30.03-Hourly-Global_Report-applicationReduction.jpeg	73766 bytes	Feb 1, 2016 12:31:55 AM
01.Feb.16-00.30.03-Hourly-Global_Report-summary.jpeg	8885 bytes	Feb 1, 2016 12:32:41 AM
01.Feb.16-00.30.03-Hourly-Global_Report-tunnelDataTransferAndReduction.jpeg	64186 bytes	Feb 1, 2016 12:32:26 AM
01.Feb.16-00.30.03-Hourly-Global_Report-tunnelsLatency.jpeg	54896 bytes	Feb 1, 2016 12:32:18 AM
01.Feb.16-00.30.03-Hourly-Global_Report-tunnelsLoss.jpeg	65079 bytes	Feb 1, 2016 12:32:03 AM
01.Feb.16-00.30.03-Hourly-Global_Report-tunnelsMaxOutOfOrderPackets.jpeg	64766 bytes	Feb 1, 2016 12:32:11 AM
01.Feb.16-00.30.03-Minutes-Global_Report-appliancesDataTransferAndReduction.jpeg	76650 bytes	Feb 1, 2016 12:31:40 AM
01.Feb.16-00.30.03-Minutes-Global_Report-applicationReduction.jpeg	82437 bytes	Feb 1, 2016 12:31:01 AM
01.Feb.16-00.30.03-Minutes-Global_Report-summary.jpeg	8852 bytes	Feb 1, 2016 12:31:47 AM
01.Feb.16-00.30.03-Minutes-Global_Report-tunnelDataTransferAndReduction.jpeg	72265 bytes	Feb 1, 2016 12:31:32 AM
01.Feb.16-00.30.03-Minutes-Global_Report-tunnelsLatency.jpeg	63719 bytes	Feb 1, 2016 12:31:24 AM
01.Feb.16-00.30.03-Minutes-Global_Report-tunnelsLoss.jpeg	72825 bytes	Feb 1, 2016 12:31:09 AM
01.Feb.16-00.30.03-Minutes-Global_Report-tunnelsMaxOutOfOrderPackets.jpeg	72648 bytes	Feb 1, 2016 12:31:17 AM
01.Feb.16-00.32.51-Daily-Global_Report.pdf	26534 bytes	Feb 1, 2016 12:32:54 AM
01.Feb.16-00.32.51-Hourly-Global_Report.pdf	26781 bytes	Feb 1, 2016 12:32:54 AM
01.Feb.16-00.32.51-Minutes-Global_Report.pdf	35769 bytes	Feb 1, 2016 12:32:54 AM
01.Feb.16-00.32.51-rawdata-Global_Report.zip	16111 bytes	Feb 1, 2016 12:32:54 AM
02.Feb.16-00.30.04-Daily-Global_Report-appliancesDataTransferAndReduction.jpeg	69697 bytes	Feb 2, 2016 12:30:49 AM
02.Feb.16-00.30.04-Daily-Global_Report-applicationReduction.jpeg	76012 bytes	Feb 2, 2016 12:30:11 AM
02.Feb.16-00.30.04-Daily-Global_Report-summary.jpeg	8806 bytes	Feb 2, 2016 12:30:56 AM
02.Feb.16-00.30.04-Daily-Global_Report-tunnelDataTransferAndReduction.jpeg		

Scheduled & Historical Jobs

Monitoring > [Status & Reporting] Scheduled & Historical Jobs

This tab has two views:

- It provides a central location for viewing and deleting **scheduled jobs**, such as appliance backup and any custom reports configured for distribution.

Job	Appliances	Description	Schedule	Last Run	Next Run	Status
Orchestrator...	Silver Peak Systems	Global Report	Every day at 0:30 starting 03-Jun-14 0:30 PDT	21-Feb-16 00:30...	22-Feb-16 00:30...	Success - Global Rep... X
QosMap Acti...	Release 8.0.2	primary map	Every day at 6:00 starting 21-Feb-16 20:38 PST		22-Feb-16 06:00...	X
QosMap Acti...	Release 8.0.2	evening map	Every day at 20:00 starting 21-Feb-16 20:38 PST		22-Feb-16 20:00...	X

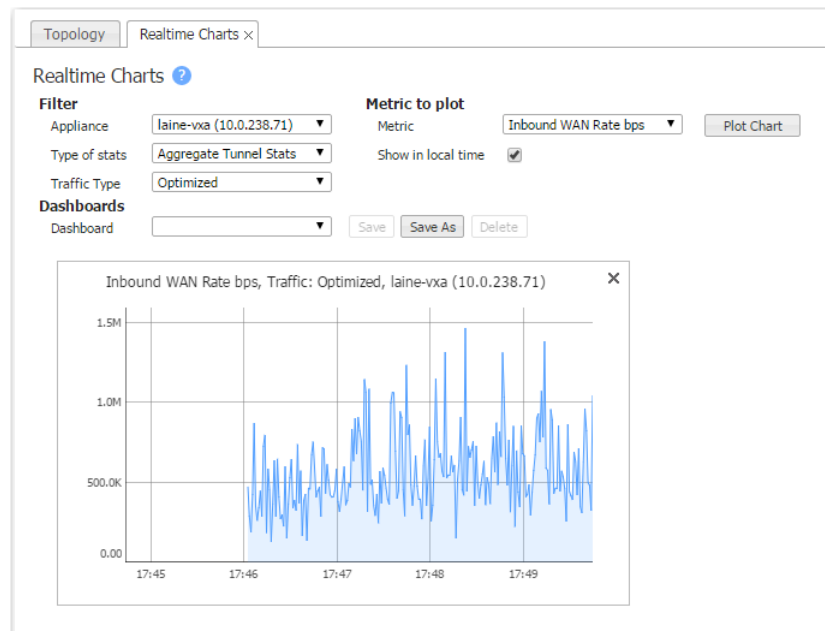
- It provides a central location for viewing **historical jobs**.

Job	Appliances	Description	Start Time	Duration	Status
Orchestrator Report	Silver Peak Systems	Global Report	21-Feb-16 00:30 PST	2m 27s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	20-Feb-16 00:30 PST	2m 28s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	19-Feb-16 00:30 PST	2m 28s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	18-Feb-16 00:30 PST	2m 26s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	17-Feb-16 00:30 PST	2m 32s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	16-Feb-16 00:30 PST	2m 25s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	15-Feb-16 00:30 PST	3m 28s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	14-Feb-16 00:30 PST	2m 27s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	13-Feb-16 00:30 PST	3m 59s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	12-Feb-16 00:30 PST	3m 25s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	11-Feb-16 00:30 PST	4m 13s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	10-Feb-16 00:30 PST	3m 25s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	09-Feb-16 00:30 PST	2m 25s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	08-Feb-16 00:30 PST	2m 46s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	07-Feb-16 00:30 PST	2m 26s	Success - Global Report Time ta...
Orchestrator Report	Silver Peak Systems	Global Report	06-Feb-16 00:30 PST	6m 16s	Success - Global Report Time ta...
Orchestrator Report	10.0.238.20(laine2-va),10.0.238.21(lai...	Global Report	05-Feb-16 00:30 PST		Success - Global Report Time ta...
Orchestrator Report	10.0.238.20(laine2-va),10.0.238.21(lai...	Global Report	04-Feb-16 00:30 PST		Success - Global Report Time ta...

Realtime Charts

Monitoring > [Status & Reporting] Realtime Charts

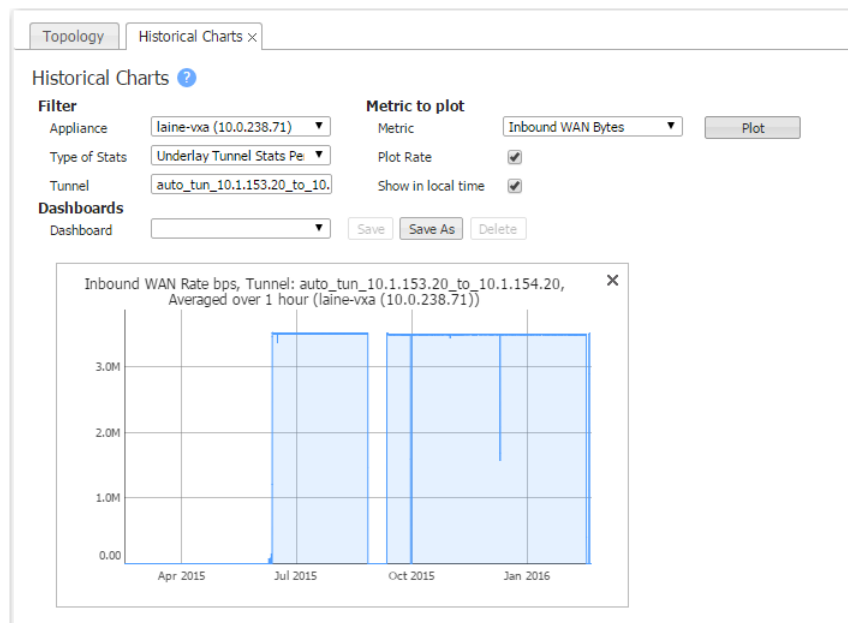
As an aid to troubleshooting, **Realtime Charts** are useful for monitoring the performance of individual appliances. You can save sets of charts as dashboards.



Historical Charts

Monitoring > [Status & Reporting] Historical Charts

As an aid to troubleshooting, **Historical Charts** are useful for reviewing the performance of individual appliances. You can save sets of charts as dashboards.



Reachability Tab

Monitoring > [Status & Reporting] Reachability

This page summarizes the status of communications in two directions -- **Orchestrator to Appliances** and **Appliances to Orchestrator**.

The top screenshot shows the 'Orchestrator to Appliances' view. The table has the following data:

Appliance Name	Admin Username	Protocol	State	Unsaved Changes
Tallinn	admin	HTTPS	Normal	Yes
laine-vxa	admin	HTTPS	Normal	No
laine-vxb	admin	HTTPS	Normal	No
laine2-vxa	admin	HTTPS	Normal	No
laine2-vxb	admin	HTTPS	Normal	No

The bottom screenshot shows the 'Appliance to Orchestrators' view. The table has the following data:

Appliance Name	Orchestrator IP	REST	SSH	HTTPS	Web Socket
Tallinn	10.0.238.26	Reachable	Reachable	Reachable	Reachable
Tallinn	10.0.238.70	Reachable	Reachable	Reachable	Reachable
laine-vxa	10.0.238.26	Reachable	Reachable	Reachable	Reachable
laine-vxa	10.0.238.70	Reachable	Reachable	Reachable	Reachable
laine-vxa	23.21.179.138	Unreachable	Reachable	Unreachable	Unreachable
laine-vxa	23.21.224.150	Unreachable	Reachable	Unreachable	Unreachable
laine-vxb	10.0.238.26	Reachable	Reachable	Reachable	Reachable
laine-vxb	10.0.238.70	Reachable	Reachable	Reachable	Reachable
laine2-vxa	10.0.238.26	Reachable	Reachable	Reachable	Reachable
laine2-vxb	10.0.238.26	Reachable	Reachable	Reachable	Reachable

- **Admin Username** is the username that an Orchestrator server uses to log into an appliance.
- An Orchestrator can use the web protocols, **HTTP**, **HTTPS**, or **Both** to communicate with an appliance. Although **Both** exists for legacy reasons, Silver Peak recommends using **HTTPS** for maximum security.
- An appliance's **State** can be Normal, Unknown, Unsupported, or Unreachable.
 - **Normal** indicates that all is well.
 - **Unknown** is a transitional state that appears when first adding an appliance to the network.
 - **Unsupported** indicates an unsupported version of appliance software.
 - **Unreachable** indicates a problem in your network. Check your ports, firewalls, and deployment configuration.
- The **Appliance to Orchestrators** table lists the protocols that the appliance uses to communicate with an Orchestrator.
 - HTTPS and Web Socket share Port 443.



Orchestrator Administration

This chapter describes the administrative tasks that directly relate to managing **Orchestrator-related events and tasks only**. These activities do not relate to managing appliances.

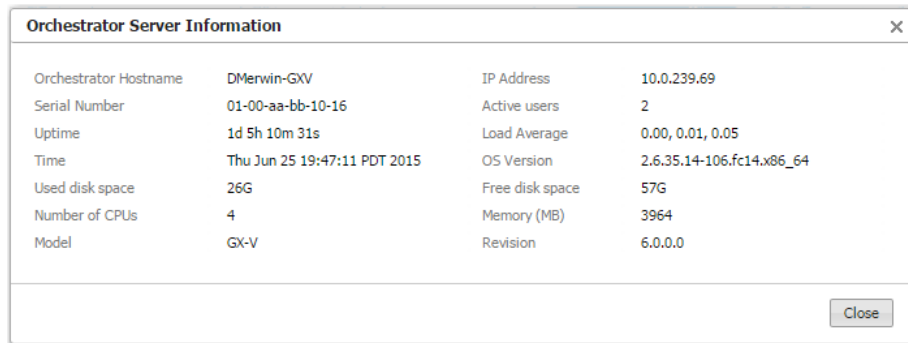
In This Chapter

- **Viewing Orchestrator Server Information** See page 212.
- **Restart, Reboot, or Shutdown** See page 212.
- **Managing the Orchestrator Server License** See page 212.
- **Managing Orchestrator Users** See page 213.
- **User Menu Access** See page 214.
- **Remote Authentication** See page 216.
- **Debug Files** See page 217.
- **Silver Peak Cloud Portal** See page 218.
- **SMTP Server Settings** See page 219.
- **Overlay Manager Settings** See page 220.
- **Schedule Timezone** See page 221.
- **Audit Logs** See page 222.
- **Getting Started Wizard** See page 223.
- **Proxy Configuration** See page 224.
- **Backing Up the Orchestrator Database** See page 227.

Viewing Orchestrator Server Information

Orchestrator Administration > [General] Server Information

This page lists specifications and data specific to this Orchestrator server.



Orchestrator Server Information			
Orchestrator Hostname	DMerwin-GXV	IP Address	10.0.239.69
Serial Number	01-00-aa-bb-10-16	Active users	2
Uptime	1d 5h 10m 31s	Load Average	0.00, 0.01, 0.05
Time	Thu Jun 25 19:47:11 PDT 2015	OS Version	2.6.35.14-106.fc14.x86_64
Used disk space	26G	Free disk space	57G
Number of CPUs	4	Memory (MB)	3964
Model	GX-V	Revision	6.0.0.0

Restart, Reboot, or Shutdown

Orchestrator Administration > [General] Restart Orchestrator Application

Orchestrator Administration > [General] Reboot Server

Orchestrator Administration > [General] Shutdown Server

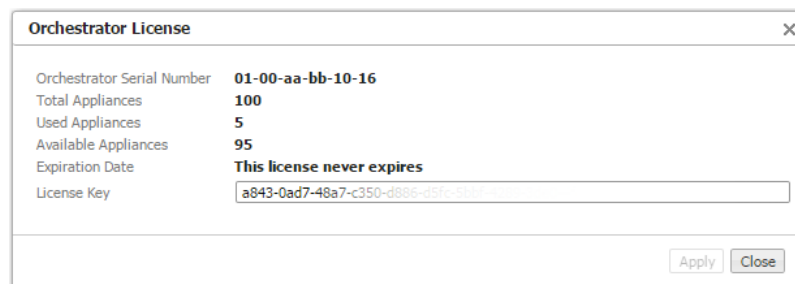
The Orchestrator provides these three actions as a convenience, in the **Orchestrator Administration** menu.

- **Restart Appliance** quickly restarts the Orchestrator software.
- **Reboot Orchestrator Server** is a more thorough restart, accomplished by rebooting the Orchestrator server.
- **Shutdown Server** results in the server being unreachable. You will have to manually power on the server to restart.

Managing the Orchestrator Server License

Orchestrator Administration > [General] License Management

Unity Orchestrator requires a license.



Orchestrator License	
Orchestrator Serial Number	01-00-aa-bb-10-16
Total Appliances	100
Used Appliances	5
Available Appliances	95
Expiration Date	This license never expires
License Key	<input type="text" value="a843-0ad7-48a7-c350-d88c-d5fc-510c-510c"/>

Managing Orchestrator Users

Orchestrator Administration > [General] User Management

The **User Management** page allows you to manage who has access to the Orchestrator server.

The screenshot shows the 'User Management' window with a table of users. The 'kalev' user is selected, and a dropdown menu is open showing the available roles: 'Admin Manager' and 'Network Monitor'.

User Name	First Name	Last Name	Phone	Email	Password	Repeat Passw...	Create time	Status	Role
admin	Admin				*****	*****	03/20/2014 2...	Active	Admin Manager
kalev					*****	*****	09/02/2014 0...	Active	Admin Ma

You cannot modify a Username. You must delete it and create a new user.

The Orchestrator has two user roles:

- **Admin Manager** has all privileges and can see/access all screens. It's the equivalent of **Superuser**.
- **Network Monitor** can view certain configuration, alarm, and report data. They can also troubleshoot network connectivity.

Authorization always maps to one of these.

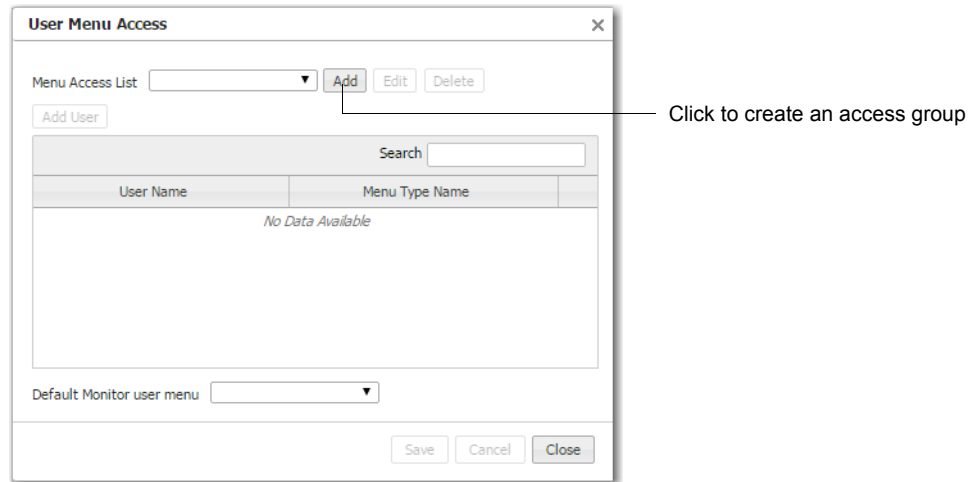
Guidelines for Creating Passwords

- Passwords should be a minimum of 8 characters.
- There should be at least one lower case letter and one upper case letter.
- There should be at least one digit.
- There should be at least one special character.
- Consecutive letters in the password should not be dictionary words.

User Menu Access

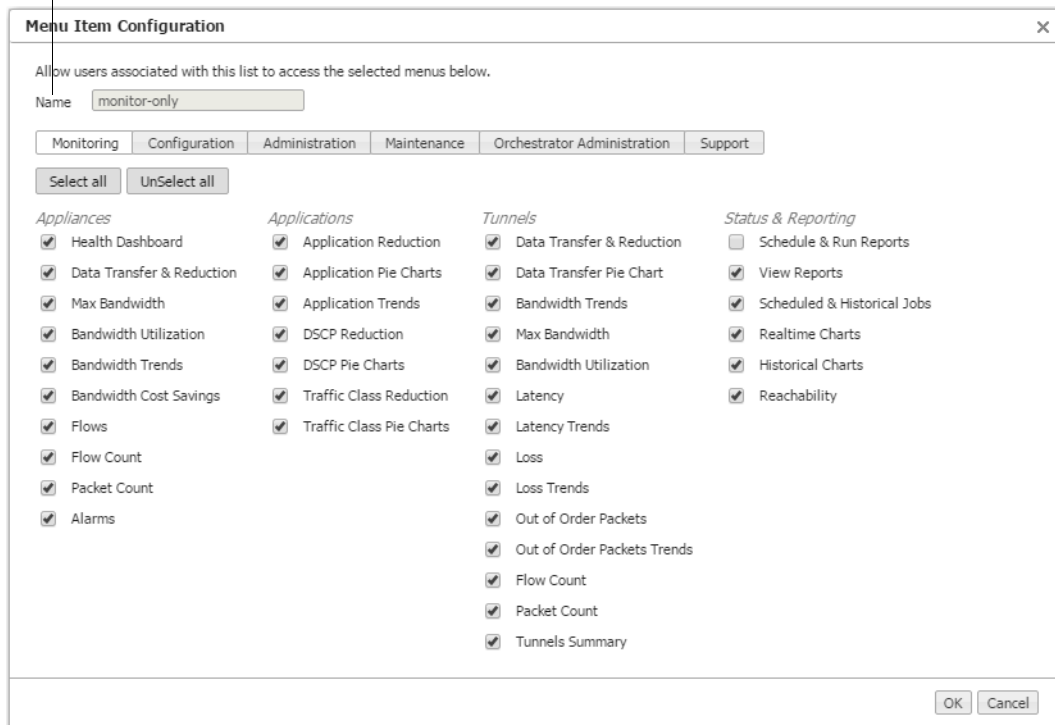
Orchestrator Administration > [General] User Menu Access

The **User Menu Access** dialog creates customized filters that restrict non-**admin** users' access to menus.



- You can create groups that restrict the access of users who aren't assigned the role of **Admin Manager** in the **Orchestrator Administration > User Management** table.

Name of the menu access list



- To assign a group to a specific user, click **Add User**. The following popup appears.

The image shows a 'User Menu Access' dialog box. At the top, there is a 'Menu Access List' dropdown menu set to 'monitor-only', with 'Add', 'Edit', and 'Delete' buttons to its right. Below this is an 'Add User' button. The main area contains a table with a search bar and a '1 Rows' indicator. The table has two columns: 'User Name' and 'Menu Type Name'. A single row is visible with 'dave' in the 'User Name' column and 'monitor-only' in the 'Menu Type Name' column. A small 'X' icon is in the rightmost cell of the row. At the bottom of the dialog, there is a 'Default Monitor user menu' dropdown menu set to 'monitor-only' and 'Save', 'Cancel', and 'Close' buttons.

User Name	Menu Type Name
dave	monitor-only

Remote Authentication

Orchestrator Administration > [General] Authentication

This **Authentication** page specifies how the Orchestrator authenticates Orchestrator users.

Local Only authenticates based on the users in the Orchestrator database.

◆ To authenticate using RADIUS or TACACS+

- 1 Select the access control protocol you want to use.
- 2 Under **Servers**, enter the information for a Primary server of that type. Entering a Secondary server is optional.

Field	Definition / Purpose
Authentication Order	Whether to use the remote map or the local map first. The default is Remote first .
Primary/Secondary Server	The IP address or hostname of the RADIUS or TACACS+ server.
Secret Key	The string defined as the shared secret on the server.
Admin Manager (Superuser) Privilege	These privilege levels must coincide with the values already configured for them at the RADIUS server.
Network Manager Privilege	
Network Monitor Privilege	
Authentication Type	When configuring to use the TACACS+ server, select either CHAP or PAP , to match what is configured on the TACACS+ server.

Debug Files

Orchestrator Administration > [General] Debug

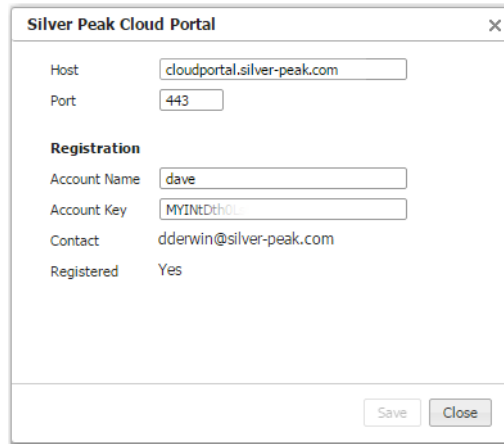
As a user, you won't need to refer to these tabs of statistics. If necessary, Silver Peak's engineers would be reviewing them in the context of a troubleshooting Webex.

Mgmt IP	Last Poll Time	Time between polls	Latest minute	Latest hour	Latest day	Minute R...	Hour Re...	Day Rec...
10.0.238.21 (1.NE)	02/21/2016 18:55:36	Average 60 Max 60 Min 60 P9...	02/21/2016 18:55:00	02/21/2016 17:00:00	02/20/2016 16:00:00	0	0	0
10.0.238.20 (0.NE)	02/21/2016 18:55:36	Average 60 Max 60 Min 60 P9...	02/21/2016 18:55:00	02/21/2016 17:00:00	02/20/2016 16:00:00	0	0	0
10.0.238.69 (11.NE)	02/21/2016 18:55:36	Average 60 Max 60 Min 60 P9...	02/21/2016 18:55:00	02/21/2016 17:00:00	02/20/2016 16:00:00	0	0	0
10.0.236.198 (12.NE)	02/21/2016 18:55:36	Average 60 Max 60 Min 60 P9...	02/21/2016 18:55:00	02/21/2016 17:00:00	02/20/2016 16:00:00	0	0	0
10.0.238.71 (10.NE)	02/21/2016 18:55:36	Average 60 Max 60 Min 60 P9...	02/21/2016 18:55:00	02/21/2016 17:00:00	02/20/2016 16:00:00	0	0	0

Silver Peak Cloud Portal

*Orchestrator Administration > [General] Silver Peak Cloud Portal
Configuration > [Unity Overlays] Silver Peak Cloud Portal*

The **Silver Peak Cloud Portal** is used to register cloud-based features and services, such as *SaaS optimization*, *EdgeConnect*, and *CPX*.



The screenshot shows a dialog box titled "Silver Peak Cloud Portal" with a close button (X) in the top right corner. The dialog contains the following fields and values:

Host	cloudportal.silver-peak.com
Port	443
Registration	
Account Name	dave
Account Key	MYINtDth0:
Contact	dderwin@silver-peak.com
Registered	Yes

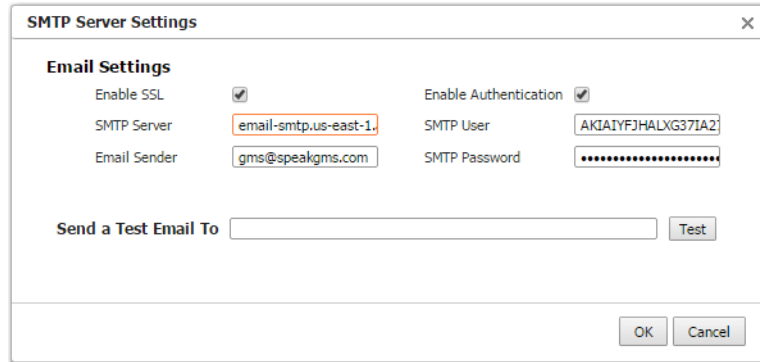
At the bottom right of the dialog, there are two buttons: "Save" and "Close".

- When you purchase one of these services, Silver Peak sends you an **Account Name** and instructions to obtain your **Account Key**. You will use these to register your appliance(s).
- The cloud portal populates the **Contact** field from information included in your purchase order.
- Use of these services requires that your appliance(s) can access the cloud portal via the Internet.

SMTP Server Settings

Orchestrator Administration > [General] SMTP Server Settings

The Orchestrator server sends **reports via email**, using a Silver Peak SMTP server in Amazon Web Services by default.



The screenshot shows a dialog box titled "SMTP Server Settings" with a close button (X) in the top right corner. The dialog is divided into two main sections: "Email Settings" and "Send a Test Email To".

Email Settings

Enable SSL	<input checked="" type="checkbox"/>	Enable Authentication	<input checked="" type="checkbox"/>
SMTP Server	<input type="text" value="email-smtp.us-east-1"/>	SMTP User	<input type="text" value="AKIAIYFJHALXG37IA2"/>
Email Sender	<input type="text" value="gms@speakgms.com"/>	SMTP Password	<input type="password" value="....."/>

Send a Test Email To

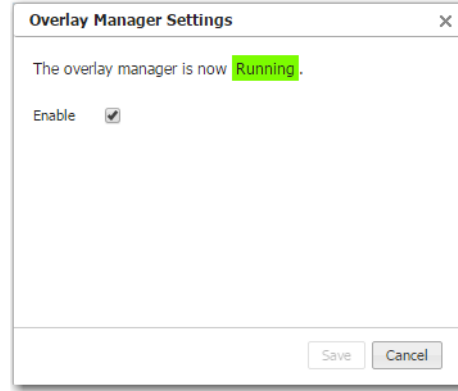
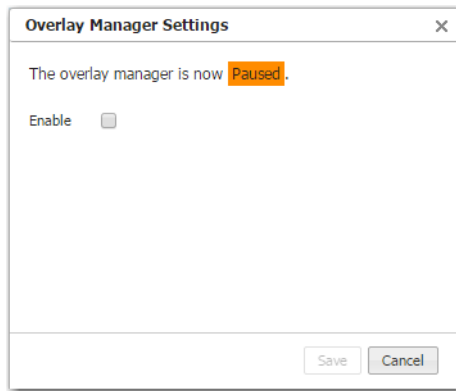
At the bottom right of the dialog are "OK" and "Cancel" buttons.

- Change the SMTP server and settings to your company's SMTP settings for permanent, private email delivery.
- If a test email doesn't arrive within minutes, check your firewall.

Overlay Manager Settings

Orchestrator Administration > [General] Overlay Manager Settings


You would only be using the **Overlay Manager Settings** dialog box at the request of Silver Peak Customer Support.



Schedule Timezone

Orchestrator Administration > [General] Schedule Timezone

Use this dialog to set the timezone for scheduled jobs and reports.



Audit Logs

Orchestrator Administration > [General] Audit Logs

The **Audit Logs** list actions initiated from the Orchestrator, whether by a user or by the system itself.

The screenshot shows the 'Audit Logs' interface with the following table structure and data:

User Name	Host Name	Action	Task Status	Results	Start Time	End Time	Queued Time	% Completed	Completion Status
admin		Overlay Manager Setting	COMPLETED	Overlay Manager Set Successfully: Overlay...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxa	Synchronize	COMPLETED	Synchronization successful. Last event on a...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	Tallinn	Synchronize	COMPLETED	Synchronization successful. Last event on a...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxb	Synchronize	COMPLETED	Synchronization successful. Last event on a...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxa	Synchronize	COMPLETED	Synchronization successful. Last event on a...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxb	Synchronize	COMPLETED	Synchronization successful. Last event on a...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxb	SaveChanges	COMPLETED	Save change on appliance successful	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxa	SaveChanges	COMPLETED	Save change on appliance successful	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxb	SaveChanges	COMPLETED	Save change on appliance successful	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxa	SaveChanges	COMPLETED	Save change on appliance successful	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	Tallinn	SaveChanges	COMPLETED	Save change on appliance successful	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	Tallinn	Delete Optimization map	COMPLETED	Optimization maps deleted. Data = {"map...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxa	Delete Optimization map	COMPLETED	Optimization maps deleted. Data = {"map...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxb	Delete Optimization map	COMPLETED	Optimization maps deleted. Data = {"map...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxa	Delete Optimization map	COMPLETED	Optimization maps deleted. Data = {"map...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxb	Delete Optimization map	COMPLETED	Optimization maps deleted. Data = {"map...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxb	Delete qos map	COMPLETED	QoS maps deleted. Data = {"map3":{"100...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine-vxa	Delete qos map	COMPLETED	QoS maps deleted. Data = {"map3":{"100...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	Tallinn	Delete qos map	COMPLETED	QoS maps deleted. Data = {"map1":{"100...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxa	Delete qos map	COMPLETED	QoS maps deleted. Data = {"map1":{"100...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success
OverlayManager	laine2-vxb	Delete qos map	COMPLETED	QoS maps deleted. Data = {"map1":{"100...	21-Feb-16 19:32	21-Feb-16 19:32	21-Feb-16 19:32	100	Success

- You can filter to determine whether actions are **Completed**, **In Progress**, or **Queued**.
- By default, the table refreshes automatically. However, you can **Pause** it to review a static set of data.

Getting Started Wizard

Orchestrator Administration > [General] Getting Started Wizard

When you first use the web browser to access the Orchestrator server's IP address, the Getting Started Wizard appears.

After initial configuration, you can always access the Getting Started Wizard from **Orchestrator Administration > Getting Started Wizard**.

The screenshot shows the 'Getting Started Wizard' window with six steps: 1. Hostname, DHCP, Password; 2. License and Registration; 3. Date/Time; 4. Email; 5. Add Appliances; 6. Backup. Step 1 is active. The configuration fields are as follows:

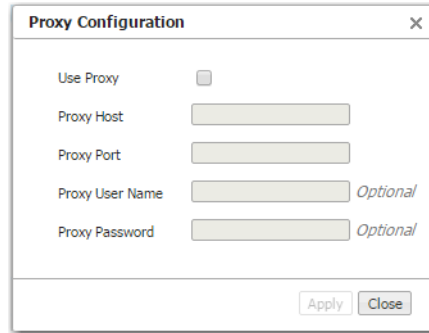
Section	Field	Value	
Orchestrator Name	Hostname	DMerwin-GW	
	Management Interface	<input checked="" type="radio"/> DHCP	
Change Admin Password (optional)	Old		
	New		
	Confirm		
	Management Interface	<input type="radio"/> Static	
	Management Interface	IP Address / Netmask	/ 1
	Management Interface	Next-hop IP Address	
Management Interface	Domain Name	speak.local	
Management Interface	DNS Primary Server	10.0.233.70	
Management Interface	DNS Secondary Server		

Buttons: Previous, Next, Apply

Proxy Configuration

Orchestrator Administration > [General] Proxy Configuration

If necessary (for example, because of firewall issues), you can configure a proxy for reaching the Silver Peak portal.



The image shows a dialog box titled "Proxy Configuration" with a close button (X) in the top right corner. The dialog contains the following fields:

- Use Proxy:** A checkbox that is currently unchecked.
- Proxy Host:** A text input field.
- Proxy Port:** A text input field.
- Proxy User Name:** A text input field with the label "Optional" to its right.
- Proxy Password:** A text input field with the label "Optional" to its right.

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

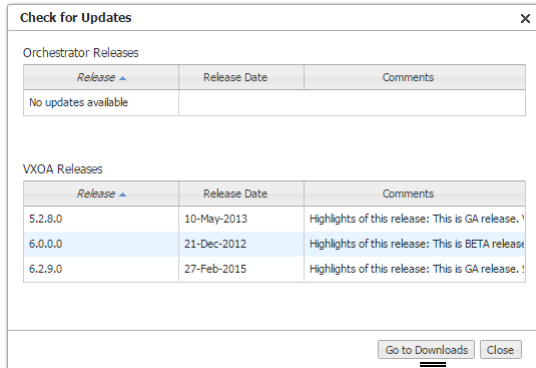
Managing Orchestrator Software

Using these screens, you can check for updated software images, upgrade the Orchestrator server software, switch to another Orchestrator software partition, backup Orchestrator software, and schedule a backup of the Orchestrator.

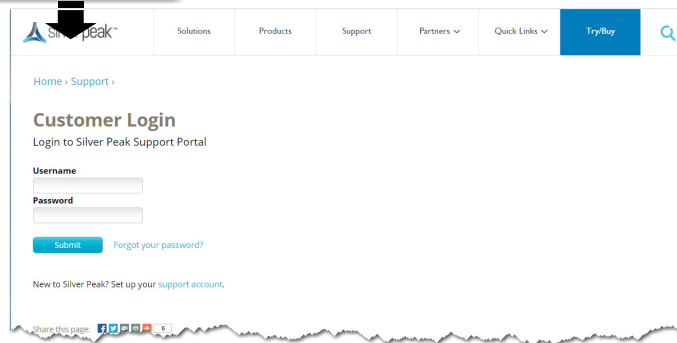
Checking for Orchestrator and Appliance Software Updates

Orchestrator Administration > [Software Management] Check for Updates

Use these screens to see what appliance and Orchestrator server software is available for download.



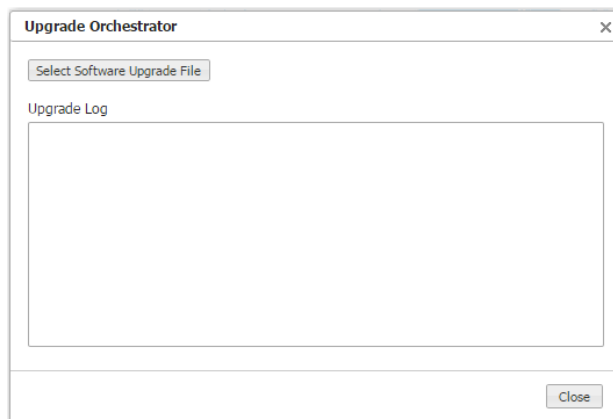
Go to Downloads takes you to the login page of the Support portal.



Upgrading Orchestrator Software

Orchestrator Administration > [Software Management] Upgrade Orchestrator Software

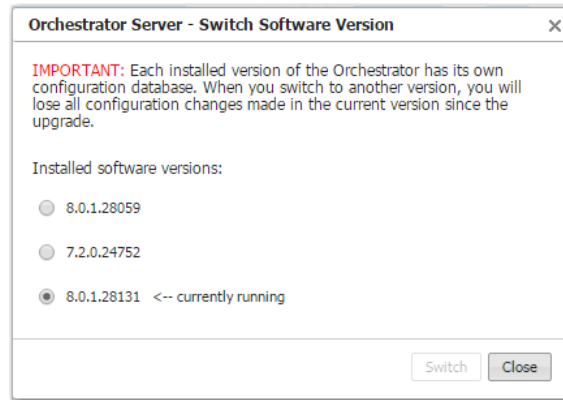
Use this screen to navigate to the file and monitor the upgrade progress.



Switching Software Versions

Orchestrator Administration > [Software Management] Switch Software Version

Each version of Orchestrator has its own separate database. If you switch to another version, then you only have access to the configuration that existed at the point you upgraded from that version.



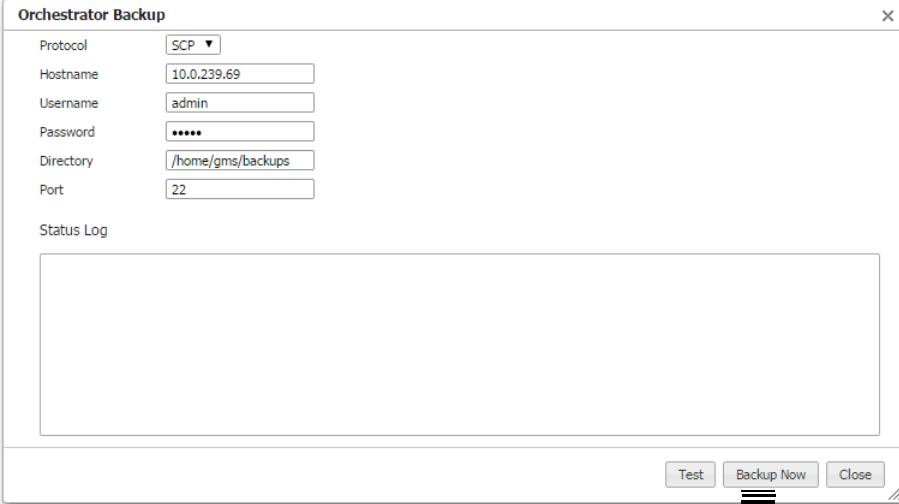
Backing Up the Orchestrator Database

Using these screens, you can backup the Orchestrator database immediately or schedule a backup.

Backing Up on Demand

Orchestrator Administration > [Software Management] Backup Now

Use this screen to backup the Orchestrator database on demand.

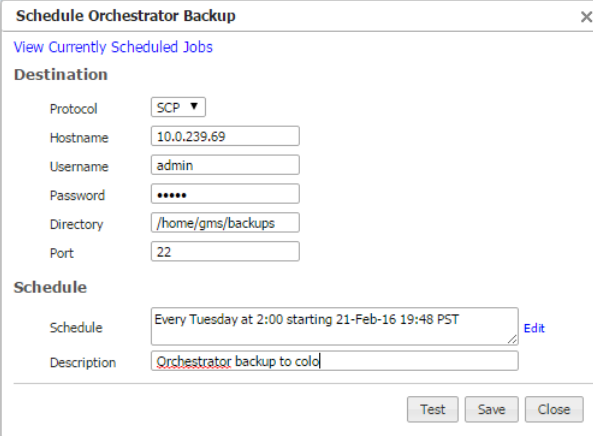


SCP Connection to Host: 10.0.239.69, Port No: 22, Directory: /home/gms/backups with username: admin was successful.

Scheduling Orchestrator Database Backup

Orchestrator Administration > [Software Management] Schedule Backup

Use this screen to schedule a backup the Orchestrator database.



Tip To specify the timezone for scheduled jobs and reports, go to **Orchestrator Administration** > [General] Schedule Timezone.



Maintenance and Support

This chapter describes operations related to appliance maintenance and support.

In This Chapter

- **Viewing System Information** See page 230.
- **Software Versions** See page 231.
- **Upgrading Appliance Software** See page 232.
- **Backing Up Appliance Configuration Files** See page 233.
- **Restoring a Backup to an Appliance** See page 234.
- **Viewing Configuration History** See page 235.
- **Disk Management** See page 236.
- **Synchronizing Appliance Configuration** See page 237.
- **Putting the Appliance in System Bypass Mode** See page 238.
- **Broadcasting CLI Commands** See page 239.
- **Testing Link Integrity** See page 240.
- **Erasing Network Memory** See page 245.
- **Rebooting or Shutting Down an Appliance** See page 246.
- **Scheduling an Appliance Reboot** See page 247.
- **Scheduling QoS Map Activation** See page 248.
- **Managing Tech Support Files** See page 249.
- **Logging in to the Support Portal** See page 251.

Viewing System Information

Maintenance > [Software & System Management] System Information

The **System Information** tab lists the appliances with their relevant information. It has a **Basic** view and an **Advanced** view.

Basic

Appliance Name	Appliance Model	Hardware Version	Bios Version	Appliance ID	Serial Number	System Bandwidth	Deploy Mode	Active S/W Release	UpTime
Tallinn	NX-3700	200400-001 Rev F	200415-002 ...	164926	001BBC02843E	20000	bridge	8.0.2.0_58491	2d 6h 7m 58s
laine-va	VX-1000	206002001000 Rev 46839	6.00	1659809	000C291953A1	3600	bridge	8.0.2.0_58453	2d 6h 47m 14s
laine-vxb	VX-1000	206002001000 Rev 46839	6.00	15046294	000C29E59696	4000	bridge	8.0.2.0_58453	2d 6h 47m 8s
laine2-va	VX-1000	207000000000 Rev 51389	6.00	16715883	000C29FF106B	4000	bridge	8.0.1.0_58416	10d 5h 59m 21s
laine2-vxb	VX-1000	207000000000 Rev 51389	6.00	7104579	000C296C6843	4000	bridge	8.0.1.0_58416	4d 7h 52m 40s

-
- Appliance Name
 - Appliance Model
 - Hardware Version
 - BIOS Version
 - Appliance ID
 - Serial Number
 - System Bandwidth
 - Deploy Mode
 - Active Software Release
 - Uptime
-

Advanced

Edit	Appliance	Optimization							Subnet Sharing			Excess Flow Handling			Miscellaneous		
		Optimize ...	IP Id aut...	TCP auto ...	Automati...	Move flow...	Hair-pin t...	Use share...	Automatic...	Metric for ...	Encrypt d...	Excess flo...	Excess flo...	SSL optim...	Bridge Lo...	Enable Sa...	Enable 1G...
✓	Tallinn	No	Yes	Yes	Yes	fail-stick	No	Yes	No	50	Yes	bypass	Yes	No	Yes	No	Yes
✓	laine-va	No	Yes	Yes	No	fail-stick	No	Yes	Yes	50	Yes	bypass	Yes	No	Yes	Yes	Yes
✓	laine-vxb	No	Yes	Yes	No	fail-stick	No	Yes	Yes	50	Yes	bypass	Yes	No	Yes	Yes	Yes
✓	laine2-va	No	Yes	Yes	No	fail-stick	No	Yes	Yes	50	Yes	bypass	Yes	No	Yes	No	Yes
✓	laine2-vxb	No	Yes	Yes	No	fail-stick	No	Yes	Yes	50	Yes	bypass	Yes	No	Yes	No	Yes

The **Advanced** table includes information specific to optimization, subnet sharing, encrypting data on the disk, excess flow handling, and miscellaneous **System** options.

Software Versions

Maintenance > [Software & System Management] Software Versions

The **Software Versions** tab lists the software installed in each appliance's two partitions.

Software Versions ?

5 Rows Search

Appliance Name	Partition 1				Partition 2			
	Build Version	Build Date	Active	Next Boot	Build Version	Build Date ▲	Active	Next Boot
laine2-vxb	8.0.1.0_58416	2016-02-10 07:50:06	Yes	Yes	7.1.0.0_53424	2014-10-13 17:22:51	No	No
laine-vxa	8.0.2.0_58453	2016-02-17 09:18:34	Yes	Yes	7.3.3.0_57797	2015-12-09 21:14:31	No	No
laine-vxb	8.0.2.0_58453	2016-02-17 09:18:34	Yes	Yes	7.3.3.0_57797	2015-12-09 21:14:31	No	No
laine2-vxa	8.0.1.0_58416	2016-02-10 07:50:06	No	No	8.0.1.0_58416	2016-02-10 07:50:06	Yes	Yes
Tallinn	0.0.0.0_53699	2014-11-18 12:24:03	No	No	8.0.2.0_58491	2016-02-18 17:14:42	Yes	Yes

Upgrading Appliance Software

Maintenance > [Software & System Management] Software Upgrade

You can download and store new appliance software from your network or computer to the Orchestrator server, staging it for installation to the appliance(s).

Use the **Maintenance > Upgrade Appliance Software** page to upload appliance software to the Orchestrator and to install appliance software from the Orchestrator server into the appliance's inactive partition.

Deletes appliance software from the Orchestrator

Displays the appliances selected before opening this window.

Name	Version	Build Date	
image-8.0.2.0_58453.img	8.0.2.0_58453	2016-02-17 09:18:34	X
image-8.0.1.0_58416.img	8.0.1.0_58416	2016-02-10 07:50:06	X
image-6.2.7.0_53789.img	6.2.7.0_53789	2014-12-03 16:08:18	X
image-7.0.0.0_51009.img	7.0.0.0_51009	2014-05-30 18:10:54	X

Appliance	Status	Progress
laine2-vxa	[Slot0: 8.0.1.0_58416], [Slot1: 8.0.1.0_58416, Current, Next Boot]	<input type="text"/>
laine2-vxb	[Slot0: 8.0.1.0_58416, Current, Next Boot], [Slot1: 7.1.0.0_53424]	<input type="text"/>
laine-vxa	[Slot0: 8.0.2.0_58453, Current, Next Boot], [Slot1: 7.3.3.0_57797]	<input type="text"/>
laine-vxb	[Slot0: 8.0.2.0_58453, Current, Next Boot], [Slot1: 7.3.3.0_57797]	<input type="text"/>
Tallinn	[Slot0: 0.0.0.0_53699], [Slot1: 8.0.2.0_58491, Current, Next Boot]	<input type="text"/>

Upload VXOA Image

Upgrade Options

Install and reboot

Install and set next boot partition

Install only

Upgrade Close

For adding new appliance software images to the Orchestrator server.

- **Install and reboot** installs the image into the appliance's inactive partition and then reboots the appliance to begin using the new software.
- **Install only** downloads the image into the inactive partition.

Backing Up Appliance Configuration Files

Maintenance > [Software & System Management] Backup Now

The Orchestrator automatically creates a weekly backup of each appliance's configuration to the Orchestrator server. Additionally, you can create an immediate backup on demand.

After selecting the appliance(s), go to **Maintenance > Backup Now**.

Appliance Backup

Comment

Search

Mgmt IP ▲	Appliance	Status	Duration (Sec)	Details
10.0.233.196	dall	Not started		
10.0.233.197	falcon	Not started		
10.0.238.135	Seattle-EC	Not started		
10.0.238.136	SanFran-EC	Not started		
10.0.238.181	Denver-EC	Not started		

Backup Close

Appliance Backup

Comment

Search

Mgmt IP ▲	Appliance	Status	Duration (Sec)	Details
10.0.233.197	falcon	Completed	7.3	Backup for Appliance 10.0.233.197 Complet...
10.0.233.196	dall	Completed	7.4	Backup for Appliance 10.0.233.196 Complet...
10.0.238.135	Seattle-EC	Completed	6.1	Backup for Appliance 10.0.238.135 Complet...
10.0.238.136	SanFran-EC	Completed	5.7	Backup for Appliance 10.0.238.136 Complet...
10.0.238.181	Denver-EC	Completed	5.6	Backup for Appliance 10.0.238.181 Complet...

Backup Close

You cannot delete an appliance backup from the Orchestrator.

Restoring a Backup to an Appliance

Maintenance > [Software & System Management] Restore

- You can restore a configuration backup from the Orchestrator to an individual appliance.
- You **cannot** restore an appliance's backup to a different appliance.

After selecting the appliance, go to **Maintenance > Restore**. Only that appliance's backups display in the table.

The screenshot shows the 'Appliance Restore and Reboot' dialog box in the Silver Peak Systems interface. The left sidebar shows the navigation tree with 'laine-wxa (10.0.238.71)' selected. The main dialog area contains a table of configuration backups and a status log.

File Name	Backup Time	Software Version	File Content	Comment
initial	20-Feb-16 12:00:05	8.0.2.0_58453	View	
initial	19-Feb-16 12:00:04	7.3.3.0_57797	View	
initial	6-Feb-16 12:00:05	7.3.3.0_57797	View	
backup.1427787900472.10.NE	31-Mar-15 00:45:03	6.2.7.0_53789		Weekly Appliance Backup
backup.1425372300056.10.NE	3-Mar-15 00:45:02	6.2.7.0_53789		Weekly Appliance Backup
backup.1424767500055.10.NE	24-Feb-15 00:45:02	6.2.7.0_53789		Weekly Appliance Backup
backup.1424162700074.10.NE	17-Feb-15 00:45:02	6.2.7.0_53789		Weekly Appliance Backup
backup.1423557900060.10.NE	10-Feb-15 00:45:02	6.2.7.0_53789		Weekly Appliance Backup

Below the table is a 'Status Log' area, which is currently empty. At the bottom right of the dialog, there are 'Restore' and 'Close' buttons.

Disk Management

Maintenance > [Software & System Management] Disk Management

The Disk Management page lists information about physical and virtual appliance disks.

- The progress bar shows what percentage of the polling is complete.
- Physical appliances use RAID arrays with encrypted disks.
- Disk failure results in a critical alarm.
- If a row shows that a disk has failed, click Edit to access the appliance, and follow directions in the local help for replacing the failed disk.
- You can view the SMART [Self-Monitoring Analysis and Reporting Technology] data from physical appliance disks.

The screenshot shows the 'Disk Management' page with a table of 10 rows. The first two rows are for 'Tallinn' appliances. The first row has an 'Edit' icon circled in red. A callout box points to this icon with the text: 'For example, to access Tallinn's own Disk Management page, click **any** of these Edit icons.'

The table columns are: Edit, Appliance Name, Appliance Model, Slot ID, Pairing Slot ID, Status, Size(GB), Serial Number, Removable, and SMART Data.

The SMART data pop-up window for 'Tallinn ID 1' shows the following data:

Attribute	Normalized Value	Worst Value	Raw Value
Read Error Rate	83	63	219,710,224
Spin up time	100	100	0
Start/stop count	100	100	133
Reallocated sector count	100	100	11
Seek error rate	80	60	399,502,215
Power on hours	81	81	17,018
Spin retry count	100	100	0
Device power cycle count	100	100	133
End-to-End error	100	100	0
Reported Uncorrectable Errors	100	100	0
Command Timeout	100	98	1,048,613
High Fly Writes	1	1	186
Temperature Difference	76	26	471,007,256

For example, to access Tallinn's own Disk Management page, click **any** of these Edit icons.

Follow this procedure when replacing a failed disk:

- 1 Log into your Support portal account, and click Open a Self Service RMA for disk replacement.
- 2 Complete the wizard, using the serial number of the appliance (not the disk).
- 3 After you receive the new disk, access Appliance Manager by clicking any **Edit** icon that belongs to the appliance in question.
- 4 Follow the instructions in that page's on-line help.

Synchronizing Appliance Configuration

Maintenance > [Tools] Synchronize

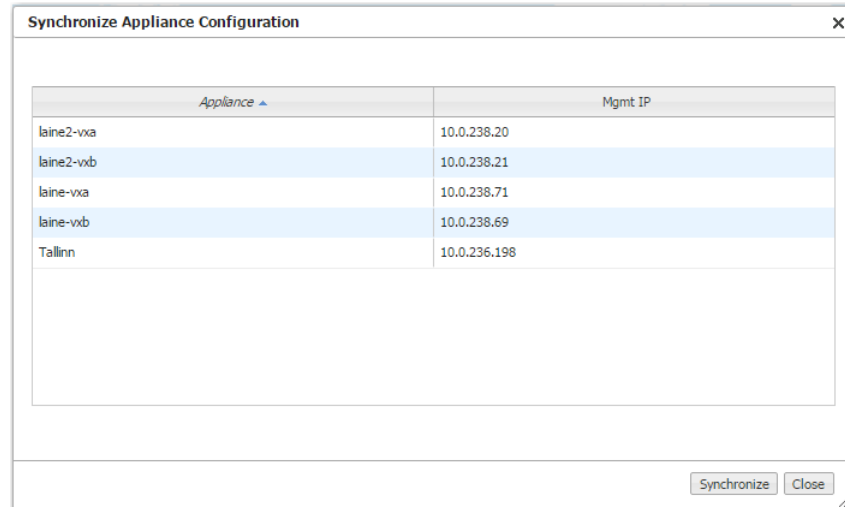
The Orchestrator keeps its database synchronized with the appliances' running configurations.

- When you use Orchestrator to make a configuration change to an appliances' running configuration, the appliance responds by sending an **event** back to the Orchestrator server to log, thereby keeping the Orchestrator and appliance in synch.
- Whenever an appliance starts or reboots, the Orchestrator automatically inventories the appliances to resync.
- Whenever the Orchestrator restarts, it automatically resyncs with the appliances.
- When an appliance is in an **OutOfSync** management state, the Orchestrator server resyncs with it as it comes back online.

If your overall network experiences problems, then you can manually resynch to ensure that the Orchestrator has an appliance's current running configuration.

- ◆ **To manually resync the Orchestrator server with the appliances' configuration database**

Select the appliance(s) and choose **Maintenance > Synchronize**.



Putting the Appliance in System Bypass Mode

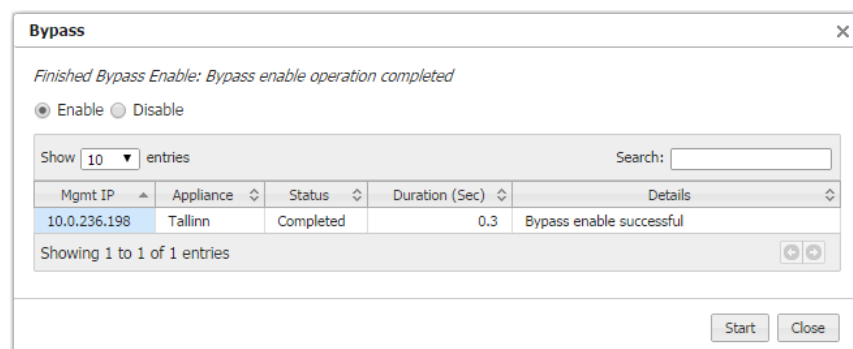
Maintenance > [Tools] Bypass

This only applies to physical appliances.


In **system bypass mode**, the fail-to-wire (or fail-to-glass) card **DOES NOT** receive or process packets:

- In an in-line deployment (Bridge mode), the **lan** interface is physically connected to the **wan** interface.
- In an out-of-path deployment (Router/Server mode), the appliance is in an open-port state.

Fail-to-wire network interfaces mechanically isolate the appliances from the network in the event of a hardware, software, or power failure. This ensures that all traffic bypasses the failed appliance and maximizes up-time.



When the appliance is in Bypass mode, a message displays in red text at the upper right corner of the user interface.

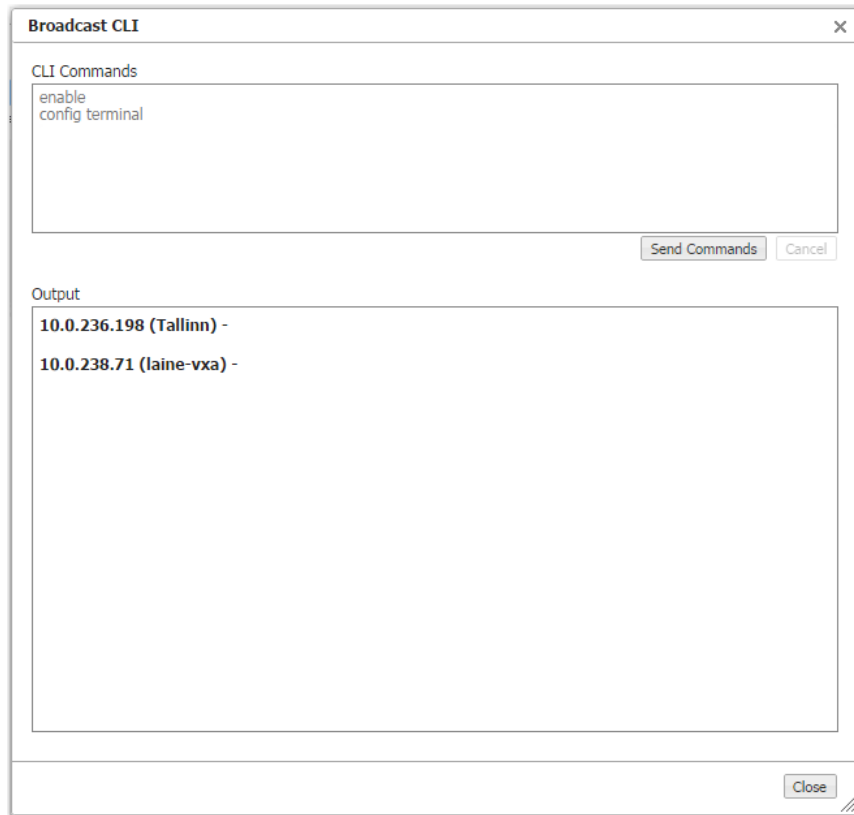


Broadcasting CLI Commands

Maintenance > [Tools] Broadcast CLI

You can simultaneously apply CLI (Command Line Interface) commands to multiple, selected appliances.

The window automatically provides you the highest user privilege level.



For more information, see the *Silver Peak Command Line Interface Reference Guide*.

Testing Link Integrity

Maintenance > [Tools] Link Integrity Test

Used for debugging, the **link integrity** test lets you measure the throughput and integrity (amount of loss) of your WAN link. You can run either **iperf** or **tcpperf** (Version 1.4.8).

The screenshot shows the 'Link Integrity Test' configuration window. On the left, the 'laine-vxb' panel shows a server listening on UDP port 5555. On the right, the 'laine-vxa' panel shows a client connecting to 10.1.154.20 on UDP port 5555. The central configuration area has the following settings: Bandwidth (laine-vxb to laine-vxa) is 1000 Kbps; Bandwidth (laine-vxa to laine-vxb) is 1000 Kbps; Duration (seconds) is 10; DSCP is 'any'; Mode is 'pass-through'; Test Program is 'iperf'; and Custom Parameters is empty. At the bottom of the configuration area, there are 'Start' and 'Clear' buttons. A line points from the text below to the 'Start' button.

The **Start** and **Stop** buttons are colocated.

- These tests run on the two selected appliances, using user-specified parameters for bandwidth, duration, DSCP marking, and type of traffic (tunnelized / pass-through-shaped / pass-through-unshaped).
- The Orchestrator runs the selected test twice -- once passing traffic from Appliance A to Appliance B, and the second run passing traffic from Appliance B to Appliance A.
- **Custom Parameters** are available for **tcpperf** and should be used cautiously, by advanced users.

TCPPERF Version 1.4.8

Basic Mode

Option	Explanation
-h	<i>help</i>
-s	<i>server</i> : Run tcpperf in server mode (not applicable for file generation). Listens on TCP port 2153 by default. [server_port [server_port [server_port]..]]
-sr	server range: <server_port_start:server_port_end>
-c	<i>client server_IP</i> : TCPperf Server's IP address (not applicable for file generation). [server_port [server_port [server_port]..]]
-cr	<server_port_start:server_port_end> <server_port_start:server_port_end>
-g	<i>generate basefilename</i> . Dump generated data to a file.
-sw	<i>sgwrite conffilename</i>

Notes:

- 1 The default server ports are 2153 and 2154.
- 2 You can specify multiple odd-numbered server ports.
- 3 The next even-numbered server ports will also be assigned automatically.
- 4 These even numbers are reserved for double connection testing (see **-I**, *interface IP*).
- 5 Generate mode generates a local file per flow with the same content that the client would have generated with the specified parameters.
- 6 SG write mode is like generate mode except that it writes to an SG device.

General Parameters

Option	Explanation
-6	<i>ip6</i> . Forces tcpperf to use IPv6 addresses only. Default is IPv4 addresses.
-I	<i>interface IP</i> : Specify source interface IP address. Default is any .
-o	<i>outname</i> : Output filename. Default is stdout .
-u	<i>update <secs></i> : Frequency of printed updates in seconds. Default is 1 .
-d	<i>duration <secs></i> : Set maximum test duration in seconds. Default is infinite .
-w	<i>wait <secs></i> : Wait until <secs> since 1970 before transmitting data.
-z	<i>realtime</i> : Elevate to realtime priority. Requires root privilege.
-cm	<i>cpu mask</i> : Specify CPU affinity. Requires root privilege.
-q	<i>quiet <level></i> : Suppresses detail based on level: 0 - None. Print results when test is complete. 1 - Default. Periodic packet/byte statistics. 2 - Verbose. Adds connection state changes. 3 - Debug. Prints everything.

TCP Parameters

Option	Explanation
-tw	<i>tcpwindow</i> : TCP window_size. Default is OS default.
-tm	<i>tcpmss</i> : TCP mss. Default is OS default.
-tn	<i>tcpnodelay</i> : TCP nodelay option. Default is nagle enabled.
-tq	<i>tcpquickack</i> : TCP quick ack option. Default is delayed acks.
-td	<i>tcpdscp</i> <cp>: Sets IP DSCP to <cp> (decimal). Default is 0.
-tr	<i>tcpretries</i> <n>: Sets number of times to retry TCP connections.
-tp	<i>tcppace</i> <n> [mode]: Pace TCP connection setup rate. Limits number of half-open connections to <n>. Valid <mode> types are: preestablish . All connections are established prior to data transmission. Default. simultaneous . Begin data transmission as soon as connection made
-ta	<i>tcpabort</i> : Sends RSTs instead of FINs on close.
-tf	<i>tcpfindelay</i> <secs>: Time to wait after all data sent before sending FIN/RST

Traffic Generation Parameters

Option	Explanation
-f	<i>file</i> . Source filename to load. Default is 10MB of random data.
-i	<i>test id</i> <i>: Set test ID. The same test ID produces the same data set. User different test IDs to generate unique data for each test run. Default is zero.
-n	<i>number</i> <n>: Generate <n> flows. Default is one.
-b	<i>begin</i> <byte>: First byte in transmission. Default is zero.
-e	<i>end</i> <byte>: End byte in transmission (number of bytes to transmit). Default is file size. Begin and end bytes can be greater than file size. The content is repeated to create extra bytes.
-a	<i>antipat</i> <mode>: Antipattern mode: default is mutate: none . Repeats same content verbatim on all flows. Repeats content if end byte exceeds content size. mutate . Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Destroys cross flow similarity. Destroys original byte code distribution. shuffle . Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Preserves cross flow similarity. Preserves original byte code distribution. fast . Ensures all flows and data repeats are unique. Does not preserve short range patterns. Destroys cross flow similarity. Destroys original byte code distribution. Uses less CPU than mutate or shuffle.

Option	Explanation (Continued)
-l	<p><i>loopback [mode]</i>: Loopback. Default is unidirectional.</p> <p>uni. Unidirectional client to server.</p> <p>rev. Unidirectional server to client.</p> <p>bidir. Bidirectional, client and server independently send data on the same TCP connection.</p> <p>bidir2. Bidirectional, client and server independently send data on secondary TCP connections.</p> <p>loop. Bidirectional, server loops data back to client on the same TCP connection.</p> <p>loop2. Bidirectional, server loops data back to client on a secondary TCP connection.</p> <p>bidir2. Bidirectional, transmits one transaction at a time. Client waits for previous transaction to be echoed. Emulates transactional data.</p> <p>NOTES:</p> <p>1. Content source for traffic originating at the server is determined by the server (not client) command line.</p> <p>2. loop2 and bidir2 modes 2 x <n> TCP connections and requires that the server has even-numbered ports available.</p>
-r	<p><i>rate <bps></i>: Limits aggregate transmission rate to <bps>. Default is no rate limit.</p>
-t	<p><i>trans <min> [max]</i>: Sets size of each socket transaction. Default is 64000.</p> <p>If <min> and <max> are specified, client generates transactions with random sizes between <min> and <max>. This feature is often used with -l and -r. Set the minimum transaction size to 100000 to improve single-flow performance.</p>
-v	<p><i>verify <mode></i>: Verify integrity of received data. Default is global.</p> <p>none. No verification. Fastest/least CPU load.</p> <p>global. Single global hash per flow. Fast, but cannot isolate an errored block.</p> <p>literal. Literal comparison of data upon reception. Fast, can isolate errors to the byte level. Requires that server has same content as client. Use random data gen or same -ffile at server.</p> <p>embedded. Embedded hashes every 4096 bytes. Slower, can isolate errors to 4096 byte block.</p>
-p	<p><i>repeat <n></i>: Repeat each content byte n times. Default is 1 (no repeats).</p> <p>Works for both random data and file content.</p>
-k	<p><i>corrupt <n> <m> <s> [<%change>[<%insert>[<%delete>]]]</i>: Corrupt 0 to n bytes of data every m bytes using seed s. Delta bytes will require 0.5*n/m percent overhead. Each corrupt may be a change, insert or delete with the probability of each being specifiable. The default is 33.3% changes, 33.3% inserts, and 33/3% deletes.</p>
-x	<p><i>excerpts <e> <l> [s]</i>: Send random excerpts of average <l> length bytes from content between egin and <e>nd bytes. The -b and -e options still specify total bytes to send. Uses random seed s.</p>

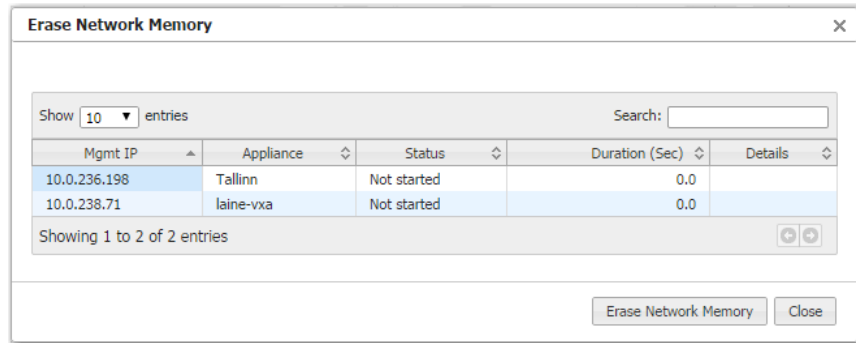
Option	Explanation (Continued)
-y	<p data-bbox="618 237 1414 289"><i>defred</i> <<i>s%</i>> <<i>m%</i>> <<i>l%</i>> <<i>sb</i>> <<i>smin</i>> <<i>smax</i>> <<i>mb</i>> <<i>mmin</i>> <<i>mmax</i>> <<i>lb</i>> <<i>lmin</i> <i>lmax</i>> :</p> <p data-bbox="618 302 1101 323">Generate content based on defined reduction model.</p> <p data-bbox="618 342 1073 363">Content is drawn from three data sets: s, m, and l:</p> <ul data-bbox="646 382 1328 485" style="list-style-type: none"> <li data-bbox="646 382 1328 403">s% specifies fraction [50%] of s-type content (short term reducible). <li data-bbox="646 422 1328 443">m% specifies fraction [30%] of m-type content (medium term reducible). <li data-bbox="646 462 1328 485">l% specifies fraction [20%] of l-type content (long term reducible). <p data-bbox="618 541 1414 594">Short term content comes from data set of <i>sb</i> Mbytes [100MB] with excerpts uniformly distributed between <i>smin</i> and <i>smax</i> bytes [10K-1M].</p> <p data-bbox="618 613 1414 665">Medium term content comes from data set of <i>mb</i> Mbytes [100GB] with excerpts uniformly distributed between <i>lmin</i> and <i>lmax</i> bytes [10K-1M].</p> <p data-bbox="618 684 1414 737">Long term content comes from data set of <i>lb</i> Mbytes [100TB] with excerpts uniformly distributed between <i>smin</i> and <i>smax</i> bytes [10K-1M].</p> <p data-bbox="618 793 1101 814">The -b and -e options still specify total bytes to send.</p> <p data-bbox="618 833 889 854">Performance is best if -b is 0.</p> <p data-bbox="618 873 808 894">Uses random seed <i>s</i>.</p>
-ssl [<i>param=value ...</i>]	<p data-bbox="618 919 1101 940">Enable SSL on connection with optional parameters.</p> <p data-bbox="618 959 1073 980">version=2 3 t10 t11 t12. Set the protocol version.</p> <p data-bbox="618 999 1214 1020">cipher=OPENSSSL-CIPHER-DESC. Set the choice of ciphers.</p> <p data-bbox="618 1039 1127 1060">ticket=yes no. Enable/disable session ticket extension.</p> <p data-bbox="618 1079 1024 1100">cert=FILENAME. Use this certificate file.</p> <p data-bbox="618 1119 1024 1140">key=FILENAME. Use this private keyfile.</p> <p data-bbox="618 1159 1268 1180">compression=none any deflate zlib rl. Set the compression method.</p> <p data-bbox="618 1199 1062 1220">sslcert. Print the SSL certificate in PEM format.</p> <p data-bbox="618 1239 1000 1260">sslkey. Print the SSL key in PEM format.</p>

Erasing Network Memory

Maintenance > [Tools] Erase Network Memory

Erasing Network Memory removes all stored local instances of data.


No reboot required.



Rebooting or Shutting Down an Appliance

Maintenance > [Tools] Appliance Reboot / Shutdown

The appliance supports three types of reboot:



- **Reboot.** Reboots the appliance gracefully. This is your typical "vanilla" restart.
Use case: You're changing the deployment mode or other configuration parameters that require a reboot.
- **Erase Network Memory and Reboot.** Erases the Network Memory cache and reboots the appliance.
Use case: You need to restart the appliance with an empty Network Memory cache.
- **Shutdown.** Shuts down the appliance and turns the power off. To restart, go to the appliance and physically turn the power on with the Power switch.

Use case:

- You're decommissioning the appliance.
- You need to physically move the appliance to another location.
- You need to recable the appliance for another type of deployment.

Behavior During Reboot

A *physical appliance* enters into one of the following states:

- *hardware bypass*, if deployed in-line (Bridge mode), or
- *an open-port state*, if deployed out-of-path (Router/Server mode).

Unless a *virtual appliance* is configured for a high availability deployment, all flows are discontinued during reboot.

Scheduling an Appliance Reboot

Maintenance > [Tools] Schedule Appliance Reboot

You can schedule an appliance for any of three types of reboot:

Displays selected appliances

Based on the Orchestrator's clock

Mgmt IP	Appliance Name
10.0.236.198	Tallinn
10.0.238.69	laine-vxb
10.0.238.71	laine-vxa

Options

Operation Reboot Shut Down

Switch partition before reboot

Erase Network Memory and Reboot

Reboot/Shutdown Schedule

Start Date

Description

Schedule Reboot Cancel

- **Reboot.** Reboots the appliance gracefully. This is your typical "vanilla" restart.
Use case: You're changing the deployment mode or other configuration parameters that require a reboot.
- **Erase Network Memory and Reboot.** Erases the Network Memory cache and reboots the appliance.
Use case: You need to restart the appliance with an empty Network Memory cache.
- **Shutdown.** Shuts down the appliance and turns the power off. To restart, go to the appliance and physically turn the power on with the Power switch.

Use case:

- You're decommissioning the appliance.
- You need to physically move the appliance to another location.
- You need to recable the appliance for another type of deployment.

Behavior During Reboot

A *physical appliance* enters into one of the following states:

- *hardware bypass*, if deployed in-line (Bridge mode), or
- *an open-port state*, if deployed out-of-path (Router/Server mode).

Unless a *virtual appliance* is configured for a high availability deployment, all flows are discontinued during reboot.



Tip To specify the timezone for scheduled jobs and reports, go to **Orchestrator Administration > [General] Schedule Timezone**.

Scheduling QoS Map Activation

Maintenance > [Tools] Schedule QoS Map Activation

You can schedule appliances to apply different QoS maps at different times.

Schedule QoSMap Activation

[View Currently Scheduled Jobs](#)

Mgmt IP/Group Name
Release 8.0.2

Add

Map	Schedule	Re-Classify	Description
map1	Every day at 6:00 starting 21-Feb-16 20:38 PST	<input checked="" type="checkbox"/>	primary map
map2	Every day at 20:00 starting 21-Feb-16 20:38 PST	<input checked="" type="checkbox"/>	evening map

Schedule QoSMap Cancel

Before using this dialog, verify the following:

- 1 The desired Template Group has the QoS maps you need.
- 2 You've applied the Template Group to the appliances that you want to schedule.



Tip To specify the timezone for scheduled jobs and reports, go to **Orchestrator Administration > [General] Schedule Timezone**.

Managing Tech Support Files

Support > Tech Support [Create Case, View Logs]

If you have a problem with an appliance, Silver Peak Support may ask you to send them specific debug files for evaluation. Listed under **Help > Tech Support**, these include log files, sysdump files, tech files, snapshots, and tcpdump results.

Filters

Mgmt IP	Appliance Name	File Type	File Name	Last Modified	File Size
10.0.236.198	Tallinn	Logs	messages	Fri, 05 Jun 2015 02:55:51 G...	40.6MB
10.0.236.198	Tallinn	Logs	web_access_log	Fri, 05 Jun 2015 02:55:33 G...	21.9MB
10.0.238.69	laine-vxb	Logs	web_access_log	Fri, 05 Jun 2015 02:41:49 G...	47.5MB
10.0.238.69	laine-vxb	Logs	messages	Fri, 05 Jun 2015 02:41:45 G...	17.1MB
10.0.238.71	laine-vxa	Logs	web_access_log	Fri, 05 Jun 2015 02:41:26 G...	49.9MB
10.0.238.71	laine-vxa	Logs	messages	Fri, 05 Jun 2015 02:41:07 G...	5.0MB
10.0.236.198	Tallinn	Sys Dump	tunbug-20150605.tar	Fri, 05 Jun 2015 02:00:05 G...	71.7KB
10.0.238.71	laine-vxa	Sys Dump	tunbug-20150605.tar	Fri, 05 Jun 2015 02:00:05 G...	71.7KB
10.0.238.69	laine-vxb	Sys Dump	tunbug-20150605.tar	Fri, 05 Jun 2015 02:00:05 G...	71.7KB
10.0.238.69	laine-vxb	Sys Dump	tunbug.lock	Fri, 05 Jun 2015 02:00:00 G...	0B
10.0.236.198	Tallinn	Sys Dump	tunbug.lock	Fri, 05 Jun 2015 02:00:00 G...	0B
10.0.238.71	laine-vxa	Sys Dump	tunbug.lock	Fri, 05 Jun 2015 02:00:00 G...	0B
10.0.238.69	laine-vxb	Logs	auditlog	Fri, 05 Jun 2015 01:01:21 G...	1.9MB
10.0.238.71	laine-vxa	Logs	auditlog	Fri, 05 Jun 2015 01:01:01 G...	2.0MB
10.0.238.71	laine-vxa	Sys Dump	tunbug-20150604.tar.gz	Thu, 04 Jun 2015 23:00:05 ...	505.1KB
10.0.236.198	Tallinn	Sys Dump	tunbug-20150604.tar.gz	Thu, 04 Jun 2015 23:00:05 ...	486.3KB
10.0.238.69	laine-vxb	Sys Dump	tunbug-20150604.tar.gz	Thu, 04 Jun 2015 23:00:05 ...	506.6KB
10.0.236.198	Tallinn	Logs	auditlog	Thu, 04 Jun 2015 17:06:41 ...	2.1MB
10.0.236.198	Tallinn	Logs	messages.1.gz	Thu, 04 Jun 2015 09:58:01 ...	3.4MB

Files you upload to Support must be associated with a Case Number.

- To open a new case, click **Create Case & Upload Diagnostics to Silver Peak**.

The highest priority is **P1**,
and the lowest is **P4**.
The default is **P3**.

Create Case & Upload Diagnostics to Silver Peak ✕

IMPORTANT: You must have a valid Silver Peak support account email address to create a case. An email will be sent to this address confirming that a case has been created.

Silver Peak Support Account Email

Contact Name

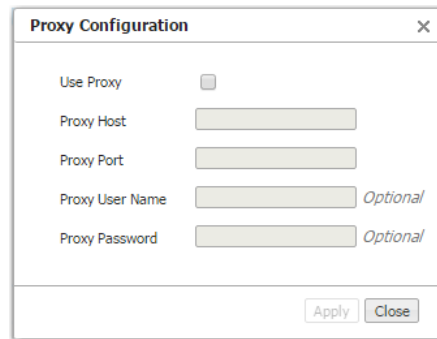
Contact Phone

Case Priority P3 - Normal

Description

This requires you to have a valid Silver Peak Support account email address. An email will be sent to this address, confirming that a case has been created and providing you with a Case Number.

- If you already have a Case Number, you'll be asked to enter it when uploading any additionally requested files.
- All debug files are stored on the appliances themselves. From the table, you can download a file to your computer or upload it to Support.
- You can upload a file from your PC to Support, using the **Advanced Options** menu.
- Although the Orchestrator logs aren't visible to you in the menus, the **Advanced Options** menu lets you upload Orchestrator logs to Support or download them to your computer.
- If necessary (for example, because of firewall issues), you can configure a proxy for uploading files to Silver Peak Support. Go to **Orchestrator Administration > Proxy Settings**.



The image shows a 'Proxy Configuration' dialog box with a close button (X) in the top right corner. It contains the following fields and controls:

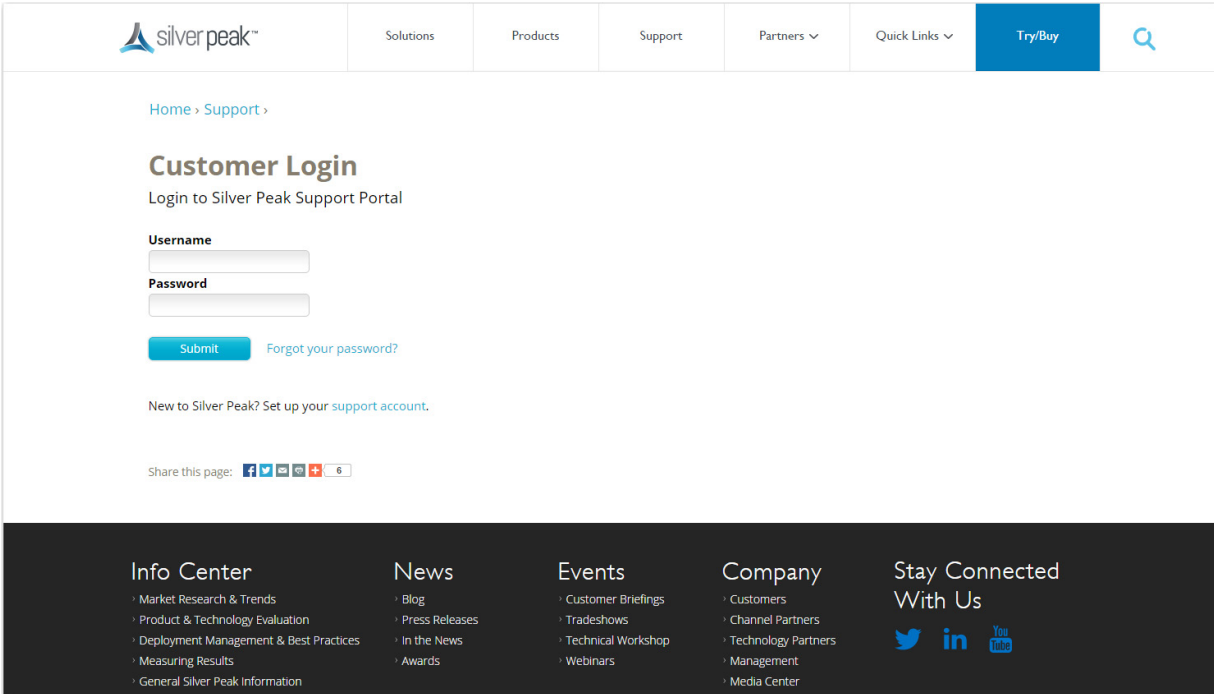
- Use Proxy:** A checkbox that is currently unchecked.
- Proxy Host:** A text input field.
- Proxy Port:** A text input field.
- Proxy User Name:** A text input field with the word 'Optional' to its right.
- Proxy Password:** A text input field with the word 'Optional' to its right.

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

Logging in to the Support Portal

Support > Support Portal Log-in

When you have a Silver Peak account and need technical or customer support, select **Support > Tech Support**. The following page opens in a separate browser tab.



You can also access this page directly by going Silver Peak’s web page and selecting **Support > Customer Login** from the menu bar.



APPENDIX A

TCP/IP Ports Used by the Orchestrator and Silver Peak Appliances

Following are lists of ports that are used by the appliances and by the Orchestrator. These are the ports used for “listening”.

If you intend to use a port, make sure that it is open in the firewall(s).

List of ports used by the Orchestrator

Following is the list of ports used by the Orchestrator. All are part of the management plane.

It is mandatory for certain ports to be open. Opening other ports is optional (opt.), depending on your network, applications, and chosen deployment.

Must open port?	TCP	UDP	Port	Application	Direction relative to the Orchestrator	Comments
yes	x		22	SSH	bidirectional	CLI (Command Line Interface) access over SSH
yes	x		443	HTTPS	bidirectional	communications between the Orchestrator and a physical or virtual appliance
opt.	x		21	FTP	outgoing	for Orchestrator backup This is the default port. If you've configured a different port, then you also need to configure the firewall with that port number.
opt.	x		22	SCP	outgoing	for Orchestrator backup This is the default port. If you've configured a different port, then you also need to configure the firewall with that port number.
opt.	x		49	TACACS+	outgoing	user authentication and authorization
opt.	x	x	53	DNS	outgoing	domain name services
opt.	x		80	HTTP	outgoing	If the appliance's web configuration is for HTTP only , then you must open this port.
opt.		x	123	NTP	outgoing	synchronizes clocks
opt.		x	1812	RADIUS	outgoing	user authentication and authorization

List of ports used by Silver Peak Appliances

Data Plane

This is for packets that traverse the optimization path. For creating tunnels, at least one of the first three applications — GRE, IPsec, or UDP — must be open.

Must open port ?	Application	Ports and Protocols	Use
yes	GRE	Protocol 47	If tunnel mode is GRE
yes	IPsec	Protocol ESP 50; UDP port 500 (for IKE key exchange)	If tunnel mode is IPsec
yes	UDP	UDP Port 4163	If tunnel mode is UDP
yes	ICMP	Protocol 1	Checks reachability of next-hop routers
opt.	flow redirection	TCP Port 4164 and UDP Port 4164	If flow direction is enabled and clustered via routers
opt.	VRRP	Protocol 112	For VRRP protocol messages
opt.	WCCP protocol	UDP Port 2048	For WCCP redirection
opt.	WCCP CRE tunnel	Protocol 47	If L3 WCCP redirection is enabled, then Protocol 47 is used to redirect traffic between WCCP router and VXOA appliance, in both directions.

Management Plane

It is mandatory for certain ports to be open. Opening other ports is optional (opt.), depending on your network, applications, and chosen deployment.

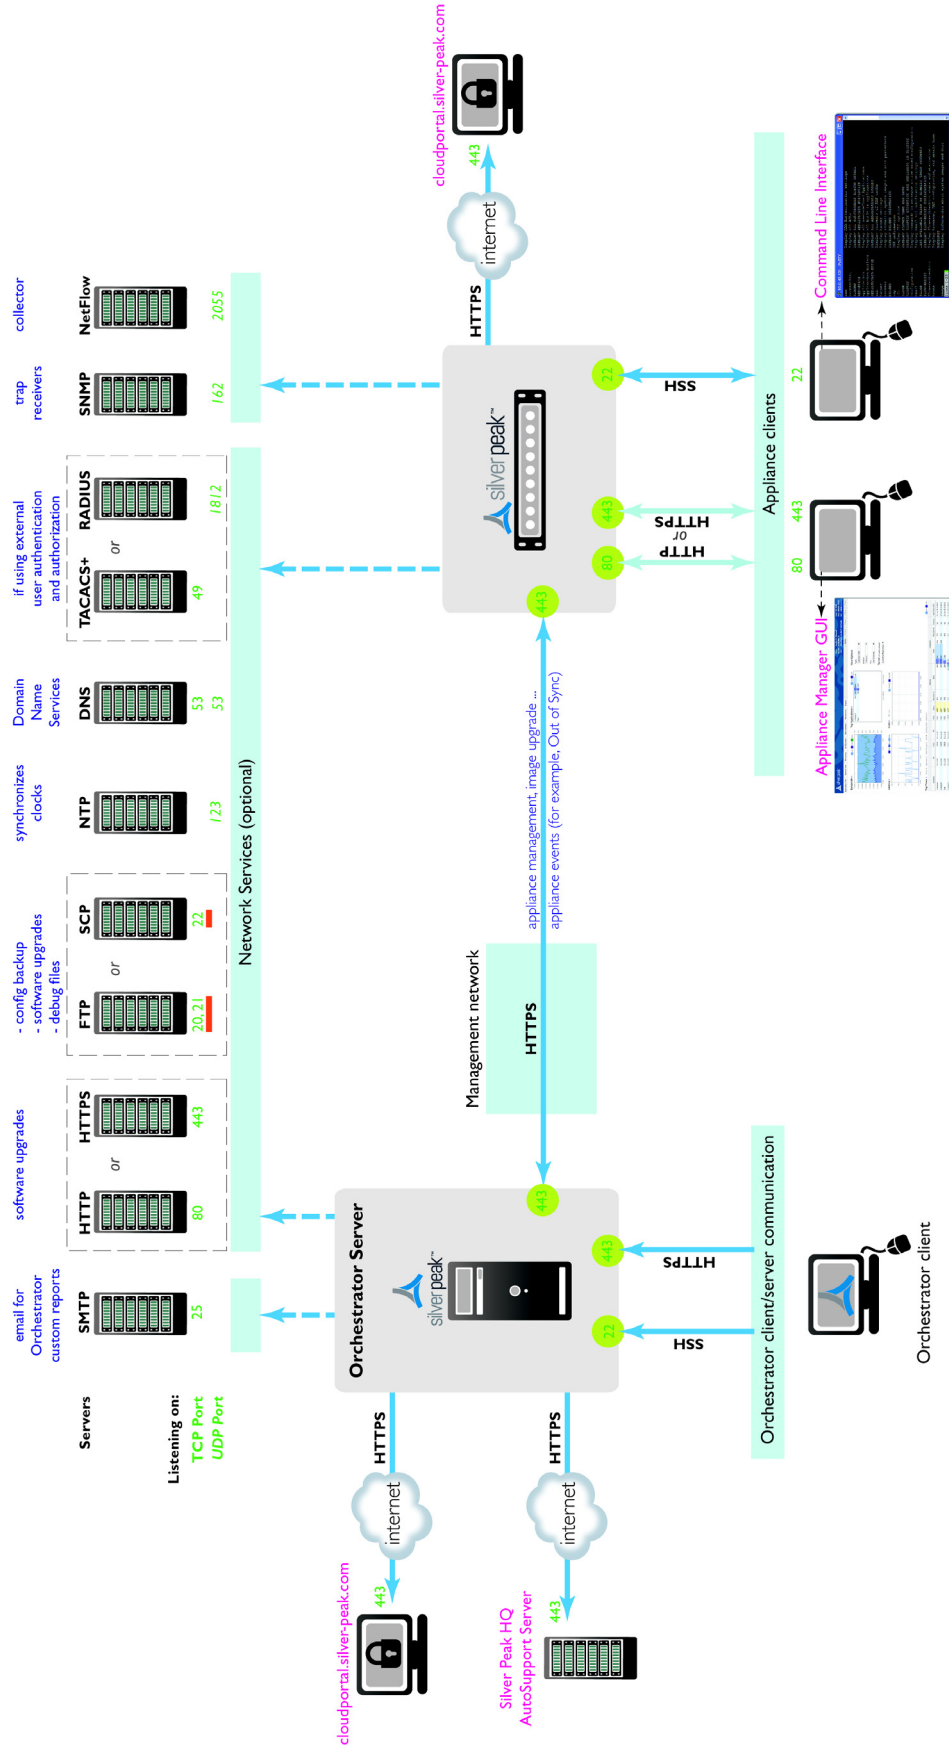
Must open port ?	TCP	UDP	Port	Application	Direction relative to the appliance	Used for ...
yes	x		22	SCP	bidirectional	<ul style="list-style-type: none"> configuration backup software upgrades
yes	x		80	HTTP	bidirectional	communication with VXOA clients and with the Orchestrator
yes	x		443	HTTPS	bidirectional	communication with VXOA clients
opt.	x		20 [data channel] 21 [control channel]	FTP	bidirectional	<ul style="list-style-type: none"> configuration backup software upgrades
opt.	x		49	TACACS+	outgoing	user authentication and authorization
opt.	x	x	53	DNS	outgoing	domain name services
opt.		x	123	NTP	outgoing	synchronizes clocks
opt.		x	1812	RADIUS	outgoing	user authentication and authorization
opt.		x	162	SNMP	outgoing	SNMP trap receivers
opt.		x	2055	Netflow	outgoing	Netflow collector

Diagrams of TCP/IP Port Use

See the following two pages.

TCP/IP ports used for listening by the Orchestrator and Silver Peak appliances
Management Plane

LEGEND
 443 = mandatory port(s)
 21 = underline indicates a user-configurable default port





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

1.877.210.7325
+1.408.935.1850

www.silver-peak.com