



Silver Peak

# Orchestrator Operator's Guide

8.8.1  
September 2019  
200095-001

# Copyright and Trademarks

Silver Peak Orchestrator Operator's Guide

Date: September 2019

Copyright © 2019 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

Silver Peak, the Silver Peak logo, and all Silver Peak product names, logos, and brands are trademarks or registered trademarks of Silver Peak Systems, Inc. in the United States and/or other countries. All trademark rights reserved. All other product names, logos, and brands are property of their respective owners.

## Warranties and Disclaimers

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

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

1.877.210.7325 (toll-free in USA)  
+1.408.935.1850

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

## Support

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

1.877.210.7325 (toll-free in USA)

+1.408.935.1850

[www.silver-peak.com/support](http://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](mailto:techpubs@silver-peak.com).
- If you have comments or feedback about the interface, please send an e-mail to [usability@silver-peak.com](mailto:usability@silver-peak.com).

# Contents

---

Copyright and Trademarks .....	2
Support .....	3
<b>Getting Started .....</b>	<b>13</b>
<b>Overview of SD-WAN Prerequisites .....</b>	<b>14</b>
<b>Unity Overlays .....</b>	<b>18</b>
Business Intent Overlays .....	19
Overview .....	19
SD-WAN traffic to internal subnets .....	19
Building SD-WAN using these interfaces .....	20
Service Level objective .....	20
Link Bonding Policy .....	21
QoS, Security, & Optimization .....	21
Breakout Traffic to Internet & Cloud Services .....	21
Hub versus branch breakout settings .....	21
Preferred Policy Order and Available Policies .....	22
Deployment Profiles .....	23
Mapping Labels to Interfaces .....	23
LAN-side Configuration: DHCP .....	23
WAN-side Configuration .....	24
Definitions .....	26
A More Comprehensive Guide to Basic Deployments .....	27
Bridge Mode .....	27
Router Mode .....	29
Server Mode .....	33
Deployment – EdgeConnect HA .....	35
Enabling EdgeConnect HA Mode .....	35
IPSec over UDP Tunnel Configuration .....	35
VRRP Configuration .....	36
LAN-Side Monitoring .....	36
Interface Labels .....	37
Firewall Zones .....	38
Apply Overlays .....	39
Internet Traffic .....	40
IPSec Pre-shared Key Rotation .....	41
Failure Handling and Orchestrator Reachability .....	41
Hubs .....	43
Discovered Appliances .....	44
Preconfigure Appliances .....	46
Appliance Configuration Wizard .....	48
Licenses .....	51

---

Cloud Portal .....	52
SSL Certificates Tab .....	53
SSL CA Certificates Tab .....	54
SSL for SaaS Tab .....	56
<b>Network Configuration Tabs .....</b>	<b>58</b>
DHCP Failover .....	59
DHCP Failover State .....	60
Regions .....	61
Regional Routing .....	61
View Status .....	62
Edit Regions .....	63
BGP Tab .....	64
BGP Edit Row .....	68
Virtual Tunnel Interface .....	71
VTI .....	72
Boost .....	73
Deployment Tab .....	74
Interfaces Tab .....	77
Terminology .....	78
Routes Tab .....	79
OSPF Tab .....	83
Enabling OSPF .....	83
Adding an Interface .....	86
Viewing the OSPF Neighbors .....	89
Multicast .....	91
Loopback .....	93
Peer Priority Tab .....	94
Admin Distance Tab .....	95
Management Routes Tab .....	96
Import .....	97
VRRP Tab .....	98
VRRP Tab Settings .....	98
WCCP Tab .....	100
WCCP Settings .....	101
Service Group Advanced Settings .....	102
PPPoE Tab .....	104
DHCP Server Defaults .....	107
DHCP Settings .....	107
DHCP/BOOTP Relay Fields .....	108
DHCP Leases .....	109
Tunnels Tab .....	110
Troubleshooting .....	112
Using Passthrough Tunnels .....	112
Tunnel Groups Tab .....	114
Topology .....	115
Interfaces .....	115
Tunnel Exception .....	116

---

Schedule Auto MTU Discovery .....	117
Zscaler Internet Access .....	118
Enabling Zscaler .....	121
Verification .....	121
NAT .....	123
NAT Rules and Pools .....	124
NAT Pools .....	124
<b>Policy Configuration Tabs .....</b>	<b>126</b>
DNS Proxy Policies .....	127
DNS Proxy Policies .....	127
Route Policies Tab .....	129
Priority .....	130
Match Criteria .....	130
Source or Destination .....	130
Wildcard-based Prefix Matching .....	131
QoS Policies Tab .....	132
Handling and Marking DSCP Packets .....	133
Applying DSCP Markings to Optimized (Tunnelized) Traffic .....	133
Applying DSCP Markings to Pass-through Traffic .....	135
Priority .....	137
Match Criteria .....	137
Source or Destination .....	138
Wildcard-based Prefix Matching .....	138
Schedule QoS Map Activation .....	139
Optimization Policies Tab .....	140
Priority .....	140
Match Criteria .....	141
Source or Destination .....	141
Wildcard-based Prefix Matching .....	141
Set Actions .....	142
TCP Acceleration Options .....	143
NAT Policies Tab .....	147
Advanced Settings .....	149
Match Criteria .....	149
Source or Destination .....	150
Wildcard-based Prefix Matching .....	150
Set Actions .....	150
Merge / Replace .....	151
Inbound Port Forwarding .....	152
Security Policies Tab .....	154
Wildcard-based Prefix Matching .....	154
Access Lists Tab .....	156
Match Criteria .....	157
Wildcard-based Prefix Matching .....	157
Shaper Tab .....	158
SaaS Optimization Tab .....	161
Configuration Tab .....	161

---

Monitoring Tab .....	162
Application Definitions .....	163
Application Groups Tab .....	165
Threshold Crossing Alerts Tab .....	166
ON by default: .....	167
OFF by default: .....	168
IP SLA Tab .....	170
Monitor Use Cases .....	171
<b>Configuration Templates .....</b>	<b>181</b>
Using Configuration Templates .....	182
System Template .....	183
Auth/Radius/TACACS+ Template .....	190
Authentication and Authorization .....	190
Appliance-based User Database .....	190
RADIUS .....	191
TACACS+ .....	191
What Silver Peak recommends .....	191
SNMP Template .....	192
Flow Export Template .....	195
DNS Proxy Policies .....	197
DNS Template .....	198
DHCP Failover State .....	199
DHCP Failover .....	200
Logging Template .....	201
Minimum Severity Levels .....	201
Configuring Remote Logging .....	202
Banner Messages Template .....	203
HTTPS Certificate Template .....	204
User Management Template .....	206
Default User Accounts .....	206
Command Line Interface privileges .....	207
Date/Time Template .....	208
Data Collection .....	209
SSL Certificates Template .....	210
SSL CA Certificates Template .....	212
SSL for SaaS Template .....	213
Tunnels Template .....	215
VRRP Template .....	218
Peer Priority Template .....	220
Admin Distance Template .....	221
Shaper Template .....	223
Dynamic Rate Control .....	224
QoS Policies Template .....	226
Priority .....	226
Match Criteria .....	227
Source or Destination .....	227
Wildcard-based Prefix Matching .....	227

---

Handling and Marking DSCP Packets .....	228
Applying DSCP Markings to Optimized (Tunnelized) Traffic .....	228
Applying DSCP Markings to Pass-through Traffic .....	230
Optimization Policies Template .....	233
Priority .....	233
Match Criteria .....	233
Source or Destination .....	234
Wildcard-based Prefix Matching .....	234
Set Actions Fields .....	235
Route Policies Template .....	236
Why? .....	237
Priority .....	237
Match Criteria .....	237
Source or Destination .....	237
Wildcard-based Prefix Matching .....	238
Set Actions Fields .....	238
Where the appliance directs traffic .....	238
How traffic is managed if a tunnel is down .....	239
NAT Policies Template .....	240
When to NAT .....	240
Advanced Settings .....	242
Match Criteria .....	242
Source or Destination .....	243
Wildcard-based Prefix Matching .....	243
Set Actions .....	243
Merge / Replace .....	244
Threshold Crossing Alerts Template .....	245
ON by default: .....	246
OFF by default: .....	246
TCA Metrics .....	247
SaaS Optimization Template .....	249
TIPS .....	251
Security Policies Template .....	252
Implicit Drop Logging .....	252
Template .....	252
Wildcard-based Prefix Matching .....	253
CLI Template .....	254
Session Management Template .....	255
Apply Template Groups .....	256
<b>Monitoring Status and Performance .....</b>	<b>257</b>
Dashboard .....	258
Topology Settings & Legend .....	259
Viewing Tunnels in the Topology Map .....	262
Live View .....	263
Historical Charts .....	264
Health Map .....	266
Alarms Tab .....	268

---

Disable Alarms .....	268
Alarm Severity .....	269
Alarm Recipients .....	269
Additional Alarm Indications .....	270
Configuring & Distributing Custom Reports .....	271
View Reports .....	273
Sample Report .....	274
Scheduled & Historical Jobs .....	275
Appliance Bandwidth .....	276
Appliance Max Bandwidth .....	277
Appliance Bandwidth Utilization .....	278
Appliance Bandwidth Trends .....	279
Appliance Packet Counts .....	280
Appliance Bandwidth Cost Savings .....	281
Application Bandwidth .....	282
Application Pie Charts .....	283
Application Trends .....	284
Firewall Drops .....	285
Top Talkers .....	286
Domains .....	288
Countries .....	289
Ports .....	290
Traffic Behavior .....	291
Overlay-Interface-Transport .....	293
Interface Bandwidth Trends .....	295
Interface Summary .....	296
Tunnels Bandwidth .....	297
Show Underlays .....	297
Traceroute .....	298
Live View .....	298
Tunnels Pie Charts .....	300
Tunnel Bandwidth Trends .....	301
Tunnel Packet Counts .....	302
DRC Bandwidth Trends .....	303
Dynamic Rate Control .....	303
Flows - Active & Recent .....	305
Reset or Reclassify Flows .....	306
Appliance Flow Counts .....	309
Appliance Flow Trends .....	310
Tunnel Flow Counts .....	311
DSCP Bandwidth .....	312
DSCP Pie Charts .....	313
DSCP Trends .....	314
Traffic Class Bandwidth .....	315
Traffic Class Pie Charts .....	316
QoS (Shaper) Trends .....	317
Works with Office 365 .....	318
Live View .....	319

---

Loss .....	320
Loss Trends .....	321
Jitter Summary .....	323
Jitter Trends .....	324
Latency .....	326
Latency Trends .....	327
Out of Order Packets .....	329
Mean Opinion Score (MOS) - Summary .....	332
Mean Opinion Score (MOS) Trends .....	333
Tunnels Summary .....	334
<b>Appliance Administration Tabs .....</b>	<b>335</b>
Appliance User Accounts Tab .....	336
Auth/RADIUS/TACACS+ Tab .....	338
Authentication and Authorization .....	338
RADIUS and TACACS+ .....	338
Date/Time Tab .....	340
DNS (Domain Name Servers) Tab .....	341
SNMP Tab .....	342
Flow Export Tab .....	344
Logging Tab .....	345
Severity Levels .....	345
Remote Logging .....	346
Banners Tab .....	347
HTTPS Certificate Tab .....	348
Orchestrator Reachability Tab .....	349
System Information .....	350
Software Versions .....	356
Upgrading Appliance Software .....	357
Appliance Configuration Backup .....	359
Viewing Configuration History .....	362
Restoring a Backup to an Appliance .....	364
Remove Appliance from Orchestrator .....	365
Remove Appliance from Orchestrator and Account .....	366
Synchronizing Appliance Configuration .....	367
Putting the Appliance in System Bypass Mode .....	368
Broadcasting CLI Commands .....	370
Link Integrity Test .....	372
TCPERF Version 1.4.8 .....	373
Disk Management .....	377
Erasing Network Memory .....	379
Rebooting or Shutting Down an Appliance .....	380
Behavior During Reboot .....	381
Scheduling an Appliance Reboot .....	382
Behavior During Reboot .....	383
Active Sessions Tab .....	385

---

<b>Orchestrator Administration</b> .....	<b>386</b>
Role Based Access Control .....	387
Assign Roles & Appliance Access .....	387
Roles .....	388
Appliance Access .....	388
Viewing Orchestrator Server Information .....	390
Restart, Reboot, or Shutdown .....	391
Managing Orchestrator Users .....	392
Adding a User .....	392
Multi-Factor Authentication .....	393
Configuring Multi-Factor Authentication through an Application .....	393
Configuring Multi-Factor Authentication through Email .....	394
Using Multi-Factor Authentication .....	394
Modify User .....	395
User Menu Access .....	398
Remote Authentication .....	401
Cloud Portal .....	403
Orchestration Settings .....	404
Audit Logs .....	405
Pause Orchestration List .....	407
Tunnel Settings Tab .....	408
Orchestrator Blueprint Export .....	412
Brand Customization .....	413
Maintenance Mode .....	414
Upgrading Orchestrator Software .....	415
Checking for Orchestrator and Appliance Software Updates .....	416
Backing Up on Demand .....	417
Scheduling Orchestrator Database Backup .....	418
SMTP Server Settings .....	419
Proxy Configuration .....	420
Orchestrator's HTTPS Certificate .....	421
Timezone for Scheduled Jobs .....	423
Orchestrator Statistics Configuration .....	424
Appliance Statistics Configuration .....	425
Orchestrator Advanced Properties .....	426
Changing Orchestrator's Log Level .....	428
Minimum Severity Levels .....	428
IP Whitelist .....	429
Orchestrator's Getting Started Wizard .....	430
<b>Customer and Technical Support</b> .....	<b>432</b>
Tech Support - Appliances .....	433
Tech Support - Orchestrator .....	434
Logging into the Support Portal .....	435
Monitoring Uploads .....	436
Packet Capture .....	437
Upload Local Files .....	438
Create a Support Case .....	439

---

Remote Access .....	440
Partition Management .....	441
Remote Log Receivers .....	442
HTTP Receiver Settings .....	442
HTTPS Receiver Settings .....	443
KAFKA Receiver Settings .....	443
SYSLOG .....	443
Routing Peers Table .....	445
RMA Wizard .....	446
Built-in Policies .....	447
Realtime Charts .....	448
Historical Charts .....	450
Appliance Charts .....	451
Internal Drop Trends .....	453
Appliance Memory Trends .....	454
System Performance .....	456
Appliance Crash Report .....	457
Orchestrator Debug .....	458
IPSec UDP Status .....	459
Unverified Emails .....	460

# Getting Started

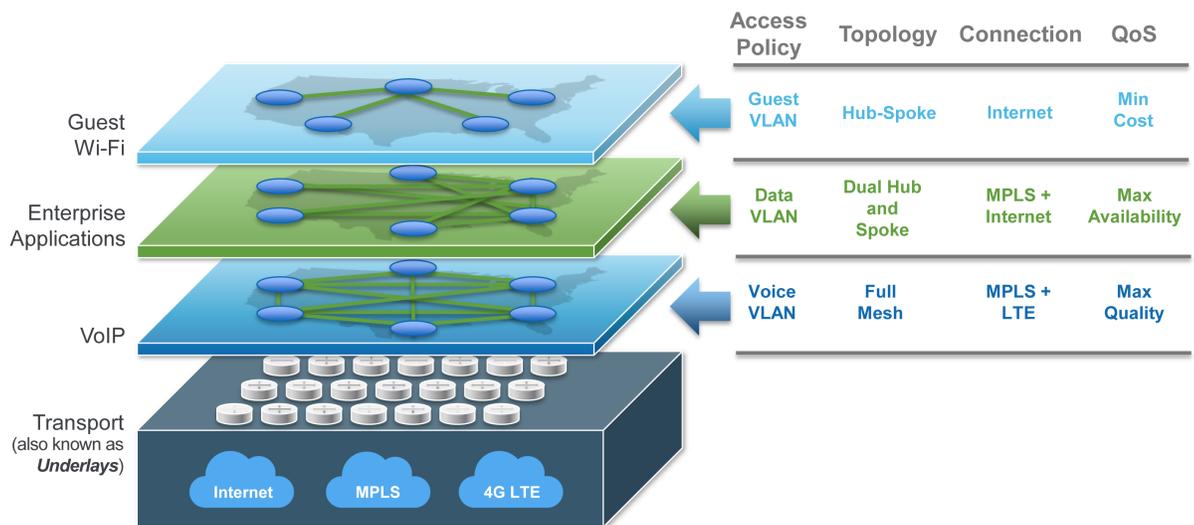
Orchestrator enables you to globally monitor performance and manage Silver Peak appliances, whether you're configuring a **WAN Optimization network** (NX, VX, or VRX appliances) or an **SD-WAN network** (EC or EC-V appliances).

# Overview of SD-WAN Prerequisites

With 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.

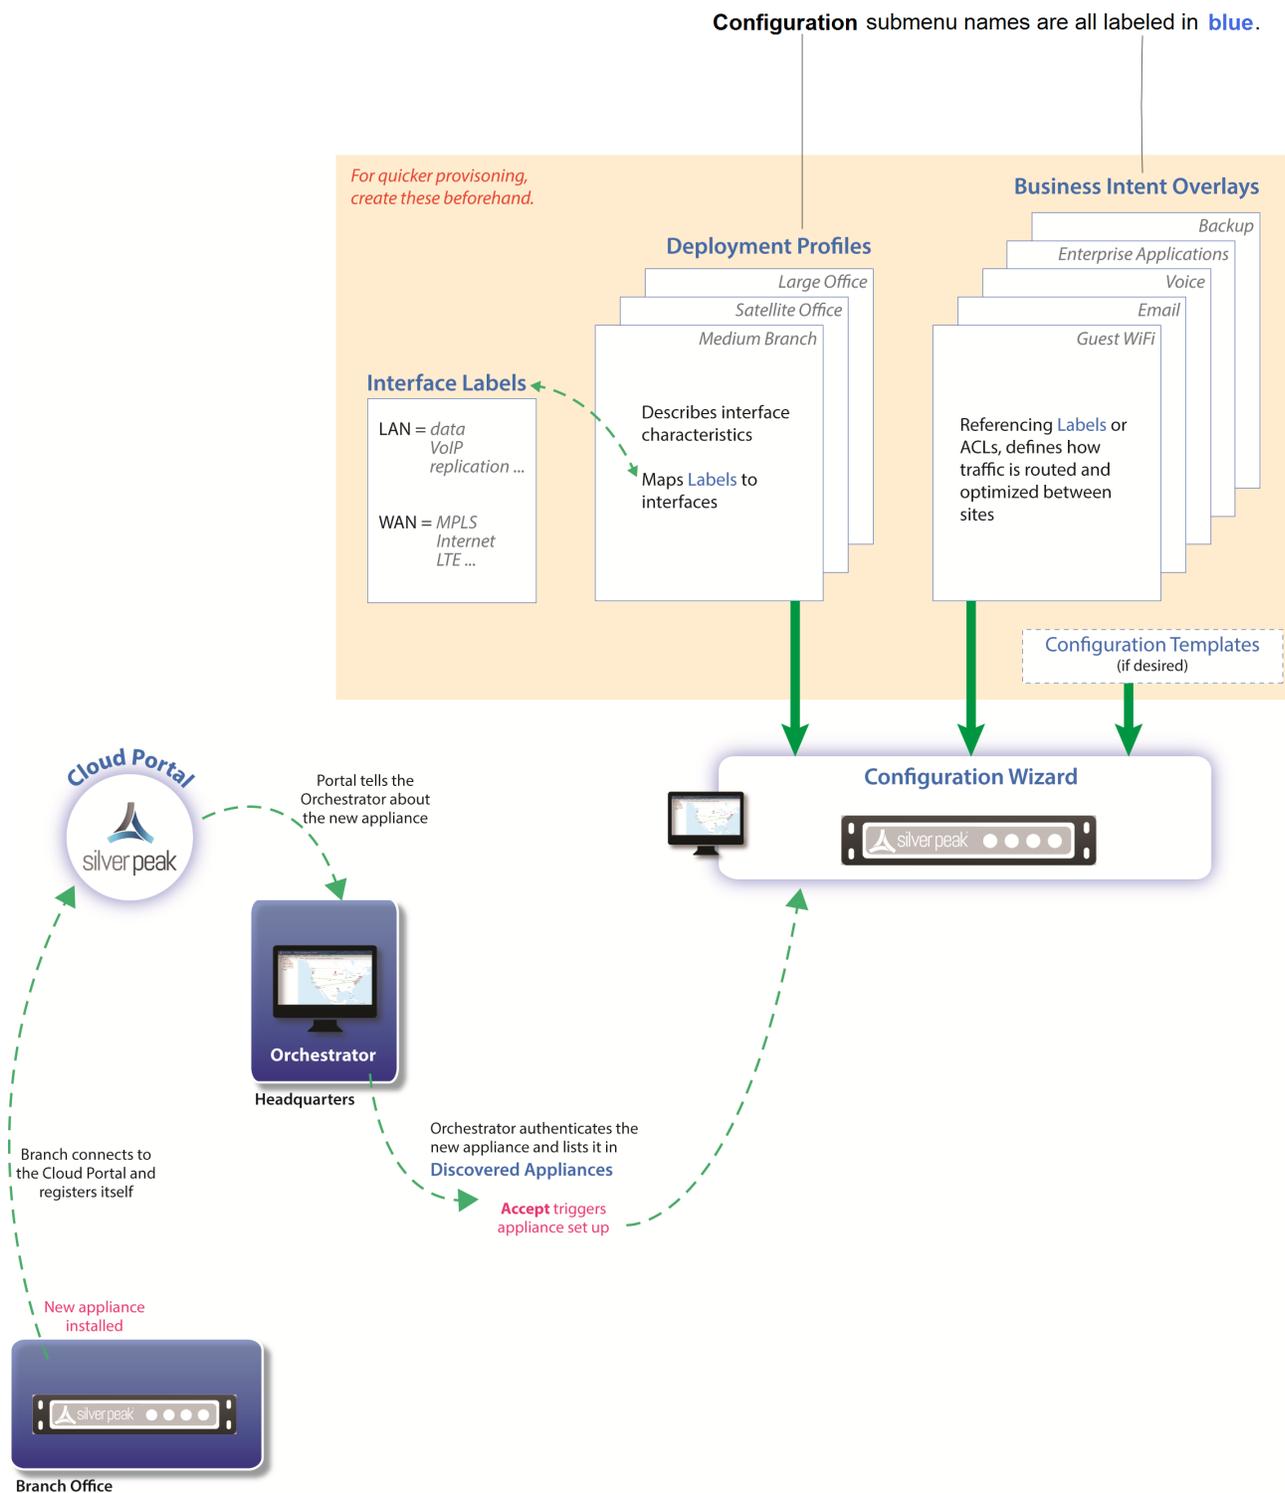
This diagram shows the basic architecture and capabilities of **Overlays**.



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

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

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



## Unity Overlays

These topics describe the pages related to deploying a WAN optimization network or a software-defined Wide Area Network (SD-WAN).

From a configuration standpoint, an SD-WAN uses Business Intent Overlays (BIOs), whereas a WANop network does not.

## Business Intent Overlays

Use the **Business Intent Overlays (BIOs)** tab to create separate, logical networks that are individually customized to your applications and requirements within your network. By default, there are several predefined overlays matching a range of traffic within your network.

The overlay summary table is used for easy comparison of values between your various configured overlays. You can select any link in the table and the **Overlay Configuration** dialog launches. You can also temporarily save your changes before officially applying those changes to your overlay. The pending configuration updates are indicated by an orange box around the edited item. Select **Save and Apply Changes to Overlays** when you are ready to apply the changes and select **Cancel** if you want to delete the changes.

### Overview

Orchestrator matches traffic to an ACL, progressing down the ordered, priority list of overlays until it identifies the first one that matches. The matched traffic is then analyzed against the overlay's Internet Traffic configuration, and forwarded within the fabric, or broken out to the internet based on the preferred policy order. If the software determines that the traffic is not destined for the internet, it refers to the **WAN Links & Bonding Policy** configuration and forwards traffic accordingly within the overlay.

### SD-WAN traffic to internal subnets

#### Overlay Configuration

You can begin to configure or modify a default overlay in the **Overlay** column. You can also select any icon on the **Business Intent Overlay** page and the selected editor or dialog opens.

Complete the following steps to configure your overlay.

1. Select the name of the overlay. The Overlay Configuration window opens. If you want to edit the default overlay or create a new overlay, enter the new name of the overlay in the **Name** field.
2. Select the **Match** field and choose the match criteria from the menu.
3. Select the **Edit** icon next to the ACL field. To apply default ACL's or create your own, select **Add Rule** in the **Associate ACL** window.
4. Select **Save**.

#### Region

To view your associated region within your overlay, select the **Regions** icon in the **Region** column in the overlay summary table. You can modify, remove, or edit overlay settings for a selected region by expanding the list at the right-top of the **Overlay Configuration** window. For more information regarding [Regions](#), refer to the help in the tab.

## Topology

Select the type of topology you want to apply to your overlay and network. You can choose between the following types of topology:

- **Mesh:** Choose **Mesh** if you want to make a local network.
- **Hub & Spoke:** Hubs are used to build tunnels in Hub & Spoke networks, and to route traffic between regions. If you choose **Hub & Spoke**, any appliance set as a hub will serve as a hub in any overlay applied to it. Hubs in different regions mesh with each other to support regional routing. To configure hubs, select the **Hubs** link at the top of the page.
- **Regional Mesh and Regional Hub & Spoke:** To streamline the number of tunnels created between groups of appliances that are geographically dispersed, you can assign appliances to **Regions** and select **Regional Mesh** or **Regional Hub & Spoke**.

1. At the top of the page, select **Regions**.
2. You can add and remove a region or view the status of each overlay within a selected region.

## Building SD-WAN using these interfaces

You can select which WAN interfaces you want to use for each device to connect to the SD-WAN. First, you assign for your traffic to go to the **Primary** interfaces. If the primary interface is unavailable or not meeting the desired Service Level Objectives configured, the **Backup** interfaces are used. Move the desired interfaces between **Primary** and **Backup**. The interfaces are grayed out until moved into the **Primary** or **Backup** boxes.

- **Cross Connect** allows you to define tunnels built between each interface label. Each appliance has a maximum number of tunnels that it can support, and using **Cross Connect** increases the number of tunnels created.
- **Add Backup if Primary Are:** Specifies when the system should use the Backup interfaces.

## Service Level objective

Traffic is routed through the primary interfaces exclusively, unless the service level thresholds for **Loss**, **Latency**, or **Jitter** have been exceeded. If this occurs, backup interfaces are added so that the service level objective can be met.

**NOTE** Primary interfaces may still be used to support the overall Service Level Objective.

## Link Bonding Policy

You can select the following Link Bonding Policies when you need to specify the criteria for selecting the best route possible when data is sent between multiple tunnels and appliances.

Field	Definition
High Availability	For critical 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. This option sends load balance information on multiple links, with no FEC or overhead.

## QoS, Security, & Optimization

To further customize your overlay configuration, enter the appropriate information for the following fields.

Field	Definition
FW Zone	Select the firewall zone you want to restrict traffic to from an overlay.
Boost	Select <b>True</b> or <b>False</b> if you want to apply any purchased Boost to your overlay.
Peer Unavailable Option	Select the following options you want your traffic to go if a peer is unavailable: <b>Use MPLS</b> , <b>Use Internet</b> , <b>Use LTE</b> , <b>Use Best Route</b> , <b>Drop</b> .
Traffic Class	Channels traffic to the desired queue based on the applied service. Select <b>Best Route</b> or Drop.
LAN DSCP	Select the DSCP you want to apply as a filter to the LAN interface.
WAN DSCP	Select the DSCP you want to apply as a filter to the WAN interface.

## Breakout Traffic to Internet & Cloud Services

You can use the **Breakout Traffic to Internet & Cloud Services** to monitor and manage traffic coming to or from the internet.

### Hub versus branch breakout settings

You can create different breakout policies for hubs. Any hub you select in the **Topology** section also displays at the top of the **Internet Traffic to Web, Cloud Services** tab. When you select an individual

hub, the **Use Branch Settings** displays, selected, to the right of the screen. Complete the following steps to create a custom breakout policy for that hub:

1. Clear the check box for **Use Branch Settings**.
2. Configure the now accessible parameters.
3. Select **OK**.

### Preferred Policy Order and Available Policies

- You can move policies back and forth between the **Preferred Policy Order** and the **Available Policies** columns. You can also change their order within a column. The defaults provided are **Backhaul via Overlay**, **Break Out Locally**, and **Drop**.
- When you choose **Break Out Locally**, confirm that any selected interface that is directly connected to the Internet has **Stateful Firewall** specified in the deployment profile.
- You can add services (such as Zscaler, Fortigate, or Palo Alto). The service requires a corresponding Internet-breakout (Passthrough) tunnel for each appliance traffic to that service. To add a service, select the **Edit** icon next to **Available Policies**.
- The **Default** policy you configure for internet breakout is pushed to all appliances that use the selected Overlay. However, you might want to push different breakout rules to your hubs.

## Deployment Profiles

*Configuration > 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 a 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**.

**NOTE** You cannot edit **IP/Mask** fields because they are appliance-specific.

### 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.
- The firewall zones you have already configured will be in the list under **FW Zone**. Select the FW zone you want to apply to the LAN you are deploying.

## WAN-side Configuration

**Firewall Zone:** Zone-based firewalls are created on the Orchestrator. A zone is applied to an **Interface**. By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. Users can create exception rules (Security Policies) to allow traffic between interfaces with different zones. The firewall zones you have already configured will be in the list under **FW Zone**. Select the FW zone you want to apply to the WAN you are deploying.

**Firewall Mode:** Four options are available at each WAN interface:

- **Allow All** permits unrestricted communication.
- **Stateful *only*** allows communication from the LAN-side to the WAN-side.  
Use this if the interface is behind the WAN edge router.
- **Stateful with SNAT** applies Source NAT to outgoing traffic.  
Use this if the interface is directly connected to the Internet.
- **Harden**
  - For traffic inbound from the WAN, the appliance accepts ***only*** IPSec tunnel packets that terminate on a Silver Peak appliance.
  - For traffic outbound to the WAN, the appliance ***only*** allows IPSec tunnel packets and management traffic that terminate on a Silver Peak appliance.



**WARNING** Activating fail-to-wire will DISABLE ALL firewall rules.

---

**NAT Settings:** 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. Select one of the following options:

- If the appliance is behind a NAT-ed interface, select **NAT**.
- If the appliance is not behind a NAT-ed interface, select **Not behind NAT**.
- **Enter an IP address** to assign a destination IP for tunnels being built from the network to this WAN interface.

**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

- For additional 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.
- Select the appropriate licensing you have applied to your EC appliance from the menu. The licenses will only display depending on the licenses you have for that particular account. You can select the following licensing options:
  - Mini
  - Base
  - Base + Plus
  - 50 Mbps
  - 200 Mbps
  - 500 Mbps
  - 1 Gbps
  - 2 Gbps
  - Unlimited

**NOTE** You must have the correct hardware to support the license selected.

#### 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. Refer to the Silver Peak user documentation.

## Definitions

### *DHCP Server*

Field Name	Description
Default gateway	When selected, indicates the default gateway is being used.
Default lease, Maximum lease	Specify, in hours, how long an interface can keep a DHCP-assigned IP address.
DNS server(s)	Specifies the associated Domain Name System server(s).
Exclude first N addresses	Specifies how many IP addresses are not available at the beginning of the subnet's range.
Exclude last N addresses	Specifies how many IP addresses are not available at the end of the subnet's range.
NetBIOS name server(s)	Used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
NetBIOS node type	The <b>NetBIOS node type</b> of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types: <ul style="list-style-type: none"> <li>■ <b>B-node</b> = 0x01 Broadcast</li> <li>■ <b>P-node</b> = 0x02 Peer (WINS only)</li> <li>■ <b>M-node</b> = 0x04 Mixed (broadcast, then WINS)</li> <li>■ <b>H-node</b> = 0x08 Hybrid (WINS, then broadcast)</li> </ul>
NTP server(s)	Specifies the associated Network Time Protocol server(s).
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.
Subnet Mask	A mask that specifies the default number of IP addresses reserved for any subnet. For example, entering <b>24</b> reserves 256 IP addresses.

*DHCP/BOOTP Relay*

Field Name	Description
Destination DHCP/BOOTP Server	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 <b>append</b> , <b>replace</b> , <b>forward</b> , or <b>discard</b> .

## 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 website for various deployment guides.

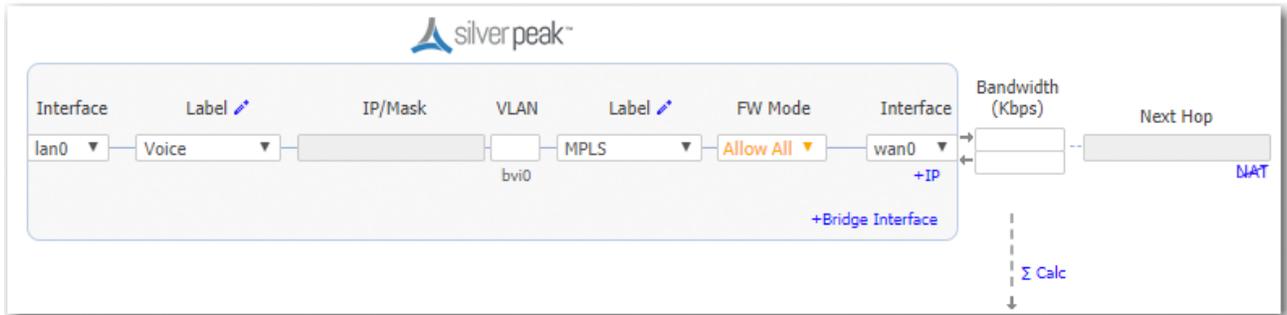
In Bridge Mode and in Router Mode, you can provide security on any WAN-side interface by **hardening the interface**. This means:

- 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.

## Bridge Mode

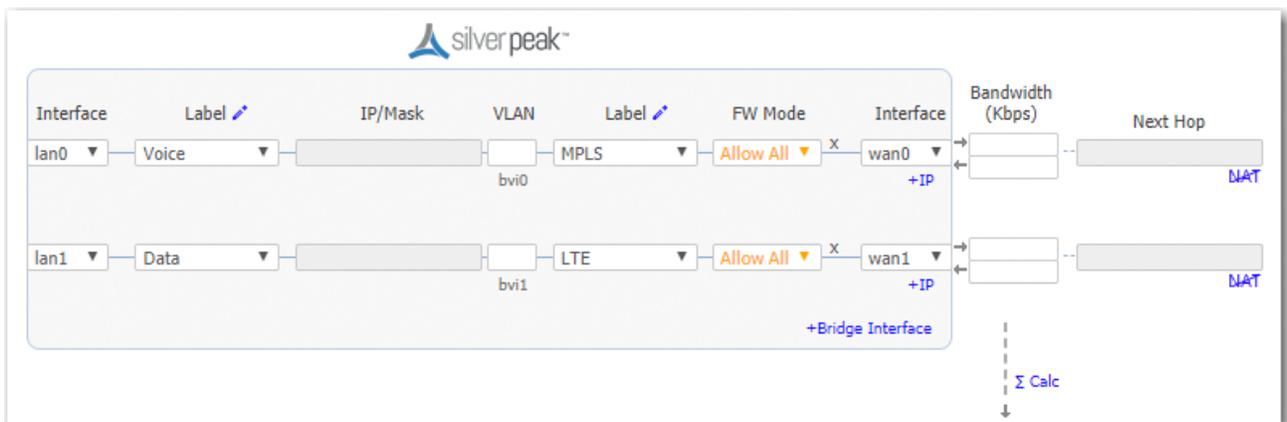
### Single WAN-side Router

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.

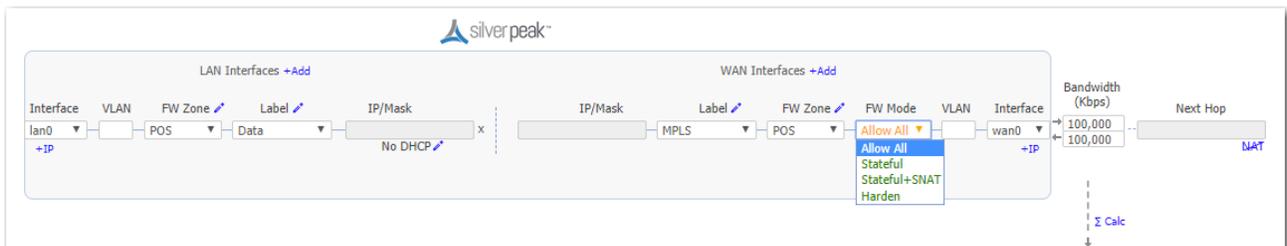
## Router Mode

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

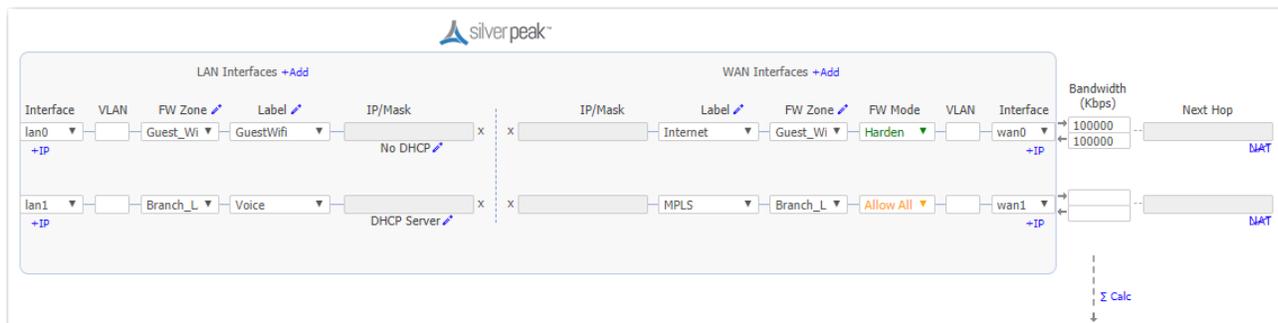


For this deployment, you have two options:

1. You can put Silver Peak *in-path*. In this case, if there is a failure, you need other redundant paths for high availability.
2. 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.

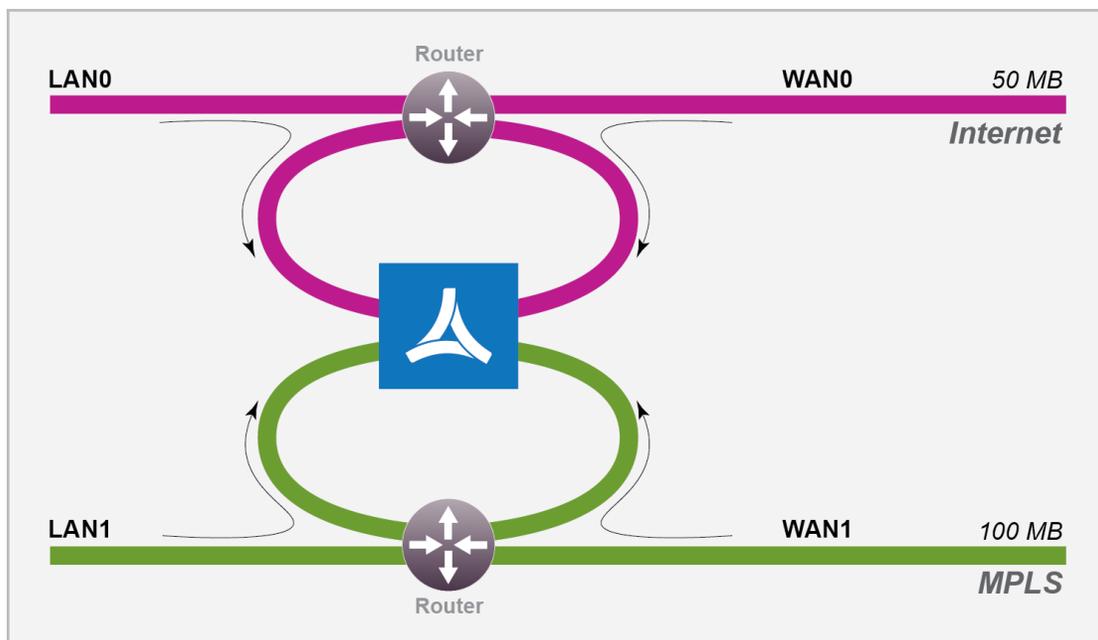
### #2 - Dual LAN Interfaces & Dual WAN Interfaces



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.)

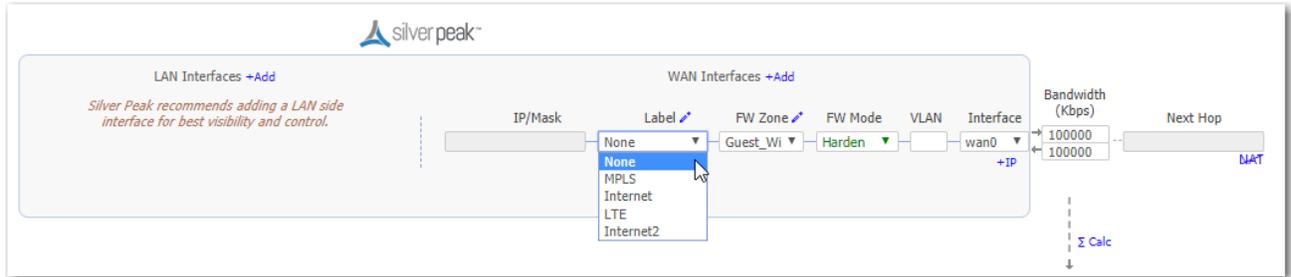
#### Out-of-path dual LAN and dual WAN interfaces



For this deployment, you have two options:

1. You can put Silverpeak *in-path*. In this case, if there is a failure, you need other redundant paths for high availability.
2. You can put Silverpeak *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).

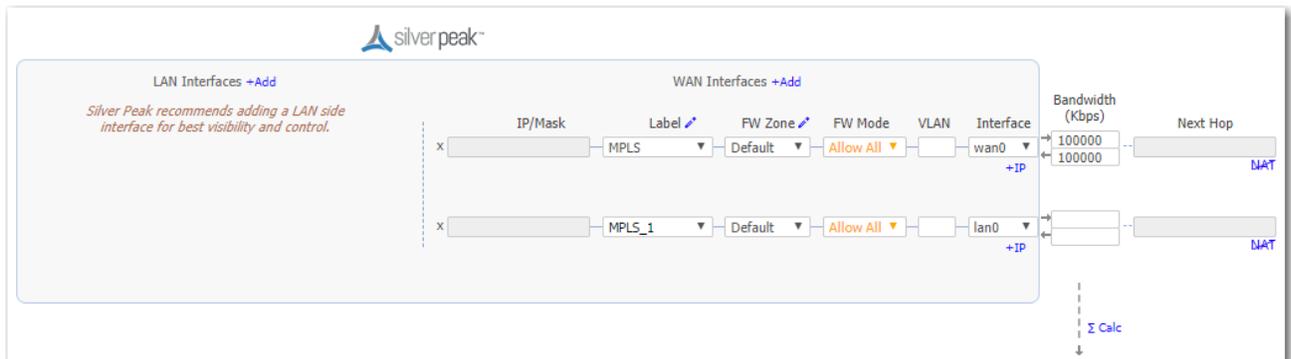
### #3 - Single WAN Interface Sharing LAN and WAN traffic



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



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**):

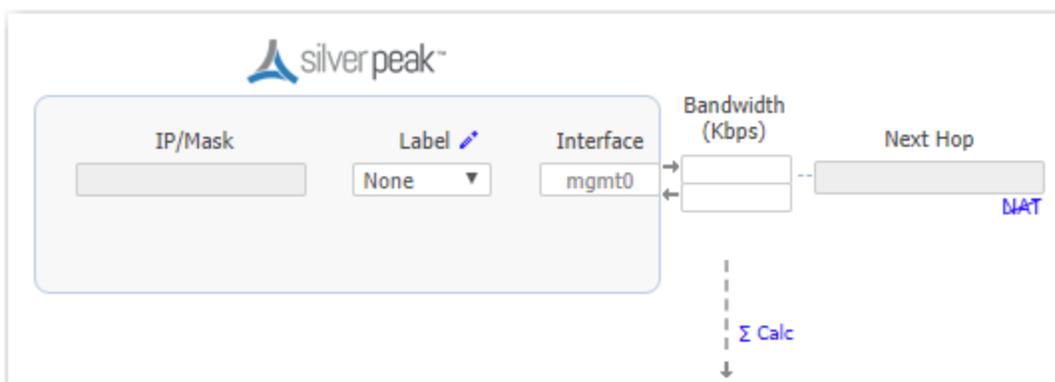
- If you use **subnet sharing** (which relies on advertising local subnets between Silver Peak appliances) or **route policies** (which specify destination IP addresses), then you only need LAN-side redirection.
- If, instead, you rely on **TCP-based** or **IP-based** auto-optimization (which relies 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.



## 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**.

## Deployment – EdgeConnect HA

The EdgeConnect HA (High Availability) mode is a high availability cluster configuration that provides appliance redundancy by pairing two EdgeConnect devices together.

When a deployment profile configures two EdgeConnect appliances in EdgeConnect HA mode the resilient cluster acts as a single logical system. It extends the robust SD-WAN multipathing capabilities such as Business Intent Overlays seamlessly across the two devices as if they were one entity.

With EdgeConnect HA mode a WAN uplink is physically plugged into a single one of the EdgeConnect appliances but is available to both in the cluster. For WAN connections that perform NAT (for example, a consumer-grade Broadband Internet connection), it means that only a single Public IP needs to be provisioned in order for both the EdgeConnect devices in the EdgeConnect HA cluster to be able to build Business Intent Overlays using that transport resource.

### Enabling EdgeConnect HA Mode

1. In the navigation pane, select the appliance and then right-click to select **Deployment** from the contextual menu. The appliance's Deployment page appears.
2. Select the **EdgeConnect HA** checkbox.
3. Configure the interfaces (LAN and WAN-side) on both EdgeConnect devices to reflect the WAN connections that are plugged into each one of the respective appliances.

**NOTE** Both EdgeConnect devices will be able to leverage all WAN connections regardless of which chassis they are physically plugged into. It is however important to match the deployment profile interface configuration to the actual chassis the WAN connection is physically, directly connected to.

4. Select the physical ports on the respective EdgeConnect appliances that you will connect to each other using an Ethernet cable (RJ-45 twisted pair or SR optical fiber)

**NOTE** You can choose any LAN or WAN port combination for this HA Link that is available on the respective EdgeConnect chassis. You must match the media type and speed for both ends of the HA link (for example, 1 Gigabit-Ethernet RJ-45 to RJ-45 or 10 Gigabit-Ethernet multimode fiber LC-connector-to-LC-connector). Also please note that you cannot use MGMT ports for the HA Link; only LAN or WAN ports.

### IPSec over UDP Tunnel Configuration

For both EdgeConnect appliances in a high availability cluster to be able to share a common transport connection, you must set the tunnel type to IPSec over UDP mode.

Please see the Overlay Tunnel Settings in the Orchestrator, under the **Orchestrator Overlay Manager and Tunnel Settings** menu.

**NOTE** *If you are deploying a network with EdgeConnect appliances running VXOA 8.1.6 or higher and Orchestrator 8.2 or higher, the tunnel type is already set to IPSec over UDP mode by default.*

## VRRP Configuration

Typically, in a branch site deployment, you will choose to configure the cluster with a VRRP protocol and assign a VIP (virtual IP) address to the cluster.

- Set the VRRP priority of the preferred LAN-side Primary EdgeConnect to **128**.
- Set the other, Secondary appliance's VRRP priority to **127**.

## LAN-Side Monitoring

The IP SLA feature should be configured to monitor the LAN-side VRRP state in order to automatically disable subnet sharing from that appliance in the case of a LAN link failure.

Please refer to the IP SLA configuration guide for more information.

## Interface Labels

*Configuration > Overlays > Interface Labels*

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

Interface Labels
×

Edit	Type	Label	
	wan	test (Hub & Spoke)	×
	wan	LTE	
	wan	Internet	
	wan	MPLS	
	lan	GuestWifi	
	lan	POS	×
	lan	Data	×
	lan	Voice	×

*Use labels to match and route traffic into overlays. Type specifies "which side" of the network the interface is on. LAN labels identify LAN-side data (subnets), and WAN labels identify the WAN service. If you edit a label, tunnels that reference that labeled interface are renamed accordingly.*  
*Labels used in overlays or tunnel groups cannot be deleted.*

Select **New Label** to add a new interface label to an appliance.

You can also use interface labels to end a connection between tunnels.

Select the **Edit** icon and the Interface Label Configuration window pops up.

Interface Label Configuration
×

wan
lan

Label Name

Topology

Enter the name of the interface label name of the tunnels to be avoided in a selected topology.

## Firewall Zones

*Configuration > Overlays > Firewall Zones*

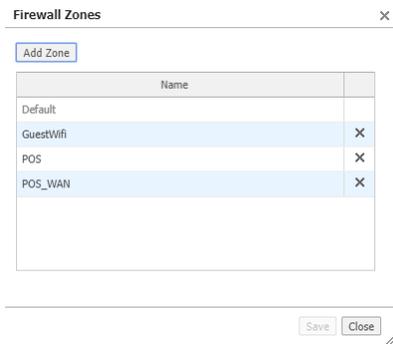
Zone-based firewalls are created on the Orchestrator.

A zone is applied to an **Interface**.

By default, traffic is allowed between interfaces labeled with the same zone.

Any traffic between interfaces with different zones is dropped.

Users can create exception rules (Security Policies) to allow or deny traffic between interfaces within the same or different zones.



Note: "Default" will always be the initial default zone. You cannot have another zone named "Default".

Note: The name of your firewall cannot exceed 16 characters and cannot contain any special characters. It can only contain alphanumeric characters and underscores only.

## Apply Overlays

*Configuration > Overlays > Overlay Prerequisites > Apply Overlays*

Use this page to **add or remove overlays** from appliances. If you select **Edit Overlays**, you will be redirected to the **Business Intent Overlay** tab for further customization. You can also view the status of the overlays if you select **View Status**.

## Internet Traffic

Configuration > [Overlays > Overlay Prerequisites] Internet Traffic

Internet Traffic is defined as any traffic the **does NOT match** the internal subnets listed on this page.

**Internet Traffic** [X]

*Internet traffic is defined as any traffic that **DOES NOT MATCH** the internal subnets below.*

**Internal Subnets**

Subnet/Mask	
192.168.0.0/16	X
172.16.0.0/12	X
10.0.0.0/8	X

## IPSec Pre-shared Key Rotation

*Configuration > [Overlays] IPSec Key Rotation*

Use this page to schedule the rotation of auto-generated IPSec pre-shared keys.

**Schedule IPSec Pre-shared Key Rotation** [X]

**Ipssec Key Rotation**

Enable

Period Every Sunday at 1:00 starting 30-Jul-18 11:09 PDT

**UDP IPSec Key Rotation**

Enable

Period Every day at 10:20 starting 18-Oct-17 10:20 PDT

Save Cancel

## Failure Handling and Orchestrator Reachability

Orchestrator distributes key material to all EdgeConnect appliances in the network. Immediately before the end of a key rotation interval, Orchestrator activates new ephemeral key material for all of the EdgeConnect appliances in the SD-WAN network. For key activation, all the appliances should be reachable to Orchestrator. However, there are two cases of unreachability:

1. **Inactive appliances:** When appliances are inactive, they exist in the Orchestrator, but don't have tunnels configured to any active appliances.
2. **Temporary unreachability:** Temporary unreachability issues occur in cases where an EdgeConnect appliance reboots or if there is a link or communication failure. In this case, Orchestrator won't activate the new key material until all active appliances are reachable and have received the new key material. If the appliance is unreachable for a period longer than the key rotation interval, it will be treated as an inactive appliance.

**Re-authorization:** Inactive appliances that become active at a later point in time will be authorized to receive the current key material. Only then will they be able to download configurations and build tunnels.

## Hubs

*Configuration > Overlays > Hubs*

In this tab, you can add, remove, and associate hubs to a specified region within the Regional Mesh or Regional Hub-and-Spoke topologies configured in the **Business Intent Overlay** tab.

Complete the following steps to add a hub:

1. Enter the name of the hub you want to add from the menu.
2. Select **Add Hub**.

To delete a hub, select the **X** icon next to the hub you want to delete.

## Discovered Appliances

Configuration > [Unity Overlays] Discovered Appliances

This page lists each appliance that Orchestrator discovers.

Discovered Appliances x											
Discovered Devices <span>?</span> <input type="button" value="Save"/>											
Discovery Email Recipients <input type="text" value="rsinha@silver-peak.com"/> <input type="button" value="Save"/>											
<a href="#">Show Denied Devices</a> <a href="#">Show Approved Orchestrators</a>											
15 Rows <span style="float: right;">Search <input type="text"/></span>											
Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Reachability	Approve	Deny	Software Version	Model
0100AAB1016	GM5-7.3		128.242.109.226	Milpitas, California, US	Unassigned	10-Feb-17 11:07		<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.0.10.31334	GX-V
001BBC090B05	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	11-Jan-17 11:50	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.4.0_62991	EC-V
001BBC090B05	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	11-Jan-17 10:50	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.4.0_62991	EC-V
001BBC090B05	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	10-Jan-17 16:58	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.4.0_62991	EC-V
0100AAB1016	GM5-7.3		128.242.109.226	Milpitas, California, US	Unassigned	06-Jan-17 16:51		<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.0.10.31334	GX-V
001BBC091F3B	CPX-1	10.0.233.189	128.242.109.226	Milpitas, California, US	Unassigned	06-Jan-17 10:31	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.4.0_62991	CPX
001BBC090B05	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	06-Jan-17 10:29	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	7.3.9.0_62228	EC-V
001BBC090B05	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	05-Jan-17 12:55	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	7.3.9.0_62228	EC-V
001BBC091AD5	ECV-A	10.8.43.10	128.242.109.226	Milpitas, California, US	Unassigned	29-Nov-16 15:14	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.4.0_62697	EC-V
000C2913A3BE	SqaSaas-VWA	172.25.40.215	128.242.109.226	Milpitas, California, US	Unassigned	21-Oct-16 22:15	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.4.0_62671	VX-1000
000C29490B05	traceroute-vwb	10.0.248.33	128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Reachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.1.0_60893	VX-1000
000C29FC7634	traceroute-vva	10.0.248.28	128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Reachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.1.0_60893	VX-1000
001BBC091E21	rsinha-ecv		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>		EC-V
000C29790872	rsinha-vxd	10.0.233.134	10.0.233.134	-,-,-	Unassigned	20-Oct-16 11:44	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.0.2.0_59017	VX-2000
000C29B9008F	rsinha-vxc	10.0.233.132			Unassigned	19-Oct-16 11:41	Unreachable	<input type="button" value="Approve"/>	<input type="button" value="Deny"/>	8.1.1.0_60810	VX-2000

- To enable 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 Devices**, locate it in the table, and click **Approve**.

Discovered Appliances x											
Denied Devices <span>?</span> <input type="button" value="Save"/>											
Discovery Email Recipients <input type="text" value="rsinha@silver-peak.com"/> <input type="button" value="Save"/>											
<a href="#">Show Discovered Devices</a>											
7 Rows <span style="float: right;">Search <input type="text"/></span>											
Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Reachability	Approve	Software Version	Model	Comment
000C290AF2A1	rsinha-vva	10.0.233.164			Unassigned	19-Oct-16 11:42	Reachable	<input type="button" value="Approve"/>	8.0.8.0_63501	VX-3000	Appliance was delet...
001BBC011E36	ECXS2	10.0.250.53	128.242.109.226	Milpitas, California, US	Unassigned	30-Jan-17 16:26	Reachable	<input type="button" value="Approve"/>	8.0.8.0_63501	EC-XS	Appliance was delet...
001BBC121FAA	ECXS1	10.0.250.52	128.242.109.226	Milpitas, California, US	EC-XS-San Jose	07-Feb-17 14:15	Reachable	<input type="button" value="Approve"/>	8.0.8.0_63501	EC-XS	Appliance was delet...
001BBC121FAA	ECXS1	10.0.250.52	128.242.109.226	Milpitas, California, US	Unassigned	21-Oct-16 11:49	Unreachable	<input type="button" value="Approve"/>	8.0.6.0_61853	EC-XS	Appliance was delet...
001BBC090B05	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	11-Jan-17 10:46	Unreachable	<input type="button" value="Approve"/>	8.1.4.0_62991	EC-V	Appliance was delet...
001BBC091F3B	silverpeak-2494a5	10.0.233.189	128.242.109.226	Milpitas, California, US	Unassigned	05-Jan-17 14:05	Unreachable	<input type="button" value="Approve"/>	8.1.4.0_62991	CPX	Appliance was delet...
000C290D1901	rsinha-vvb	10.0.233.169			Unassigned	19-Oct-16 11:46	Reachable	<input type="button" value="Approve"/>	8.1.4.0_62779	VX-2000	Appliance was delet...

- As a security measure to prevent unauthorized management of your network, any Orchestrator with your Account Name and Account Key must be approved by the originally deployed Orchestrator.

To view the approved Orchestrators, click **Show Approved Orchestrators**.

Discovered Appliances x

Approved Orchestrators 🔍 📄 Discovery Email Recipients  Save Show Discovered Devices

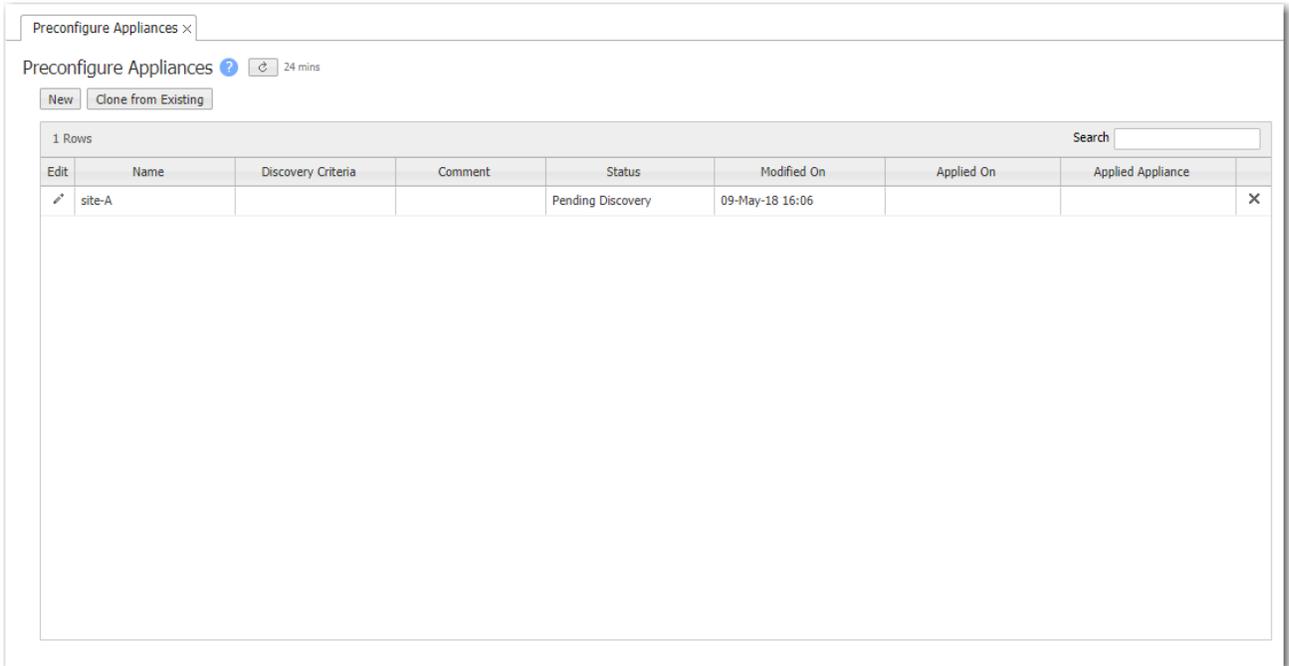
3 Rows Search

Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Deny	Software Version	Model
0030AAB1016	GMS		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Deny	8.1.3.30972	GX-V
0030AAB1016	GMS		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Deny	99.99.99.30654	GX-V
0100AAB1016	GMS-7.3		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Deny	7.3.10.30880	GX-V

## Preconfigure Appliances

*Configuration > [Overlays > Discovery] Preconfiguration*

You can use this page to prepopulate flat data files that are matched with appliances as you add them to your network.



The screenshot shows the 'Preconfigure Appliances' interface. At the top, there is a breadcrumb trail: 'Preconfigure Appliances > [Overlays > Discovery] Preconfiguration'. Below this, the page title 'Preconfigure Appliances' is displayed with a help icon and a refresh icon. There are two buttons: 'New' and 'Clone from Existing'. A search bar is located on the right side of the table header. The table has the following columns: Edit, Name, Discovery Criteria, Comment, Status, Modified On, Applied On, and Applied Appliance. The first row contains the following data: Edit (pencil icon), Name (site-A), Discovery Criteria (empty), Comment (empty), Status (Pending Discovery), Modified On (09-May-18 16:06), Applied On (empty), and Applied Appliance (empty). There is a close icon (X) in the bottom right corner of the table.

1 Rows								Search
Edit	Name	Discovery Criteria	Comment	Status	Modified On	Applied On	Applied Appliance	
	site-A			Pending Discovery	09-May-18 16:06			

The information in the files is a combination of items found in the Appliance Configuration Wizard, along with site-specific information (such as BGP, OSPF, IP SLA rules, VRRP, interfaces, and addressing).

You can create a new file or clone (and rename) an existing one. Make any changes with the built-in editor.

After the appliance is discovered and approved, software upgrade and configuration push are done automatically.

## New or Clone

**Appliance Preconfiguration**
✕

Name\*

Comment

Auto Approve when Discovered

**Discovery Criteria**

Serial

Appliance Tag

```

1 # Preconfiguration Template for Appliance Setup
2 # If any fields aren't needed then remove them.
3 # applianceInfo:
4 # softwareVersion - Version of the software to upgrade the appliance to,
5 #   this will happen before the rest of the preconfiguration
6 #   is applied. The valid versions to use can be seen in the
7 #   "Appliance Upgrade" dialog in Orchestrator
8 #   Values: string software version. ex: "8.4.0.0_12345"
9 # hostname - Hostname to set for this appliance
10 #   Values: Maximum 60 characters
11 # group - The group in the Orchestrator tree where this appliance should be
12 #   added
13 #   Values: Orchestrator tree group name
14 # site - Name for the site where this appliance resides.
15 #   Site is mainly used to prevent building tunnels
16 #   between appliances with the same site name
17 #   Values: any string
18 # networkRole - Role that appliance plays in the network, specifically
19 #   tunnel topology.
20 #   Values: "spoke", "hub", "mesh"
21 # region - Region name used to connect hubs between regions.
22 #   Values: any string
23 # location - These are put into the appliance SNMP configuration.
24 #   This section is required.
25 #   Location comprises the following key/value pairs:
26 #   address - Address 1
27 #     Values: any string
28 #   address2 - Address 2
29 #     Values: any string
30 #   city - City
31 #     Values: any string
32 #   state - State
33 #     Values: any string
34 #   zipCode - Zip Code
35 #     Values: any string

```

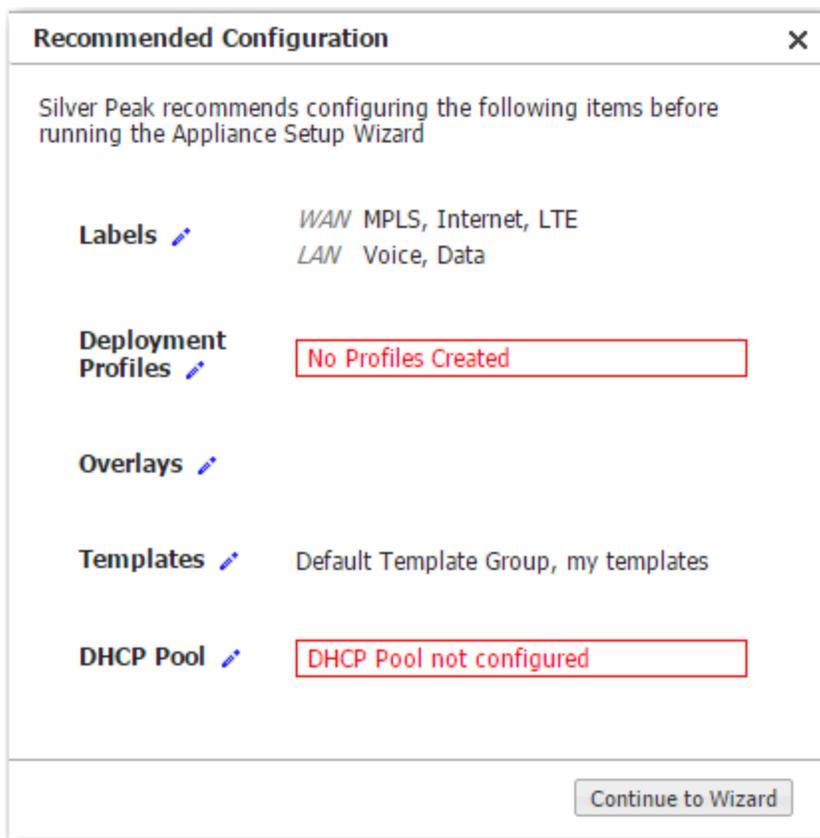
Field	Definition
<b>Name</b>	Assigns a name to the preconfiguration file.
<b>Comment</b>	An optional descriptive field
<b>Auto Approve when Discovered</b>	When <b>selected</b> , the Orchestrator finds the appliance that matches the Discovery Criteria and automatically loads it without the need for user intervention. When <b>deselected</b> , the user will be prompted to manually approve the association of preconfiguration file to the appliance.
<b>Serial</b>	The serial number associated with the appliance that is to receive this configuration.
<b>Appliance Tag</b>	A free-form text or unique identifier that an administrator can associate with the appliance. Available as a discovery criteria for EC-Vs.

# Appliance Configuration Wizard

*Configuration > [Overlays > Discovery] Configuration Wizard*

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

**NOTE** Orchestrator assumes that you'll be pushing many of the same configuration items to each appliance. To that end, it surveys the templates and Overlay prerequisite items and displays the **Recommended Configuration** list, showing what comprehensive items you have and have not yet configured.

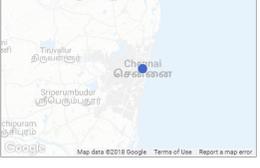


**Appliance Wizard**

Appliance Setup 1 2 3 4

Hostname\*  Site Name  Used to identify appliances at the same location. Tunnels will NOT be built between appliances with identical Site Names.  
 Group\*  Contact Name   
 Admin Password  Contact Email   
 Confirm Password  Serial Number\*   
 Software Version

Location  
 Address 1   
 Address 2   
 City  State   
 Zip Code  Country   
 Region



**Appliance Wizard**

Appliance Setup 1 2 3 4

Deployment Profile  ⓘ

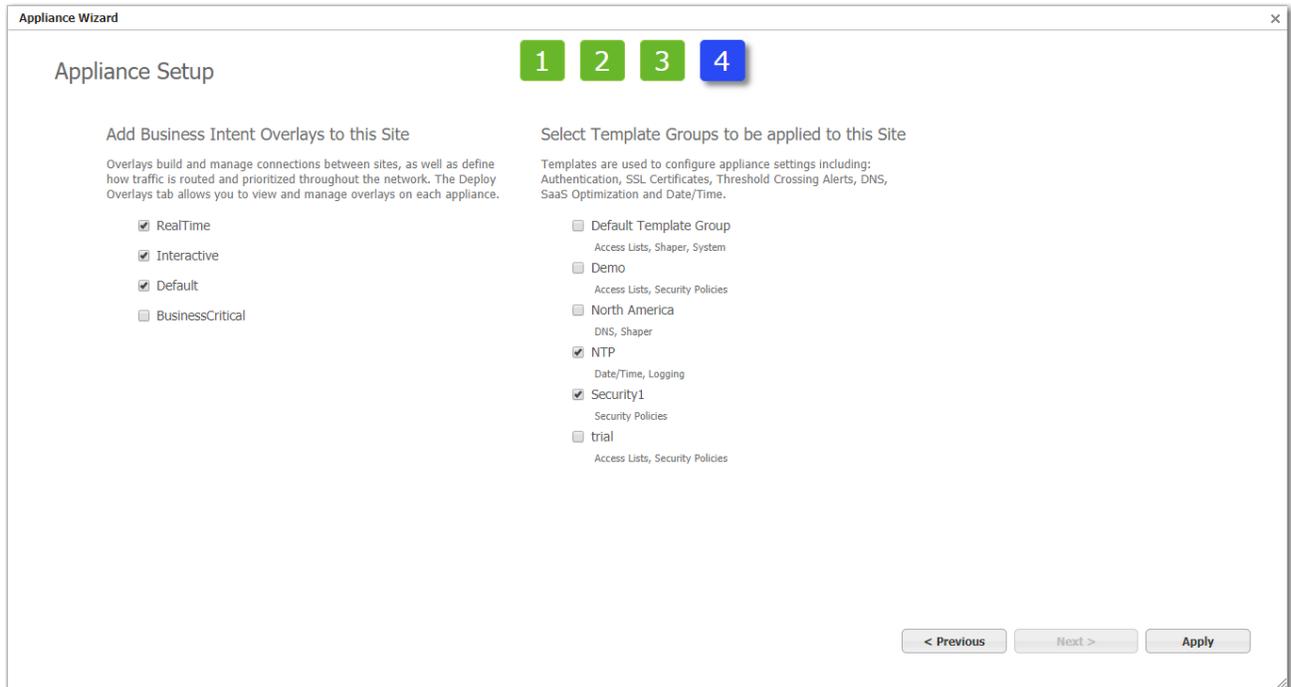
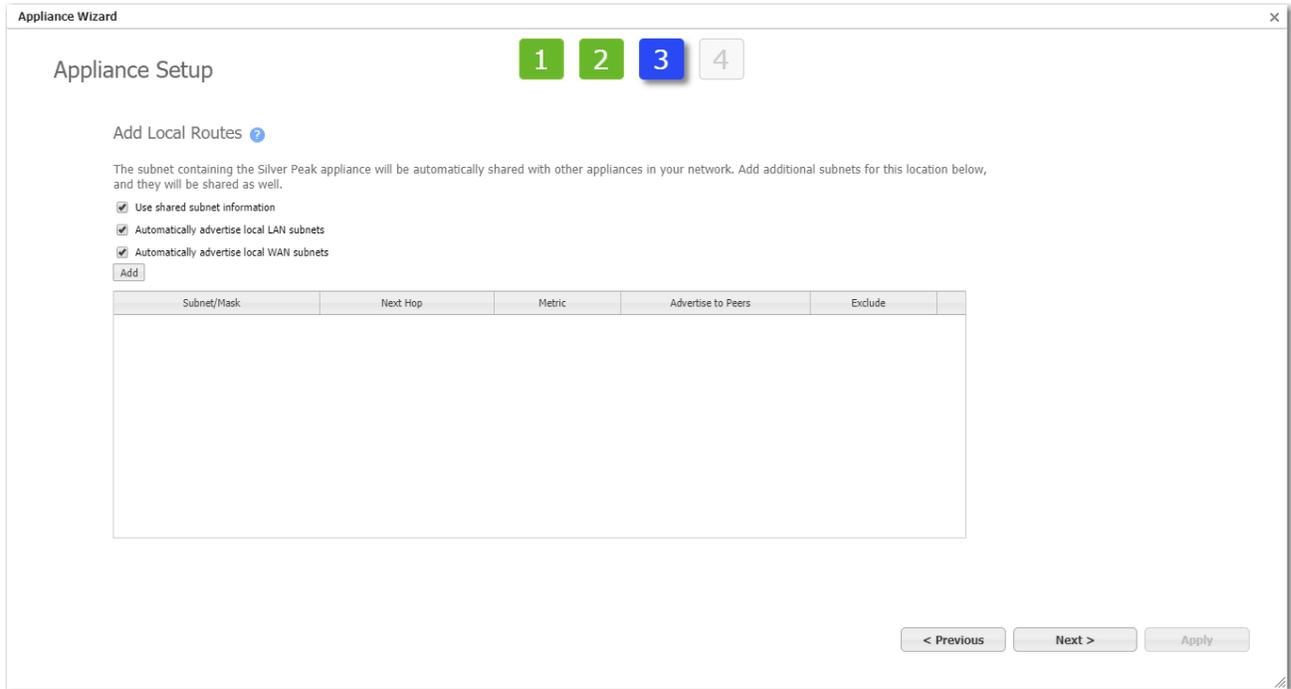


Next Hop	Interface	VLAN	FW Zone	Label	IP/Mask	IP/Mask	Label	FW Zone	FW Mode	VLAN	Interface	Bandwidth (Kbps)	Next Hop
<input type="text"/>	lan0 +IP		Default	None	10.17.17.20/24 No DHCP	10.17.18.20/24	MPLS	Default	Allow All		wan0 +IP	100,000	10.17.18.1 MAT
						DHCP10.0.184.154/24	Internet	Default	Stateful+		wan1 +IP	100,000	DHCP10.0.184.1 MAT

Σ Calc

Total Outbound →  Kbps ≤ 200,000 Kbps  
 Total Inbound ←  Kbps  Shape Inbound Traffic

EdgeConnect Licensing   Kbps



## Licenses

*Configuration > Overlays > Licensing > Licenses*

This page lists each appliance's make, model, license terms, and registered services. You can also revoke or regrant a metered license to/from an appliance or change the EdgeConnect (EC) license settings and RMA your device. Complete the following steps to configure or modify your EC license.

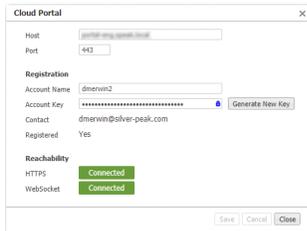
1. Select the **Edit** icon next to a selected appliance in the table. The **Configure EdgeConnect License** window opens.
2. Check **Grant**, **Revoke**, or **No Change**.
3. Select the following EC size options from the menu: **Mini**, **Base**, **Base + Plus**, **50 Mbps**, **200 Mbps**, **500 Mbps**, **1 Gbps**, **2 Gbps**, and **Unlimited**.
4. Check **Enable Boost** if you want to enable the boost you have purchased with your license.
5. Enter the amount of boost you have applied to your EC.
6. Select **Apply**.

**NOTE** EdgeConnect stops passing traffic if your license has expired.

## Cloud Portal

*Configuration > [Overlays > Licensing] Cloud Portal*  
*Orchestrator > [Orchestrator Server > Licensing] Cloud Portal*

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



- 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.

## SSL Certificates Tab

Configuration > [Overlays > SSL] SSL Certificates

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

The screenshot shows the 'SSL Certificates' tab in the Silver Peak Orchestrator interface. At the top, there are tabs for 'Topology' and 'SSL Certificates'. Below the tabs, there is a section titled 'Manage SSL Certificates with Templates' with an 'Export' button and a refresh icon. The main content area is titled 'SSL Certificates' and features a search bar and a 'Show 25' dropdown menu. A table displays the following data:

Edit	Appliance Name ▲	Issuer	Issued To	Certificate	Expiration Date
	DM-VX-B	dm	dm	<a href="#">View</a>	Jan 1 23:19:39 2015 G...

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

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



**TIP** For a historical matrix of Silver Peak security algorithms, click [here](#).

## SSL CA Certificates Tab

*Configuration > [Overlays > SSL] 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 shows the 'SSL CA Certificates' tab in the Silver Peak Orchestrator. The interface includes a search bar, a 'Show' dropdown set to 25, and a table with the following data:

Edit	Appliance Name	Issuer	Issued To	Certificate	Expiration Date
	chateau	Silver Peak SSL Proxy	Silver Peak SSL Proxy	<a href="#">View</a>	Dec 31 23:59:59 2034 GMT

Below the table, a 'View Certificate Content' dialog box is open, displaying 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:d0:cb:19:45:
        12:55:68:cd:fa:8d: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:
        6F:BC:F0:A1:FE:A8
  
```

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.



**TIP** For a historical matrix of Silver Peak security algorithms, click [here](#).

---

## SSL for SaaS Tab

Configuration > [Overlays > SSL] SSL for SaaS

This report lists the appliances' signed substitute certificates.

The screenshot shows the 'SSL for SaaS' configuration page in the Silver Peak Orchestrator. The page has a top navigation bar with 'Topology' and 'SSL for SaaS x' tabs. Below the navigation bar, there are buttons for 'Manage SSL for SaaS with Templates', 'Export', and a refresh icon with '3 mins' next to it. The main content area is titled 'SSL for SaaS' and features a 'Show' dropdown set to '25' and a search box. A table with the following columns is displayed: 'Appliance Name', 'Decrypt SSL', 'Issuer', 'Issued To', 'Certificate', and 'Expiration Date'. The table contains two entries: 'castle' with 'Decrypt SSL' set to 'No', and 'chateau' with 'Decrypt SSL' set to 'Yes', 'Issuer' as 'Silver Peak SSL Proxy', 'Issued To' as 'Silver Peak SSL Proxy', and 'Expiration Date' as 'Dec 31 23:59:59 203...'. A 'View' link is present in the 'Certificate' column for the 'chateau' entry. A dialog box titled 'View Certificate Content' is open, showing the following certificate details:

```

Certificate:
  Data:
    Versions: 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:dc:
        6f:a1:c1:c9:2a:d3:10:c3:26:dc:b3:5e 73:ba:5c:
        ce:ce:d1:d7:13:9d:82:90:9a:58:ca:cb 79:a6:0f:
        f9:0e:2e:43:e1:fc:60:55:4b:4b:05:fa 83:97:b2:
        fe:3f:35:64:23:38:18:72:db:73:8f:7b 98:58:9e:
        rr:e8:20:5f:2b:de:5f:e1:ae:6a:20:ae a2:a5:4b:
        8f:81:47:15:0f:43:8f:1f:4d:7c:db:54 8b:3d:d2:
        93:51:9f:8a:43:45:9a:87:20:23:b8:f9 42:5f:bf:
        bf:52:a2:98:78:8b:4c:60:cc:fb:05:52 53:38:56:
        75:2e:3b:72:22:04:fb:45:1b:b5:63 ce:ad:14:
        af:29:11:19:92:a8:12:00:a9:39:64:bd 1e:ea:2b:
        5a:48:41:73:22:82:4a:7f:fa:a7:a5:b4 50:50:4c:
        a0:5b:32:f2:57:67:87:1a:5d:1f:9d:3a 73:12:df:
        bf:92:4a:80:97:ead:90:d1:b7:e3:e9 8d:1f:d1:
        0f:ec:1b:31:b9:74:31:1b:0c:69:e9:79 5e:ef:b4:
        f8:a3:cc:c0:87:8dc3:54:91:06:60:5d c8:5b:75:
        ba:c1
      Exponent: 65537 (0x10001)
    X509v3 extensions:
      X509v3 Basic Constraints: critical
        CA:TRUE
      X509v3 Key Usage: critical
        Certificate Sign
  
```

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 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).



**TIP** For a historical matrix of Silver Peak security algorithms, click [here](#).

---

# Network Configuration Tabs

These topics describe the pages related to configuring and managing the network.

## DHCP Failover

Configure the following settings to apply to your DHCP failover servers.

1. Check the DHCP Failover box to enable the DHCP Failover feature.
2. Select whether you are configuring the failover settings for either the Primary or Secondary server.
3. Complete configuring the remaining settings in the table below.

### *DHCP Failover Fields*

Field Name	Description
<b>My IP</b>	The IP address of the LAN interface.
<b>My Port</b>	The port number of the LAN interface.
<b>Peer IP</b>	The IP address of the DHCP peer.
<b>Peer Port</b>	The port number of the DHCP peer.
<b>MLCT</b>	Optional. If selected, the default is 60 minutes. This field cannot be zero.
<b>SPLIT</b>	Optional. If selected, determines which peer (primary/secondary) should process the DHCP requests.
<b>Max Response Delay</b>	Optional. If selected, determines how many seconds the DHCP server may pass without receiving a message from its failover peer before it assumes the connection has failed.
<b>Max Unacked Updates</b>	Tells the remote DHCP server how many BNDUPD messages it can send before it receives a BNDACK from the local system.
<b>Load Balance Max Seconds</b>	Optional. Allows you to configure a cutoff after which load balancing is disabled. The cutoff is based on the number of seconds since the client sent its first DHCPDISCOVER or DHCPREQUEST message, and only works with clients that correctly implement the secs field

## DHCP Failover State

EdgeConnect appliances can act as a DHCP server for clients on the LAN side. DHCP failover allows redundancy by creating failover groups when two appliances are combined in an HA configuration. DHCP failover also provides stability if one EdgeConnect dies by allowing the other EdgeConnect HA pair to take over as the DHCP server. To do so, the primary and secondary servers must be completely synchronized so each server can reply on the other if one fails.

This tab displays the DHCP failover peer states of each server for troubleshooting purposes.

### *DHCP Failover Fields*

Field Name	Description
Appliance Name	The name of the Silver Peak appliance that is part of the DHCP failover configuration.
Interface Name	The failover group name that is the same for all the tagged and untagged interfaces corresponding to one physical interface.
My State	The failover endpoint state of the selected primary appliance. The three states are: <b>Normal, Communications-Interrupted, Partner-Down, Recover, Recover-wait, Recover-done.</b>
My State Time	The date and time the selected appliance's DHCP server entered the specified state in the table.
Partner State	The failover endpoint state of the partner appliance. The three states are: <b>Normal, Communications-Interrupted, Partner-Down, Recover, Recover-wait, Recover-done.</b>
Partner State Time	The date and time the partner appliance entered the specified state in the table.
MCLT	The maximum client lead time: the maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner.

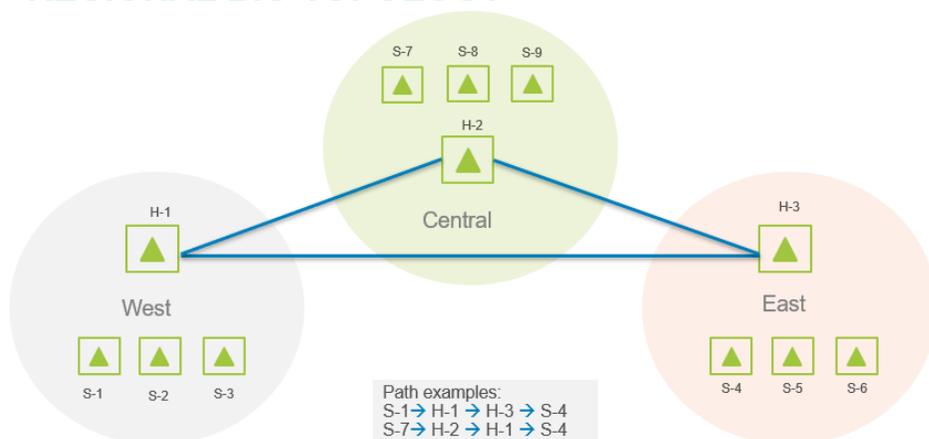
## Regions

Use this page to add or remove regions from the SD-WAN fabric and configure regional routing. The regions within your SD-WAN fabric can represent geographical regions, administrative regions, or a set of sites in the network that have common business goals.

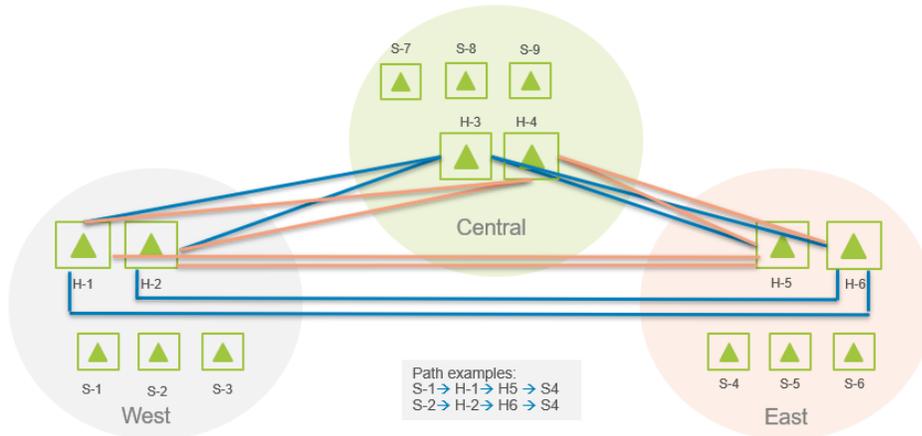
### Regional Routing

Regional routing when enabled, allows you to manage your SD-WAN fabric by regions. It involves intra-region and inter-region route distribution across the SD-WAN fabric. The regions within your network can represent geographical regions, administrative regions, or a set of sites in the network that have common business goals. You can provide different Business Intent Overlay for each region by enabling regional routing and customizing BIODs per region. The following diagrams show examples of different regional network topologies you can build by enabling regional routing.

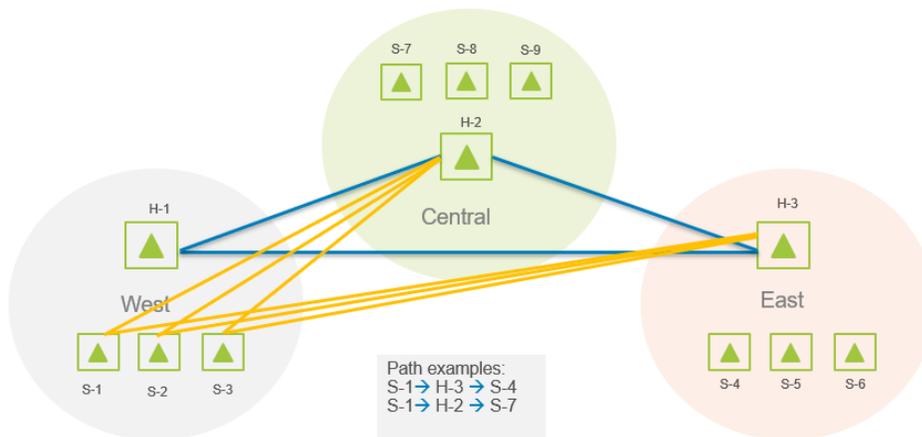
### REGIONAL BIOD TOPOLOGY



## REGIONAL MULTI-HUB BIO TOPOLOGY



## OPTIMIZED REGIONAL BIO TOPOLOGY



You can enable regional routing within your Orchestrator UI. Select the **Enable Regional Routing** icon in the header and move the toggle.

View Status

Select **View Status** to view the status of the added or updated appliances to regions.

## Edit Regions

Complete the following steps to add a region or edit existing regions that you want to add to your overlays.

1. Select **Edit Regions**.
2. Select **New Region**.
3. Enter the name of your new region in the **Region Configuration** window.
4. Select **Save**.

You can also edit an existing region.

1. Select the **Edit** icon next to the region you want to edit.
2. Enter the region name.
3. Select **Save**.

Navigate to the **Business Intent Overlay** tab to make further customizations to your regions and overlays.

## BGP Tab

*Configuration > Networking > BGP*

Use this page to configure **BGP (Border Gateway Protocol)** for appliances, and to add their BGP peers (also known as BGP "neighbors").

- Configuration of advertisement and redistribution rules is peer-based, with the exception of the rules to redistribute BGP routes to BGP peers.
- Silver Peak has the following behaviors relative to **communities**:
  - Although Silver Peak does not configure BGP communities, it does propagate existing communities so that they stay attached to routes, from end to end.
  - Appliances can display up to 10 communities per route.
  - Appliances subnet share communities with their Silver Peak peers.
  - Appliances advertise communities to remote peers, if learned from Silver Peak peers.
  - Appliances advertise communities to BGP neighbors.
- All BGP-learned subnets also appear in the appliance Routes table, displayed on the Routes configuration page. In addition, any AS Path or BGP Community information learned with a particular subnet will also be displayed with that subnet entry in the table.
- BGP route updates aren't refreshed unless the peer specifically asks for it. To update the BGP routes, go to the **Peers** table and select **Soft Reset** in the desired row.

To configure the BGP information for an appliance, select the **Edit** icon, select BGP, and complete the dialog box:

Field	Description
Enable BGP	Allows the exchange of Silver Peak known subnets with BGP peers.
Autonomous System Number (ASN)	Configure this number as needed for your network.
Router ID	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of BGP.

Field	Description				
<b>Enable Graceful Restart</b>	Enable receiver-side graceful restart capability. Silver Peak retains routes learned from the peer and continues to use it for forwarding (if possible) if/when a BGP peer goes down. The retained routes are considered stale routes and will be deleted and replaced with newly received routes.				
	<table border="1"> <tr> <td><b>Max Restart Time</b></td> <td>Specifies the maximum time (in seconds) to wait for a Graceful Restart capable peer to come back after a peer restart or peer session failure.</td> </tr> <tr> <td><b>Stale Path</b></td> <td>Specifies maximum time (in seconds) following a peer restart that SP waits before removing stale routes associated with that peer.</td> </tr> </table>	<b>Max Restart Time</b>	Specifies the maximum time (in seconds) to wait for a Graceful Restart capable peer to come back after a peer restart or peer session failure.	<b>Stale Path</b>	Specifies maximum time (in seconds) following a peer restart that SP waits before removing stale routes associated with that peer.
<b>Max Restart Time</b>	Specifies the maximum time (in seconds) to wait for a Graceful Restart capable peer to come back after a peer restart or peer session failure.				
<b>Stale Path</b>	Specifies maximum time (in seconds) following a peer restart that SP waits before removing stale routes associated with that peer.				
<b>AS Path Prepend</b>	The learned path from an external prepend between a remote BGP site to local BGP peers.				
<b>Redistribute OSPF routes to BGP</b>	Enables subnet sharing of OSPF routes (subnets) to BGP peers.				
	<table border="1"> <tr> <td><b>Filter Tag</b></td> <td>If this tag matches the Route Tag associated with the OSPF route, then that route is not advertised to BGP peers. It's filtered out.</td> </tr> </table>	<b>Filter Tag</b>	If this tag matches the Route Tag associated with the OSPF route, then that route is not advertised to BGP peers. It's filtered out.		
<b>Filter Tag</b>	If this tag matches the Route Tag associated with the OSPF route, then that route is not advertised to BGP peers. It's filtered out.				

To add a **BGP Peer**, select **Add Peer** and complete the following parameters:

Field	Definition
<b>Admin Status</b>	Choose to admin the peer UP or DOWN.
<b>AS Prepend Count</b>	When a BGP router advertises to a neighbor in a different autonomous system (for example, an external or eBGP neighbor), it adds its own AS number to the front (left side) of the AS path. The AS path lists all the ASes that need to be traversed to reach the destination. Preference is given to the shortest AS path.
<b>Enable Imports</b>	Allows the learning of routes from this specific BGP peer.
<b>Enable MD5 Password</b>	Optionally, you can add a password to authenticate the TCP session with the peer.
<b>Hold Timer</b>	When availability to a peer is lost, this specifies how long to wait before dropping the session.

Field	Definition
<b>Input Metric</b>	The metric that is advertised with the route when shared.
<b>Keep Alive Timer</b>	This specifies the interval, in seconds, between keep alive signals to a peer.
<b>Local Preference</b>	The local preference is the first attribute a Cisco router looks at to determine which route towards a certain destination is the “best” one. This value is not exchanged between external BGP routers. Local preference is a discretionary BGP attribute. Default value is 100. The path with the highest local preference is preferred.
<b>Local Interface</b>	A list of the interfaces that can be chosen: <b>Any, lan0, wan0, wan1.</b>
<b>MED</b>	<i>Multi Exit Discriminator.</i> When BGP needs to choose the best route to reach a certain destination, it first looks at the local preference and AS path attributes. When the local preference and AS path length are the same for two or more routes towards a certain prefix, the Multi Exit Discriminator (MED) attribute comes into play. Unlike with the local preference and weight, where higher is more preferred, with MED, the lowest value is preferred. <b>NOTE:</b> If you configured the Metric Delta parameter in an earlier version of our software, this value has been translated into a MED value.
<b>Next-Hop-Self</b>	The advertised route connected to a CE router that an EdgeConnect appliance learns from the eBGP with a PE router.
<b>Peer ASN</b>	Peer's <i>Autonomous System Number.</i>
<b>Peer IP</b>	IP address of the Silver Peak peer.
<b>Peer Type</b>	Governs what kinds of routes the appliance is allowed to advertise to this BGP peer. These routes are itemized as Route Export Policies.

Currently, there are three peer types:

- **Branch** - All route types are permitted
- **Branch-transit** --- A **branch-transit** peer can reach another peer through a "back door" via routes shared through another protocol such as OSPF, ISIS, or BGP. All route types are permitted **except** Remote BGP branch-transit routes (type 7).
- **PE (Provider Edge) Router** – Only BGP branch and BGP branch-transit (types 1, 3, and 4) are permitted.

The **Route Export Policies** are as follows:

1. Locally configured
2. Learned via subnet sharing (from a non-BGP source)
3. Learned from a local BGP branch peer
4. Learned from a local BGP branch-transit peer
5. Learned from a local BGP PE peer
6. Remote BGP (learned via subnet sharing, but originally from a BGP peer)
7. Remote BGP branch-transit (learned via subnet sharing, but originally a BGP branch-transit peer)

**State Details** are used by Silver Peak Support for troubleshooting. A peer can be in one of six possible States:

State	Explanation
<b>Idle</b>	This is the first state where BGP waits for a start event, which occurs when someone configures a new BGP neighbor or when you reset an established BGP peering. BGP initiates a TCP connection with the peer. It also returns to this state when an error occurs or the connection is disrupted, in an effort to bring the session back up.
<b>Connect</b>	BGP is waiting for the BGP peer to finish establishing the TCP connection.
<b>Active</b>	The TCP connection with the BGP peer has begun but has not completed
<b>OpenSent</b>	The TCP connection with the peer has successfully completed. It now waits for a BGP Open message from the peer.
<b>OpenConfirm</b>	An Open message has been received from the peer and successfully validated. It now waits for a keepalive message from the peer.
<b>Established</b>	This is the final state, and indicates that the BGP neighbor adjacency is complete. At this point, the BGP routers send and receive BGP Update packets for the purpose of exchanging routing information. If an error occurs in the session, it returns to the Idle state to try to return the peer session to the Established state.

## BGP Edit Row

Use this page to configure **BGP (Border Gateway Protocol)** for appliances, and to add their BGP peers (also known as BGP "neighbors").

- Configuration of advertisement and redistribution rules is peer-based, with the exception of the rules to redistribute BGP routes to BGP peers.
- Silver Peak has the following behaviors relative to **communities**:
  - Although Silver Peak doesn't configure BGP communities, it does propagate existing communities so that they stay attached to routes, from end to end.
  - Appliances can display up to 10 communities per route.
  - Appliances subnet share communities with their Silver Peak peers.
  - Appliances advertise communities to remote peers, if learned from Silver Peak peers.
  - Appliances advertise communities to BGP neighbors.
- All BGP-learned subnets also appear in the appliance Routes table, displayed on the Routes configuration page. In addition, any AS Path or BGP Community information learned with a particular subnet will also be displayed with that subnet entry in the table.
- BGP route updates are not refreshed unless the peer specifically asks for it. To update the BGP routes, go to the **Peers** table and select **Soft Reset** in the desired row.

To configure BGP for an appliance:

1. Select the **Edit** icon
2. Select **Enable BGP**.
3. Complete the fields in the dialog box:

Field	Description
<b>Enable BGP</b>	Allows the exchange of Silver Peak known subnets with BGP peers.
<b>Autonomous System Number (ASN)</b>	Configure this number as needed for your network.
<b>Router ID</b>	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of BGP.
<b>Redistribute OSPF routes to BGP</b>	Enables subnet sharing of OSPF routes (subnets) to BGP peers. <b>Filter Tag</b> If this tag matches the Route Tag associated with the OSPF route, then that route is not advertised to BGP peers. It's filtered out.

To add a **BGP Peer**, select **Add** and complete the following parameters:

Field	Definition
<b>Admin Status</b>	Choose to admin the peer UP or DOWN.
<b>AS Prepend Count</b>	When a BGP router advertises to a neighbor in a different autonomous system (for example, an external or eBGP neighbor), it adds its own AS number to the front (left side) of the AS path. The AS path lists all the ASes that need to be traversed to reach the destination. Preference is given to the shortest AS path.
<b>Enable Imports</b>	Allows the learning of routes from this specific BGP peer.
<b>Enable MD5 Password</b>	Optionally, you can add a password to authenticate the TCP session with the peer.
<b>Hold Timer</b>	When reachability to a peer is lost, this specifies how long to wait before dropping the session.
<b>Input Metric</b>	The metric that is advertised with the route when shared.
<b>Keep Alive Timer</b>	This specifies the interval, in seconds, between keep alive signals to a peer.
<b>Local Interface</b>	You can specify the source address or interface for a specific BGP peer. Select the interface from the menu: <b>Any</b> , <b>lan0</b> , <b>wan0</b> , <b>wan1</b> .
<b>Local Preference</b>	The local preference is the first attribute a Cisco router looks at to determine which route towards a certain destination is the “best” one. This value is not exchanged between external BGP routers. Local preference is a discretionary BGP attribute. Default value is 100. The path with the highest local preference is preferred.
<b>MED</b>	<i>Multi Exit Discriminator.</i> When BGP needs to choose the best route to reach a certain destination, it first looks at the local preference and AS path attributes. When the local preference and AS path length are the same for two or more routes towards a certain prefix, the Multi Exit Discriminator (MED) attribute comes into play. Unlike with the local preference and weight, where higher is more preferred, with MED, the lowest value is preferred. <b>NOTE:</b> If you configured the Metric Delta parameter in an earlier version of our software, this value has been translated into a MED value.
<b>Peer ASN</b>	Peer's <i>Autonomous System Number</i>
<b>Peer IP</b>	IP address of the Silver Peak peer

Field	Definition
Peer Type	<p>Governs what kinds of routes the appliance is allowed to advertise to this BGP peer. These routes are itemized as Route Export Policies.</p> <p>Currently, there are three peer types:</p> <ul style="list-style-type: none"> <li>■ <b>Branch</b> - All route types are permitted</li> <li>■ <b>Branch-transit</b> --- A <b>branch-transit</b> peer can reach another peer through a "back door" via routes shared through another protocol such as OSPF, ISIS, or BGP. All route types are permitted <b>except</b> Remote BGP branch-transit routes (type 7).</li> <li>■ <b>PE (Provider Edge) Router</b> – Only BGP branch and BGP branch-transit (types 1, 3, and 4) are permitted.</li> </ul> <p>The <b>Route Export Policies</b> are as follows:</p> <ol style="list-style-type: none"> <li>1. Locally configured</li> <li>2. Learned via subnet sharing (from a non-BGP source)</li> <li>3. Learned from a local BGP branch peer</li> <li>4. Learned from a local BGP branch-transit peer</li> <li>5. Learned from a local BGP PE peer</li> <li>6. Remote BGP (learned via subnet sharing, but originally from a BGP peer)</li> <li>7. Remote BGP branch-transit (learned via subnet sharing, but originally a BGP branch-transit peer)</li> </ol>

**State Details** are used by Silver Peak Support for troubleshooting. A peer can be in one of six possible States:

State	Explanation
Idle	This is the first state where BGP waits for a start event, which occurs when someone configures a new BGP neighbor or when you reset an established BGP peering. BGP initiates a TCP connection with the peer. It also returns to this state when an error occurs or the connection is disrupted, in an effort to bring the session back up.
Connect	BGP is waiting for the BGP peer to finish establishing the TCP connection.
Active	The TCP connection with the BGP peer has begun but has not completed
OpenSent	The TCP connection with the peer has successfully completed. It now waits for a BGP Open message from the peer.
OpenConfirm	An Open message has been received from the peer and successfully validated. It now waits for a keepalive message from the peer.
Established	This is the final state, and indicates that the BGP neighbor adjacency is complete. At this point, the BGP routers send and receive BGP Update packets for the purpose of exchanging routing information. If an error occurs in the session, it returns to the Idle state to try to return the peer session to the Established state.

## Virtual Tunnel Interface

A VTI (Virtual Tunnel Interface) is a tunneling protocol that does not require a static mapping of IPsec sessions to a physical interface. The tunnel endpoint is associated with a tunnel interface that enables a constant secure and stable connection throughout your network.

Select the **Edit** icon to get started configuring your VTIs.

## VTI

Complete the following steps to configure a VTI with an associated tunnel in Orchestrator.

1. Select **Add**.
2. The **Add Interface** window appears. Complete the following fields with the appropriate information.

Field	Definition
<b>Interface</b>	The name of the VTI interface.
<b>IP/Mask</b>	The IP and mask address of your VTI.
<b>Admin</b>	Select whether the admin is up or down.
<b>Passthrough Tunnel</b>	The name of the passthrough tunnel associated with your VTI.
<b>Direction</b>	The direction of the label for your VTI: LAN or WAN.
<b>Label</b>	The label on your VTI interface.
<b>Zone</b>	Select the zone from the drop-down list that you are applying the VTI to.

3. Select **Add**.

## Boost

This page shows you various details regarding your boost. You can purchase additional boost for the traffic within your network. You can also search for the boost used per appliance by the hour or specify a time frame within the **Range** field at the top of the page.

The following table shows the fields and definitions regarding your boost.

Field	Definition
Appliance	The name of the appliance you are applying boost to.
% Time Insufficient Boost	The percent of time the appliance did not have enough boost configured to boost the traffic.
Minutes Insufficient Boost	The amount of time (in minutes) the appliance did not have enough boost configured to boost the traffic.
Configured Boost (Kbps)	The amount of boost configured on the appliance.
Average Boost Bytes	The average boost in bytes.
Trends	A graph displaying the trends of your boost.

To configure or update your boost:

1. Select the appliance you want to add more or less boost to from the table in the **Boost** tab.
2. Select **Increase 20%**, **Decrease 20%**, or **Set to this Value**.
3. If you select **Set to this Value**, enter the exact amount in the field.
4. Select **Apply**.

# Deployment Tab

*Configuration > Networking > Deployment*

This page summarizes the appliance **Deployment** settings in the following two views:

## Summary view

Deployment x
2 mins

Edit Deployment Profiles
Summary
Details
Export

Deployment ?

9 Rows, 1 Selected
Search

Edit	Appliance Name ▲	Mode	WAN Labels Used	LAN Labels Used	Details
	Chennai	Bridge	MPLS	VoIP	
	Chicago	Bridge	MPLS	VoIP	
	London	Bridge	MPLS	VoIP	
	Los-Angeles	Bridge	MPLS, Internet	VoIP, Data	
	Miami	Inline Router	MPLS	VoIP	
	Minneapolis	Inline Router	MPLS	VoIP	
	Mumbai	Bridge	MPLS	VoIP	
	Munich	Bridge	MPLS	VoIP	
	Portland	Bridge	MPLS, Internet	VoIP, Data	

Deployment Details
x

Los-Angeles
Export

4 Rows
Search

Interface	Label	IP/Mask	WAN/LAN St...	Next Hop	Public IP	Shaping (Kbps)		NAT	Firewall	DHCP
						Inbound	Outbound			
wan0	MPLS	10.1.11.10/24	WAN	10.1.11.1		0	20000	No	Allow All	No
wan1	Internet	10.1.44.10/24	WAN	10.1.44.1		0	20000	No	Allow All	No
lan0	VoIP		LAN			0	0	No	N/A	No
lan1	Data		LAN			0	0	No	N/A	No

## Details view

Deployment ×

Edit Deployment Profiles

Summary **Details** Export ↻

Deployment ?

22 Rows Search

Edit	Appliance ...	Interface	Label	IP/Mask	WAN/LAN S...	Next Hop	Public IP	Shaping (Kbps)		NAT	Firewall	DHCP
								Inbound	Outbound			
✎	Chennai	lan0	VoIP		LAN			0	0	No	N/A	No
✎	Chennai	wan0	MPLS	10.1.83.20/24	WAN	10.1.83.1		0	4000	No	Allow All	No
✎	Chicago	lan0	VoIP		LAN			0	0	No	N/A	No
✎	Chicago	wan0	MPLS	10.1.12.20/26	WAN	10.1.12.1		50000	50000	No	Allow All	No
✎	London	lan0	VoIP		LAN			0	0	No	N/A	No
✎	London	wan0	MPLS	10.1.185.20/24	WAN	10.1.185.1		0	4000	No	Allow All	No
✎	Los-Angeles	lan0	VoIP		LAN			0	0	No	N/A	No
✎	Los-Angeles	lan1	Data		LAN			0	0	No	N/A	No
✎	Los-Angeles	wan0	MPLS	10.1.11.10/24	WAN	10.1.11.1		0	20000	No	Allow All	No
✎	Los-Angeles	wan1	Internet	10.1.44.10/24	WAN	10.1.44.1		0	20000	No	Allow All	No
✎	Miami	lan0	VoIP	10.1.104.10/24	LAN	0.0.0.0		0	0	No	N/A	No
✎	Miami	wan0	MPLS	10.1.103.10/24	WAN	10.1.103.1		0	4000	No	Allow All	No
✎	Minneapolis	lan0	VoIP	10.1.102.10/24	LAN	0.0.0.0		0	0	No	N/A	No
✎	Minneapolis	wan0	MPLS	10.1.101.10/24	WAN	10.1.101.1		0	4000	No	Allow All	No
✎	Mumbai	lan0	VoIP		LAN			0	0	No	N/A	No
✎	Mumbai	wan0	MPLS	10.1.84.20/24	WAN	10.1.84.1		0	4000	No	Allow All	No
✎	Munich	lan0	VoIP		LAN			0	0	No	N/A	No
✎	Munich	wan0	MPLS	10.1.186.20/24	WAN	10.1.186.1		0	4000	No	Allow All	No
✎	Portland	lan0	VoIP		LAN			0	0	No	N/A	No
✎	Portland	lan1	Data		LAN			0	0	No	N/A	No

## Summary View

Field	Definition
Appliance Name	The name of the appliance that was deployed.
Mode	One of four deployment <b>Modes</b> displays: <ul style="list-style-type: none"> <li><b>Router</b>                      Single or dual WAN interfaces share LAN and WAN data traffic.</li> <li><b>InLine Router</b>              Uses separate LAN and WAN interfaces to route data traffic.</li> <li><b>Bridge</b>                        Uses a virtual interface, <b>bvi</b>, created by binding the WAN and LAN interfaces.</li> <li><b>Server</b>                         Both management and data traffic use the <b>mgmt0</b> interface.</li> </ul>
WAN Labels Used	Identify the service, such as <b>MPLS</b> or <b>Internet</b> .

Field	Definition
LAN Labels Used	Identify the data, such as <i>data</i> , <i>VoIP</i> , or <i>replication</i> .
Details	Select the <b>information</b> icon to view further deployment details of an appliance.

## Interfaces Tab

Configuration > Networking > Interfaces

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

Interfaces ?

All Hardware Dynamic

48 Rows Search

Edit	Appliance Name	Name	Status	IP Address/Mask	Public IP	DHCP	Speed	Duplex	MTU	MAC Address
✓	Chennai	lan0	up	10.17.17.20/24		No	10000Mb/s (auto)	full (auto)	1500	00:0C:29:71:58:36
✓	Chennai	lan1	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Chennai	lan2	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Chennai	mgmt0	up	10.0.185.29/24		Yes	10000Mb/s (auto)	full (auto)	1500	00:0C:29:71:58:2C
✓	Chennai	mgmt1	down	169.254.0.1/16		No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Chennai	wan0	up	10.17.18.20/24		No	10000Mb/s (auto)	full (auto)	1500	00:0C:29:71:58:40
✓	Chennai	wan1	up	10.0.184.154/24		Yes	10000Mb/s (auto)	full (auto)	1500	00:0C:29:71:58:4A
✓	Chennai	wan2	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Mumbai	lan0	up	10.17.47.20/24		No	10000Mb/s (auto)	full (auto)	1500	00:0C:29:09:61:92
✓	Mumbai	lan1	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Mumbai	lan2	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Mumbai	mgmt0	up	10.0.185.44/24		Yes	10000Mb/s (auto)	full (auto)	1500	00:0C:29:09:61:88
✓	Mumbai	mgmt1	down	169.254.0.1/16		No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Mumbai	wan0	up	10.17.48.20/24		No	10000Mb/s (auto)	full (auto)	1500	00:0C:29:09:61:9C
✓	Mumbai	wan1	up	10.0.184.226/24		Yes	10000Mb/s (auto)	full (auto)	1500	00:0C:29:09:61:A6
✓	Mumbai	wan2	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Osaka	lan0	up	10.17.43.20/24		No	10000Mb/s (auto)	full (auto)	1500	00:0C:29:51:20:CE
✓	Osaka	lan1	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Osaka	lan2	up			No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Osaka	mgmt0	up	10.0.185.42/24		Yes	10000Mb/s (auto)	full (auto)	1500	00:0C:29:51:20:C4
✓	Osaka	mgmt1	down	169.254.0.1/16		No	1000Mb/s (auto)	full (auto)	1500	Unassigned
✓	Osaka	wan0	up	10.17.44.20/24		No	10000Mb/s (auto)	full (auto)	1500	00:0C:29:51:20:D8
✓	Osaka	wan1	up	10.0.184.224/24		Yes	10000Mb/s (auto)	full (auto)	1500	00:0C:29:51:20:E2

Please refer to the following table for the Interfaces field descriptions.

Field Name	Description
Appliance Name	The name of the appliance of an interface.
Name	The name of the LAN/WAN interface selected.
Status	The status of the interface (up or down.)
IP Address/Mask	The IP address.
Public IP	The public IP address.
DHCP	Whether this interface's IP address is obtained from the DHCP server. Displays as <b>Yes</b> , <b>No</b> , <b>No data</b> (not configured), or <b>Invalid data</b> (error condition).
Speed	The current interface speed.
Duplex	The current interface duplex.
MTU	The maximum number of packets being transmitted.

Field Name	Description
MAC Addresses	The MAC addresses applied to an interface.

- **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, select **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

Interface	Description
<b>blan</b>	Bonded LAN interfaces (as in <b>lan0 + lan1</b> )
<b>bwan</b>	Bonded WAN interfaces (as in <b>wan0 + wan1</b> )
<b>bvi</b>	Bridge Virtual Interface. When the appliance is deployed in-line (Bridge mode), it's the routed interface that represents the bridging of <b>wan0</b> and <b>lan0</b> .
<b>tlan</b>	10-Gbps fiber LAN interface
<b>twan</b>	10-Gbps fiber WAN interface

## Routes Tab

*Configuration > Networking > Routes*

Each appliance builds a routes table from one of the three following ways: entries added automatically by the system, added manually by a user, or learned from BGP (Border Gateway Protocol) or OSPF (Open Shortest Path First). When two appliances are connected by a tunnel, they exchange this information and use it to route traffic.

You can filter between All, Local / Static, SD-WAN Fabric, BGP, and OSPF routes. You can also import or export subnets to a .csv file.

### Filter by Subnet

Filter by subnet is a filtering tool that can be used to filter all existing routes and the results are populated in the **Routes** tab.

A **Very Large Query Response** pop-up will display if the number of the routes filtered exceeds 500,000. You can filter by subnet, cancel, or continue waiting to help mitigate this issue.

**Note:** If the number of the routes filtered is greater than 500,000 the following pop-up will display.

**Very Large Query Response**
×

---

More than 100,000 responses have been returned.  
This could make the user interface unresponsive.

Filter query to this subnet

---

Apply Filter
Cancel Query
Continue Waiting

You can filter by subnet, cancel, or continue waiting to help mitigate the above issue.

The following represents the content in the Routes table.

Field	Definition
Appliance Name	The name of the appliance.
Subnet/Mask	Actual subnet to be shared or learned.
Next Hop	The deployment interface's next-hop address.

Field	Definition
<b>Interface</b>	The interface for outgoing traffic. Display only.
<b>State</b>	Shows if the route is either up or down.
<b>Metric</b>	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 lower numerical value.
<b>Exclude</b>	Use this option to prevent optimization of more specific subnets from a wider advertised subnet range.
<b>Advertise to Peers</b>	<p>Select to share subnet information with categories of peers. Options are:</p> <ul style="list-style-type: none"> <li>■ <b>Advertise to Silver Peak Peers</b></li> <li>■ <b>Advertise to BGP Peers</b></li> <li>■ <b>Advertise to OSPF Peers</b></li> </ul> <p>Peers then learn the subnets. To add a subnet to the table without divulging it to peers, deselect this option.</p>

Field	Definition	
Type (of route)	<b>Auto (System)</b>	Automatically added subnets of interfaces on this appliance
	<b>Auto (Saas)</b>	Automatically added subnets from SaaS services
	<b>Added by user</b>	Manually added/configured subnets for this appliance
	<b>SP: Hostname</b>	Subnets added as a result of exchanging information with peer appliances. If the peer has learned the subnet from a remote BGP or OSPF peer, that information is appended.
	<b>&lt;BGP peer Type&gt;: &lt;BGP peer ip&gt;</b>	Subnets added as a result of exchanging information with local BGP peers
	<b>OPSF: OPSF neighbor IP</b>	Subnets added as a result of exchanging information with local OSPF peers.



## OSPF Tab

*Configuration > [Networking] OSPF*

Use this page to manage **OSPF (Open Shortest Path First)** on LAN and WAN interfaces.

This protocol learns routes from routing peers, and then subnet shares them with Silver Peak peers and/or BGP neighbors.

### Summary View

OSPF						
1 Rows						
Edit	Appliance Name	Enable	Router ID	Redistribute from BGP	Redistribute from Local	Redistribute from Silver Peak peers
<input checked="" type="checkbox"/>	spucha-ecvd	Yes	10.18.79.10	Redist:Yes, Metric:0, Tag:0, Metric Type:2	Redist:Yes, Metric:0, Tag:0, Metric Type:2	Redist:Yes, Metric:0, Tag:0, Metric Type:2

## Enabling OSPF

To enable or disable OSPF for an appliance, select a table row, click the Edit icon, and complete the fields.

OSPF - spucha-ecvd

**OSPF** ?

Enable OSPF

Router ID

Redistribute BGP routes to OSPF

Metric Type

Metric

Tag

Redistribute Silver Peak peers routes to OSPF

Metric Type

Metric

Tag

Redistribute local routes to OSPF

Metric Type

Metric

Tag

**Interfaces**

Edit	Interface	Area ID	Cost	Retransmit Interval	Transmit Delay	Priority	Admin Status	Hello Interval	Dead Interval	Comment
<input checked="" type="checkbox"/>	lan0	0.0.0.0	1	4	1	1	Up	10	40	X

### Enable OSPF

Field	Definition
<b>Enable OSPF</b>	When enabled, the appliance has access to use the OSPF protocol.
<b>Metric</b>	[Route Metric] The cost associated with a route. The higher the value, the less preferred.
<b>Metric Type</b>	<p><b>E1</b> and <b>E2</b> are two categories of external routes. The difference between the two is in the way the cost (metric) of the route is being calculated.</p> <p><b>E1</b> Sum of the external cost and the internal cost used to reach a route.</p> <p><b>E2</b> The external cost to reach a route, irrespective of the interior cost. Unless otherwise specified, the default external type given to external routes is E2.</p>
<b>Route redistribution</b>	<p>Redistributing routes into OSPF from other routing protocols or from <b>static</b> will cause these routes to become OSPF external routes. Silver Peak summarizes the following:</p> <ul style="list-style-type: none"> <li>• <b>Redistribute BGP routes to OSPF</b></li> <li>• <b>Redistribute Silver Peak peers routes to OSPF</b></li> <li>• <b>Redistribute local routes to OSPF</b></li> </ul>

Field	Definition
Router ID	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of OSPF.
Tag	A route tag is applied to a route to better identify the source of the network it originated from. It is primarily used to filter routes from being redistributed in a routing loop.

### OSPF Detail

Field	Definition
Backup Designated Router	Indicates whether a Backup Designated Router (BDR) is specified for the Designated Router (DR). Options are <b>Yes</b> or <b>No</b> .
External learned OSPF Routes	Number of external OSPF routes advertised to the Silver Peak.
External LSAs	<i>External Link-State Advertisement</i> . Also known as <i>Type 5 LSAs</i> , these LSAs contain information imported into OSPF from other routing processes. They are flooded to all areas unchanged (except stub and NSSA areas).
External LSA Refresh Interval	When an OSPF link state advertisement (LSA) age reaches the link-state refresh time (1800 seconds), the OSPF updates the LSAs for advertisement. Defined in RFC 2328, the 1800-second interval cannot be changed.
Link state hold interval	The minimum hold time interval between two LSA generations. The default value is 5000 milliseconds and it doubles every time the same LSA has to be regenerated.
Minimum LSA interval	The OSPF Link-State Advertisement (LSA) Throttling feature provides a dynamic mechanism to slow down link-state advertisement (LSA) updates in OSPF during times of network instability. It also allows faster Open Shortest Path First (OSPF) convergence by providing LSA rate limiting in milliseconds.  The first LSA is always generated immediately upon an OSPF topology change, and the next LSA generated is controlled by the <b>minimum start interval</b> . The subsequent LSAs generated for the same LSA are rate-limited until the maximum interval is reached. The <b>same LSA</b> is defined as an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID.
Minimum LSA arrival	Controls the minimum interval for accepting the same LSA. If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped. It is recommended that the arrival interval be less than or equal to the hold-time interval of the <b>timers throttle lsa all</b> command.
Originated new LSAs	The number of new link state advertisements that have been originated by the router.

Field	Definition
Route calculation max delay	Maximum wait between two SPF calculations. Range is 1 to 600000 milliseconds.

## Adding an Interface

### Interfaces View

OSPF x

Routes | BGP | OSPF | Peer Priority | Admin Distance

Summary Interfaces Neighbors Export

OSPF

1 Rows 50 % Search

Edit	Appliance Name	Interface	Admin Status	Interface State	Area ID	Cost	Priority	Comment	Details
<input checked="" type="checkbox"/>	spucha-ecvd	lan0	Up	Backup Designated Router	0.0.0.0	1	1		

To add an interface to OSPF, select a table row, click the Edit icon, and then click **Add** above the Interfaces table. After completing the fields in the **Add Interface** dialog, click **Add**.

**Add Interface**
✕

---

Interface ▼  
wan0

Area ID = 0  
0.0.0.0  
(Area ID is the same for all interfaces)

Cost 1

Priority 1

Admin Status UP DOWN

Hello Interval 10 (0..65535) Sec

Dead Interval 40 (0..65535) Sec

Transmit Delay 1 (0..65535) Sec

Retransmit Interval 4 (0..65535) Sec

Authentication None ▼

Comment

Add
Close

### Add Interface

Field	Definition
Admin Status	Indicates whether the interface is set to admin <b>UP</b> or <b>DOWN</b> .
Area ID	The number of the area in which to locate the interface. The Area ID is the same for all interfaces. It can be an integer between 0 and 4294967295, or it can take a form similar to an IP address, A.B.C.D.

Field	Definition
<b>Authentication</b>	<b>None</b> No authentication
	<b>Text</b> Simple password authentication allows a password (key) to be configured per area. Routers in the same area that want to participate in the routing domain will have to be configured with the same key. The drawback of this method is that it is vulnerable to passive attacks.
	<b>MD5</b> Message Digest authentication is a cryptographic authentication. A key (password) and key-id are configured on each router. The router uses an algorithm based on the OSPF packet, the key, and the key-id to generate a "message digest" that gets appended to the packet. Unlike the simple authentication, the key is not exchanged over the wire. A non-decreasing sequence number is also included in each OSPF packet to protect against replay attacks.  This method also allows for uninterrupted transitions between keys. This is helpful for administrators who wish to change the OSPF password without disrupting communication.
<b>Comment</b>	Any information you want to include for your own use.
<b>Cost</b>	The cost of an interface in OSPF is an indication of the overhead required to send packets across a certain interface. It's used in the OSPF path calculation to determine link preference.
<b>Dead Interval</b>	Number of seconds that a router's Hello packets have not been seen before its neighbors declare the OSPF router down.
<b>Hello Interval</b>	Specifies the length of time, in seconds, between the hello packets that a router sends on an OSPF interface.
<b>Interface</b>	The interface on which OSPF is configured.
<b>Priority</b>	Router priority [If two or more best routes are subnet shared, then peer priority is used as the tie-breaker.]
<b>Retransmit Interval</b>	When OSPF sends an advertisement to an adjacent router, it expects to receive an acknowledgment from that neighbor. If no acknowledgment is received, the router will retransmit the advertisement to its neighbor. The retransmit-interval timer controls the number of seconds between retransmissions.
<b>Transmit Delay</b>	Number of seconds required to transmit a link state update packet. Valid values are 1 to 65535.

## Viewing the OSPF Neighbors

### Neighbors View

The screenshot shows the OSPF Neighbors View. At the top, there are navigation tabs: Routes | BGP | OSPF | Peer Priority | Admin Distance. Below these are sub-tabs: Summary | Interfaces | Neighbors | Export. The main content area displays a table with the following data:

Edit	Appliance Name	Local IP	Neighbor IP	Router ID	State	Role	Area	Dead Time	Up Time	Details
<input checked="" type="checkbox"/>	spucha-ecvd	10.18.75.10 (lan0)	10.18.75.253	10.18.75.253	Full	Designated Router	0.0.0.0	30s	1h 1m 20s	

### Neighbors View

Field	Definition
<b>Appliance Name</b>	Name of the appliance
<b>Area</b>	Identifies the area to which the packet belongs. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.0.0.0.
<b>Dead Time</b>	The time in seconds after which a non-responding neighbor is considered dead.
<b>Details</b>	Active state information of the routing process.
<b>Local IP</b>	The IP address of the Silver Peak appliance interface.
<b>Neighbor IP</b>	The IP address of the directly connected neighbor.
<b>Role</b>	<p><b>Designated Router</b> Highest router priority on the interface.</p> <p><b>Backup Designated Router</b> Usually, the second highest router priority on the interface.</p> <p><b>Other</b> Cannot be a Designated Router. Indicated by Router Priority = 0.</p>
<b>Router ID</b>	This router identifier is the IPv4 address by which the remote peer can identify this device for purposes of OSPF.

Field	Definition
<b>State</b>	<p>When OSPF adjacency is formed, a router goes through several state changes before it becomes fully adjacent with its neighbor. They are as follows:</p> <p><b>Down</b> This is the first OSPF neighbor state. It means that no information (hellos) has been received from this neighbor, but hello packets can still be sent to the neighbor in this state.</p> <p><b>Attempt</b> This state is only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends multicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.</p> <p><b>Init</b> This state specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet.</p> <p><b>2-Way</b> This state designates that bi-directional communication has been established between two routers. Bi-directional means that each router has seen the other's unicast hello packet.</p> <p><b>Exstart</b> In this state, the routers and their DR and BDR establish a master-slave relationship and choose the initial sequence number for adjacency formation. The router with the higher router ID becomes the master and starts the exchange, and as such, is the only router that can increment the sequence number.</p> <p><b>Exchange</b> OSPF routers exchange database descriptor (DBD) packets. Database descriptors contain link-state advertisement (LSA) headers only and describe the contents of the entire link-state database.</p> <p><b>Loading</b> Actual exchange of link state information occurs. Based on the information provided by the DBDs, routers send link-state request packets. The neighbor then provides the requested link-state information in link-state update packets.</p> <p><b>Full</b> Routers are fully adjacent with each other. All the router and network LSAs are exchanged and the routers' databases are fully synchronized. In other words, the neighbor is "up".</p>

## Multicast

Orchestrator supports multicast routing, a method of sending data from a single IP address to a larger group of recipients. This is only supported in Inline Router mode. There are three different ways you can display the status of multicast: **Interfaces**, **Neighbors**, and **Routes**.

From Summary, Interfaces, Neighbors, or Routes view:

1. Select the **Edit** icon.
2. Select **Enable Multicast**.
3. Enter the Rendezvous Point IP Address.

### Interfaces

Select **Add** to add an interface. The **Add Interface** window appears.

1. Select the **Interface** field and select the desired interface from the list.
2. Select if you want to **Enable PIM**.
3. Select if you want to **Enable IGMP**.
4. Select **Add**.

Field	Definition
Interface	The name of the interfaces you want to connect.
PIM Enabled	Enabling the Protocol Independent Multicast.
IGMP	Enabling Internet Group Management Protocol.
DR Priority	The designated router priority of the given interface.
DR Router IP	The IP address of the designated router within your network.

### Neighbors

Field	Definition
Appliance Name	The name of the appliance you are using for multicast.
Interface	The name of the interfaces you want to connect.
Neighbor DR Priority	The designated router priority of the neighbor.
Neighbor IP	The IP address of the neighbor.

### Routes

Field	Definition
Source	The transmitter of the multicast data.
Group	The IP address of the multicast group.
Incoming Interface	The interface that receives inbound traffic.
Outgoing Interfaces	The interface that receives outbound traffic.

You can also export an excel file of the multicast report, as well as refresh the page and the information from each appliance.

## Loopback

*Configuration > Routing > Loopback*

The loopback features enhances reliability and security by allowing you to access your network using one, static, IP address. If one interface goes down, You can access all interfaces through the one, static IP address. To add a loopback interface to your network, do the following:

1. Navigate to the **Loopback** tab in Orchestrator.
2. Select the **Edit** icon.
3. Select **Add**.
4. Enter the appropriate information for your loopback interface in the **Add Interface** window.

## Peer Priority Tab

Configuration > [Networking] Peer Priority

When an appliance receives a **Subnet** with the same **Metric** from multiple remote/peer appliances, it uses the Peer Priority list as a tie-breaker.

- If a **Peer Priority** is not configured, then the appliance randomly distributes flows among multiple peers.
- The lower the number, the higher the peer's priority.

Peer Priority ×

Routes | BGP | OSPF | Peer Priority | Admin Distance
Export
↺ ▼
Manage Peer Priority with Templates

### Peer Priority ?

30 Rows
Search

Edit	Appliance Name	Peer Name	Priority ▼
	Mumbai	Peer Priority not configured on the Appliance.	
	Osaka	Peer Priority not configured on the Appliance.	
	Chennai	Peer Priority not configured on the Appliance.	
	Seoul	Peer Priority not configured on the Appliance.	
	Singapore	Peer Priority not configured on the Appliance.	
	Tokyo	Peer Priority not configured on the Appliance.	
	Barcelona	Peer Priority not configured on the Appliance.	
	Edinburgh	Peer Priority not configured on the Appliance.	
	Frankfurt	Peer Priority not configured on the Appliance.	
	Geneva	Peer Priority not configured on the Appliance.	
	Milan	Peer Priority not configured on the Appliance.	
	Paris	Peer Priority not configured on the Appliance.	
	London	Peer Priority not configured on the Appliance.	
	Boston	Peer Priority not configured on the Appliance.	

**Note:** By default, the peer priority range starts at 1.

## Admin Distance Tab

*Configuration > Networking > Admin Distance*

This table shows the values associated with various types of **Admin Distance**.

Edit	Appliance Name	Local	BGP Branch	BGP Transit	BGP PE	Subnet Shared - Static Routes	Subnet Shared - BGP Remote	OSPF	Subnet Shared - OSPF Remo...
✓	Toronto	1	20	20	25	10	20	20	20
✓	Tokyo	1	20	20	25	10	20	20	20
✓	Singapore	1	20	20	25	10	20	20	20
✓	Seoul	1	20	20	25	10	20	20	20
✓	San-Jose	1	20	20	25	10	20	20	20
✓	San-Antonio	1	20	20	25	10	20	20	20
✓	Salt-Lake-City	1	20	20	25	10	20	20	20
✓	Portland	1	20	20	25	10	20	20	20
✓	Pittsburgh	1	20	20	25	10	20	20	20
✓	Paris	1	20	20	25	10	20	20	20
✓	Osaka	1	20	20	25	10	20	20	20
✓	New-York	1	20	20	25	10	20	20	20
✓	New-Orleans	1	20	20	25	10	20	20	20
✓	Mumbai	1	20	20	25	10	20	20	20
✓	Mississippi	1	20	20	25	10	20	20	20

Admin Distance (AD) is the route preference value assigned to dynamic routes, static routes, and directly connected routes. When the appliance's Routes table has multiple routes to the same destination, the appliance uses the route with the lowest administrative distance.

Field	Description
Local	A manually configured route, or one learned from locally connected subnets.
BGP Branch	A type of dynamic route learned from a local BGP branch peer.
BGP PE	A type of dynamic route learned from a local BGP PE (Provider Edge) router.
BGP Remote	A route learned from a BGP peer.
BGP Transit	A type of dynamic route learned from a local BGP branch-transit peer.
OSPF	A route learned from an OSPF (Open Shortest Path First) neighbor.
Subnet Shared	A route learned from a Silver Peak peer.

## Management Routes Tab

Configuration > [Networking] Management Routes

Use this tab to configure **next-hops** for management interfaces.

Management Routes ? ↻

Add new route

6 Rows Search

Subnet ▼	Next-hop IP	Interface	Source IP	Metric
10.17.46.0/24	0.0.0.0	wan0	0.0.0.0	100
10.17.45.0/24	0.0.0.0	lan0	0.0.0.0	100
10.0.185.0/24	0.0.0.0	mgmt0	10.0.185.43	
10.0.184.0/24	0.0.0.0	wan1	0.0.0.0	100
0.0.0.0/0	10.17.46.1	wan0	0.0.0.0	253
0.0.0.0/0	10.0.185.1	mgmt0	0.0.0.0	252

- Management routes specify the **default gateways** and local IP subnets for the management interfaces.
- In a Dual-Homed Router Mode configuration, you may need to add a static management route for flow redirection between appliances paired for redundancy at the same site.
- The management routes table shows the configured static routes and any dynamically created routes. If you use **DHCP**, then the appliance automatically creates appropriate dynamic routes. A user cannot delete or add dynamic routes.
- If the **Source IP** is listed as **0.0.0.0**, then packets sent using this route use the **Interface's** IP address as the Source IP address. If the **Source IP** lists a specific IP address, then that IP address is used instead.

## Import

**Import** allows you to import a CSV file (Comma Separate Values) into a pair of appliances used at the same site. Before you import, you must remove the header row and save the files on your computer. Complete the following steps to begin your import.

1. Select the appliance you want to upload the routes to.
2. Select **Import** in the **Routes** page.
3. Select **Choose File**.
4. Locate the file you want to import on your desktop.
5. Select **Open**.
6. Select **Import**. Orchestrator will begin generating a CSV file.

The following table is an example of what the CSV file will look like before you import a file.

Subnet	Max Length	Is Local	Adv to SP Peers	Exclude	Adv to BGP Peers	Next Hop	Adv to OSP Neighbors	Interface Name
10.1.0.0	16	TRUE	TRUE	FALSE	10.1.0.1	FALSE	Ian0	wan0

**NOTE** You can limit the file to only the Subnet and Mask Length columns. Orchestrator then uses the default values for the five unlisted columns.

**NOTE** No table cell can be blank when using seven columns.

## VRRP Tab

Configuration > [Networking] 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	Advertisi...	Priority	Preem...	Master IP	Virtual MAC Addr...	State Uptime	Master State Transiti...	IP Add...
	Chicago			No data available										
	Dallas			No data available										
	Denver-EC			Not applicable as this appliance is in Bridge mode.										
	Los-Angeles			Not applicable as this appliance is in Bridge mode.										
	Seattle-EC			Not applicable as this appliance is in Bridge mode.										

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**.

### VRRP Tab Settings

Field Name	Definition
Admin	The options are up (enable) and <b>down</b> (disable).
Advertisement Timer	The default is <b>1 second</b> .
Group ID	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 <b>1 - 255</b> .

Field Name	Definition
<b>Interface</b>	The interface that VRRP is using for peering.
<b>IP Address Owner</b>	A Silver Peak appliance cannot use one of its own IP addresses as the VRRP IP, so this will always be <b>No</b> .
<b>Master IP</b>	Current VRRP Master's Interface or local IP address.
<b>Master State Transitions</b>	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.
<b>Preemption</b>	Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
<b>Priority</b>	The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.
<b>State Uptime</b>	Time elapsed since the VRRP instance entered the state it's in.
<b>State</b>	The VRRP instance has three options: <ul style="list-style-type: none"> <li>■ <b>Backup</b> = Instance is in VRRP backup state.</li> <li>■ <b>Init</b> = Instance is initializing, it's disabled, or the interface is down.</li> <li>■ <b>Master</b> = Instance is the current VRRP master.</li> </ul>
<b>Virtual IP</b>	The IP address of the VRRP instance. VRRP instances may run between two or more appliances, or an appliance and a router.
<b>Virtual MAC address</b>	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 <b>wan0</b> ).

## WCCP Tab

Configuration > [Networking] WCCP

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](#) and the [SD-WAN Deployment Guide](#) for examples, best practices, and deployment tips.

## WCCP Settings

Field Name	Definition
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 <a href="#">Silver Peak Network Deployment Guide</a> .
Compatibility Mode	Select the option appropriate for your router. If a WCCP group is peering with a router running <b>Nexus</b> OS, then the appliance must adjust its WCCP protocol packets to be compatible. By default, the appliance is <b>IOS</b> -compatible.
Forwarding Method	<p>Also known as the <i>Redirect Method</i>. 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 <b>Generic Route Encapsulation (GRE)</b> and <b>L2 redirection</b>.</p> <ul style="list-style-type: none"> <li>■ <b>either</b> allows the appliance and the router to negotiate the best option. You should always select <b>either</b>. During protocol negotiation, if the router offers both GRE and L2 as redirection methods, the appliance will automatically select L2.</li> <li>■ <b>GRE</b> (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.</li> <li>■ <b>L2</b> (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.</li> </ul>
Group ID	Refers to the Service Group ID.
Interface	The default value is <b>wan0</b> .

Field Name	Definition
Oper Status	<p>Common states:</p> <ul style="list-style-type: none"> <li>■ <b>INIT</b>. Initializing or down</li> <li>■ <b>ACTIVE</b>. This indicates that the protocol is established and the router has assigned hash/mask buckets to this appliance.</li> <li>■ <b>BACKUP</b>. 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 <b>0</b>.</li> <li>■ <b>Designated</b>. 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.</li> </ul>
Protocol	Although many more protocols are supported, generally <b>TCP</b> and <b>UDP</b> are the focus. For troubleshooting, you may consider adding a group for <b>ICMP</b> 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

Field Name	Definition
Assignment Detail	<ul style="list-style-type: none"> <li>■ 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 <b>LAN-ingress</b> setting.</li> <li>■ <b>WAN-ingress</b> and <b>LAN-ingress</b> are not applicable if there is only one active appliance.</li> <li>■ <b>WAN-ingress</b> and <b>LAN-ingress</b> are also not applicable if you are using route-map or subnet sharing to tunnelize.</li> <li>■ If there is more than one active appliance and you're using <b>TCP-IP auto-optimization</b>: <ul style="list-style-type: none"> <li>● Use <b>LAN-ingress</b> for WCCP groups that are used to redirect outbound traffic.</li> <li>● Use <b>WAN-ingress</b> for WCCP groups that are used to redirect inbound traffic.</li> </ul> </li> <li>■ This ensures that a connection will go through the same appliance in both inbound and outbound directions and avoid asymmetry.</li> <li>■ <b>custom</b> provides granular control of the distribution of flows. Contact Silver Peak Technical Support for assistance.</li> </ul>

Field Name	Definition
<b>Assignment Method</b>	<p>Determines how redirected packets are distributed between the devices in a Service Group, effectively providing load balancing among the devices. The options are:</p> <ul style="list-style-type: none"> <li>■ <b>either</b>, which lets the appliance and router negotiate the best method for assignment. This is preferred. If the router offers both <b>hash</b> and <b>mask</b> methods, then the appliance will select the <b>mask</b> assignment method.</li> <li>■ <b>hash</b>, for hash table assignment</li> <li>■ <b>mask</b>, for mask/value sets assignment</li> </ul>
<b>Force L2 Return</b>	<p>Generally is 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. <b>Force L2 Return</b> 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.</p>
<b>Password</b>	<p>This field is <i>optional</i>.</p>
<b>Priority</b>	<p>The lowest priority is <b>0</b>, and the default value is <b>128</b>. 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.</p>
<b>Weight</b>	<p>The default value is <b>100</b>. 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 <b>0</b> on the backup appliance.</p>

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

## PPPoE Tab

*Configuration > Networking > PPPoE*

Point-to-Point Protocol over Ethernet (**PPPoE**) is a network protocol for encapsulating PPP frames inside Ethernet frames. It is used mainly with DSL services where individual users connect to a DSL modem over Ethernet.

PPPoE ?

7 Rows Search

Edit	Appliance Name ▼	PPPoE Name	Ethernet Device	Details
	Paris	<i>PPPoE was not configured on the Appliance.</i>		
	Milan	<i>PPPoE was not configured on the Appliance.</i>		
	London	<i>PPPoE was not configured on the Appliance.</i>		
	Geneva	<i>PPPoE was not configured on the Appliance.</i>		
	Frankfurt	<i>PPPoE was not configured on the Appliance.</i>		
	Edinburgh	<i>PPPoE was not configured on the Appliance.</i>		
	Barcelona	<i>PPPoE was not configured on the Appliance.</i>		

When configuring a PPPoE connection, complete the following fields:

Field	Definition
Ethernet Device	Specifies which physical interface to use for sending the protocol. Generally, this is a WAN-side interface.
Password	This is set up with your Internet Service Provider (ISP).
PPPoE Name	The name is <b>ppp</b> , followed by a numerical suffix from <b>0</b> to <b>9</b> .
User Name	This is set up with your Internet Service Provider (ISP).

Generally, this is all the configuration required. If your ISP is fine-tuning the access, you may be asked to configure some of the **Optional Fields**, below.

**Add PPPoE - Paris** ✕

PPPoE Name  ▼

User Name

Password  🔒

Ethernet Device  ▼

**Optional Fields**

UNIT	<input type="text" value="0"/>	Connect Timeout	<input type="text" value="30"/>
LCP Failure	<input type="text" value="3"/>	Connect Poll	<input type="text" value="2"/>
LCP Interval	<input type="text" value="20"/>	Service Name	<input type="text"/>
DNS Type	<input type="text" value="NOCHANGE"/> <span style="margin-left: 20px;">▼</span>	ACNAME	<input type="text"/>
DNS1	<input type="text" value="0.0.0.0"/>	Default Route	<input type="checkbox"/>
DNS2	<input type="text" value="0.0.0.0"/>		

Field	Definition
ACNAME	<i>Access Concentrator Name.</i> Provided by ISP.
Connect Poll	Specifies how many times to try to establish the link. The default value is <b>2</b> .
Connect Timeout	When trying to establish the link, this specifies how many seconds until the effort times out. The default value is <b>30</b> seconds.
Default Route	If the checkbox is selected, the connection uses the default gateway provided by the ISP.

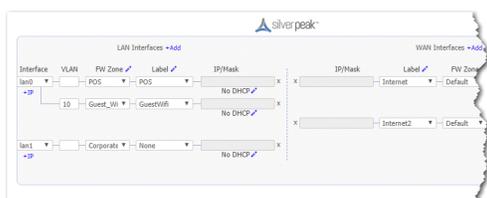
Field	Definition
DNS Type	This specifies which resolver to use: <ul style="list-style-type: none"><li>■ <b>NOCHANGE</b> - Don't accept or configure the ISP's Domain Name Server (DNS). Use the DNS configured in the Orchestrator <b>Administration &gt; [General Settings &gt; Setup] DNS</b> tab.</li><li>■ <b>SERVER</b> - Accept the ISP's DNS. This then overrides Silver Peak's DNS configuration.</li><li>■ <b>SPECIFY</b> - Use <b>DNS1</b> and <b>DNS2</b> to resolve domain names.</li></ul>
LCP Failure	<i>Link Control Protocol Failure</i> . Specifies the number of times the keep-alive can fail before the link goes down. The default value is <b>3</b> .
LCP Interval	The default value for this keep-alive interval is <b>20</b> seconds.
Service Name	Provided by ISP.

## DHCP Server Defaults

*Configuration > Networking > DHCP > DHCP Server Defaults*

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, select **Monitoring**.

## DHCP Settings

### *DHCP Server Fields*

Field Name	Description
<b>Default gateway</b>	When selected, indicates the default gateway is being used.
<b>Default lease, Maximum lease</b>	Specify, in hours, how long an interface can keep a DHCP-assigned IP address.

Field Name	Description
DNS server(s)	Specifies the associated Domain Name System server(s).
Exclude first N addresses	Specifies how many IP addresses are not available at the beginning of the subnet's range.
Exclude last N addresses	Specifies how many IP addresses are not available at the end of the subnet's range.
NetBIOS name server(s)	Used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.
NetBIOS node type	The <b>NetBIOS node type</b> of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types: <ul style="list-style-type: none"> <li>■ <b>B-node</b> = 0x01 Broadcast</li> <li>■ <b>P-node</b> = 0x02 Peer (WINS only)</li> <li>■ <b>M-node</b> = 0x04 Mixed (broadcast, then WINS)</li> <li>■ <b>H-node</b> = 0x08 Hybrid (WINS, then broadcast)</li> </ul>
NTP server(s)	Specifies the associated Network Time Protocol server(s).
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.
Subnet Mask	A mask that specifies the default number of IP addresses reserved for any subnet. For example, entering <b>24</b> reserves 256 IP addresses.
DHCP Failover	

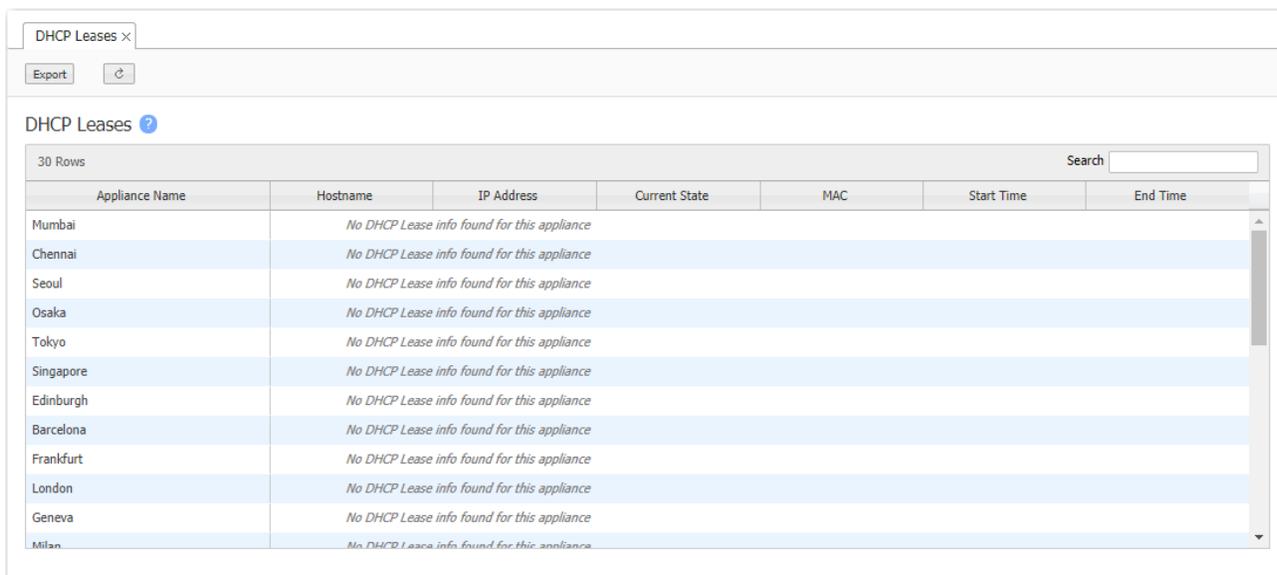
### DHCP/BOOTP Relay Fields

- **DHCP/BOOT:** If enabled, the DHCP settings will only be applied to the configured LAN interface. Enter the specific IP address you want your DHCP settings to be applied in the **Destination DHCP/BOOTP Server** field.
- **Global - All LAN Interfaces on this Appliance:** If enabled and **Enable Option 82** is selected, the DHCP settings will be applied to every appliance. The choices are **append**, **replace**, **forward**, or **discard**.

## DHCP Leases

Configuration > [Networking > DHCP] DHCP Leases

This page lists which IP addresses are currently being leased from the DHCP pool.



DHCP Leases x

Export 

DHCP Leases 

30 Rows Search

Appliance Name	Hostname	IP Address	Current State	MAC	Start Time	End Time
Mumbai	<i>No DHCP Lease info found for this appliance</i>					
Chennai	<i>No DHCP Lease info found for this appliance</i>					
Seoul	<i>No DHCP Lease info found for this appliance</i>					
Osaka	<i>No DHCP Lease info found for this appliance</i>					
Tokyo	<i>No DHCP Lease info found for this appliance</i>					
Singapore	<i>No DHCP Lease info found for this appliance</i>					
Edinburgh	<i>No DHCP Lease info found for this appliance</i>					
Barcelona	<i>No DHCP Lease info found for this appliance</i>					
Frankfurt	<i>No DHCP Lease info found for this appliance</i>					
London	<i>No DHCP Lease info found for this appliance</i>					
Geneva	<i>No DHCP Lease info found for this appliance</i>					
Milan	<i>No DHCP Lease info found for this appliance</i>					

# Tunnels Tab

Configuration > Templates & Networking > Tunnels

Use this page to **view**, **edit**, **add**, or **delete** tunnels. This tab has separate tables for **Overlay**, **Underlay**, and **Passthrough** tunnels.

- If you've deployed an SD-WAN network, then **Business Intent Overlays (BIOs)** govern tunnel creation and properties.
  - Overlay tunnels consist of bonded underlay tunnels.
- If you're not using Overlays, then use the **Tunnels** configuration template to assign tunnel properties. In general, accepting the defaults is sufficient and appropriate.
  - To create tunnels, use **Tunnel Groups**.
  - These tunnels display in the **Underlays** table.

. **Status:** You can also filter by the following statuses: All, Up, or Down.

## Add a Tunnel

Complete the following to add a tunnel to an Overlay or Passthrough Tunnel.

Tunnels ⓘ

Overlay Underlay Passthrough Status All ▾

3482 Rows, 1 Selected Search

Edit	Appliance	Overlay Tunnel	Overlay	Admin Status	Status	MTU	Uptime	Underlay Tunnels	Live Vie...	Historic...
✓	Mumbai	to_Milan_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 19s	to_Milan_MPLS-MPLS, to_Milan_Internet-Inter...	👁	👁
✓	Mumbai	to_Paris_CriticalApps	CriticalApps	up	up - idle	1488	5d 20h 20m 57s	to_Paris_Internet-Internet, to_Paris_MPLS-MP...	👁	👁
✓	Mumbai	to_Seoul_Default	Default	up	up - idle	1488	5d 20h 20m 19s	to_Seoul_MPLS-MPLS, to_Seoul_Internet-Inte...	👁	👁
✓	Mumbai	to_Mexico-City_CriticalApps	CriticalApps	up	up - idle	1488	5d 20h 20m 19s	to_Mexico-City_MPLS-MPLS, to_Mexico-City_I...	👁	👁
✓	Mumbai	to_Toronto_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 19s	to_Toronto_MPLS-MPLS, to_Toronto_Internet...	👁	👁
✓	Mumbai	to_Los-Angeles_CriticalApps	CriticalApps	up	up - idle	1488	5d 20h 19m 54s	to_Los-Angeles_MPLS-MPLS, to_Los-Angeles...	👁	👁
✓	Mumbai	to_Chicago_CriticalApps	CriticalApps	up	up - idle	1488	5d 20h 20m 19s	to_Chicago_Internet-Internet, to_Chicago_MP...	👁	👁
✓	Mumbai	to_Dallas_Default	Default	up	up - idle	1488	5d 20h 20m 19s	to_Dallas_Internet-Internet, to_Dallas_MPLS...	👁	👁
✓	Mumbai	to_Miami_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 39s	to_Miami_Internet-Internet, to_Miami_MPLS...	👁	👁
✓	Mumbai	to_Denver_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 45s	to_Denver_Internet-Internet, to_Denver_MPL...	👁	👁
✓	Mumbai	to_Pittsburgh_Default	Default	up	up - idle	1488	5d 20h 20m 19s	to_Pittsburgh_Internet-Internet, to_Pittsburgh...	👁	👁
✓	Mumbai	to_Minneapolis_CriticalApps	CriticalApps	up	up - idle	1488	5d 20h 20m 19s	to_Minneapolis_Internet-Internet, to_Minneap...	👁	👁
✓	Mumbai	to_Dallas_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 19s	to_Dallas_Internet-Internet, to_Dallas_MPLS...	👁	👁
✓	Mumbai	to_San-Antonio_Default	Default	up	up - idle	1488	5d 20h 20m 30s	to_San-Antonio_MPLS-MPLS, to_San-Antonio...	👁	👁
✓	Mumbai	to_Osaka_Default	Default	up	up - idle	1488	5d 20h 20m 19s	to_Osaka_MPLS-MPLS, to_Osaka_Internet-Int...	👁	👁
✓	Mumbai	to_Chennai_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 22s	to_Chennai_MPLS-MPLS, to_Chennai_Internet...	👁	👁
✓	Mumbai	to_New-Orleans_Default	Default	up	up - idle	1488	5d 20h 20m 19s	to_New-Orleans_MPLS-MPLS, to_New-Orleans...	👁	👁
✓	Mumbai	to_Toronto_Default	Default	up	up - idle	1488	5d 20h 20m 19s	to_Toronto_MPLS-MPLS, to_Toronto_Internet...	👁	👁
✓	Mumbai	to_New-York_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 1s	to_New-York_Internet-Internet, to_New-York...	👁	👁
✓	Mumbai	to_London_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 40s	to_London_MPLS-MPLS, to_London_Internet...	👁	👁
✓	Mumbai	to_New-Orleans_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 19s	to_New-Orleans_MPLS-MPLS, to_New-Orleans...	👁	👁
✓	Mumbai	to_Miami_Default	Default	up	up - idle	1488	5d 20h 20m 31s	to_Miami_Internet-Internet, to_Miami_MPLS...	👁	👁
✓	Mumbai	to_Minneapolis_RealTime	RealTime	up	up - idle	1488	5d 20h 20m 19s	to_Minneapolis_Internet-Internet, to_Minneap...	👁	👁
✓	Mumbai	to_Frankfurt_CriticalApps	CriticalApps	up	up - idle	1488	5d 20h 20m 19s	to_Frankfurt_Internet-Internet, to_Frankfurt...	👁	👁

Field	Definition
Appliance	The name of the selected appliance.
Overlay Tunnel	The designated overlay tunnel.
Overlay	The tunnels are applied to this designated overlay.
Admin Status	Indicates whether the tunnel has been set to admin <b>Up</b> or <b>Down</b> .
Status	<p>The indications are as follows:</p> <p><b>Down</b> 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:</p> <ul style="list-style-type: none"> <li>■ Lack of end-to-end connectivity / routability (test with <i>iperf</i>)</li> <li>■ Intermediate firewall is dropping the packets (open the firewall)</li> <li>■ Intermediate QoS policy (<b>be</b> packets are being starved. Change control packet DSCP marking)</li> <li>■ Mismatched tunnel mode (udp / gre / ipsec / ipsec_ _udp)</li> <li>■ IPsec is misconfigured: (1) enabled on one side (see <i>show int tunnel configured</i>), or mismatched pre-shared key</li> </ul> <p><b>Down - In progress</b> The tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel.</p> <p><b>Down - Misconfigured</b> The two appliances are configured with the same System ID (see <i>show system</i>)</p> <p><b>Up - Active</b> The tunnel is up and active. Traffic destined for this tunnel will be forwarded to the remote appliance.</p> <p><b>Up - Active - Idle</b> 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.</p> <p><b>Up - Reduced Functionality</b> The tunnel is up and active, but the two endpoint appliances are running mismatched software releases that give no performance benefit.</p> <p><b>UNKNOWN</b> The tunnel status is unknown. This can be because the appliance is unable to retrieve the current tunnel status. Try again later.</p>
MTU	<i>Maximum Transmission Unit</i> . The largest possible unit of data that can be sent on a given physical medium. Silver Peak provides support for MTUs up to 9000 bytes. <b>Auto</b> allows the tunnel MTU to be discovered automatically, and it overrides the MTU setting.
Uptime	How long since the tunnel has been up.

Field	Definition
Underlay Tunnels	The designated underlay tunnel.
Live View	A live view of the status of your selected tunnel. You can view by bandwidth, loss, jitter, latency, MOS, chart, traceroute, inbound or outbound, and lock the scale.
Historical Charts	A display of the historical charts for the selected appliance.

## Troubleshooting

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 **Administration > Software Versions**.

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 **WAN Links & Bonding Policy** 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 **Sites** at the top.

## Using Passthrough Tunnels

You would add a passthrough tunnel under the following circumstances:

- For internet breakout to a trusted SaaS application, like Office 365
- For service chaining to a cloud security service, like Zscaler or Symantec
  - This requires building secure, compatible third-party IPsec tunnels from Silver Peak devices to non-Silver Peak devices in the data center or cloud.
  - When you create the tunnel, the **Service Name** in the **Business Intent Overlay's** Internet Traffic **Policies** must exactly match the **Peer/Service** specified in the **Passthrough** tunnel configuration.
  - To load balance, create two or more passthrough IPsec tunnels and, in the Business Intent Overlay, ensure that they all specify the same **Service Name** in the Internet Traffic **Policies**.

## Tunnel Groups Tab

*Configuration > [Templates & Networking > Tunnels] Tunnel Groups*

If you are **not** using Business Intent Overlays (BIOs) to deploy an SD-WAN network, then you would use Tunnel Groups to create the links.

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 is a breadcrumb trail: 'Tunnel Groups > Manage Appliances and Tunnel Groups'. Below this, the 'Tunnel Group Name' is set to 'TG\_1'. The 'Topology' section has 'Mesh' selected, and a 'Connect to Hubs' button is visible. The 'Interfaces' section has 'Only Connect These Labels' selected, with checkboxes for MPLS1, INET1, LTE, INET2, and MPLS2. The 'Cross Connect' section has dropdown menus for each interface, all set to 'None'. 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 **Orchestrator > Tunnels Settings**.
- 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.

Orchestrator builds the topology when you apply a Business Intent Overlay to appliances that have already been assigned a Deployment Profile.

## 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 Orchestrator to assign labels to interfaces. Generally, WAN interfaces are named according to the service or service provider.

## Tunnel Exception

Orchestrator includes a tunnel exception feature that allows you to specify tunnel transactions between overlays. There are two ways you can enable this feature in Orchestrator.

You can configure tunnel exceptions through the *Tunnel Exception* tab.

*Configuration > Tunnels > Tunnel Exception*

1. Select the two appliances that you do not want to connect to via a tunnel.
2. Enter the Interface Labels.



Tunnel Exception ? ⌂ Tunnels

Appliance   Appliance  Description

Interface Label  Interface Label

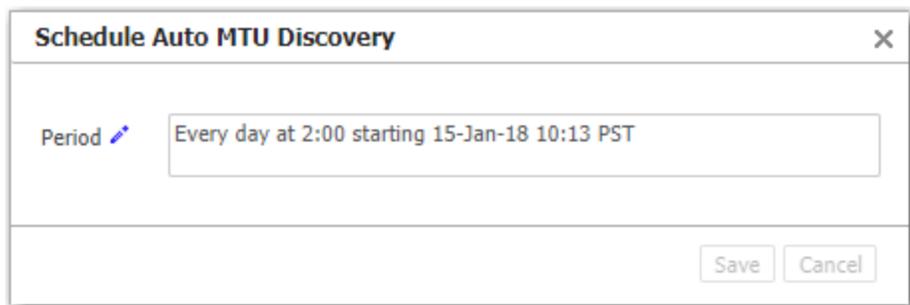
The interface label can be any type of connection, such as **any**, **MPLS**, **Internet**, or **LTE**. Specifying the label excludes appliances within a given network to communicate with that particular appliance.

**Note:** The description field allows you to add a comment if you want to list why you are adding an exception.

## Schedule Auto MTU Discovery

*Configuration > [Networking > Tunnels] Auto MTU Discovery*

Use this screen to schedule when to discover Auto MTU.



The screenshot shows a dialog box titled "Schedule Auto MTU Discovery" with a close button (X) in the top right corner. Inside the dialog, there is a label "Period" with a blue pencil icon to its right. To the right of the label is a text input field containing the text "Every day at 2:00 starting 15-Jan-18 10:13 PST". At the bottom right of the dialog, there are two buttons: "Save" and "Cancel".

## Zscaler Internet Access

Configuration > Third Party Service Orchestration > Zscaler Internet Access

Zscaler Internet Access helps inspect web traffic and enforce security policies defined within Orchestrator.

Zscaler Internet Access ? **Tunnels** ↻

Subscription Tunnel Settings Interface Labels Remote Endpoint Exception

Appliance	Interface Label	VPN Credentials and Location Status	Zscaler ZENS
<i>No Data Available</i>			

Field	Definition
<b>Appliance</b>	The name of the appliance you want to connect with Zscaler.
<b>Interface Label</b>	The name of the interfaces you want to connect with Zscaler.
<b>VPN Credentials and Location Status</b>	The VPN credentials and location status of your subscription with Zscaler.
<b>Zscaler ZENS</b>	Zscaler Enforcement Nodes: the Zscaler endpoints where the tunnels connect.

Before you begin Zscaler configuration, you need to create a Zscaler account and ensure you have an established connection with Zscaler.

### Subscription

1. Go to <https://help.zscaler.com/zia/sd-wan-api-integration>.
2. Once you have completed the steps in the above URL to configure your Zscaler account, navigate to the **Zscaler Internet Access** tab in Orchestrator.
3. Select the **Subscription** tab to get started with Zscaler.
4. Enter the information in the Subscription fields that reflect your Zscaler account.
5. Select **Save** once you have completed entering the information in the table below. The Zscaler field should reflect **Connected**.

The following table represents the values in the **Subscription** window.

Field	Definition
<b>Zscaler</b>	This field displays if you are connected or not connected to your Zscaler account.

Field	Definition
Zscaler Cloud	The Zscaler cloud URL. Ex: admin.zscalerthree.net
Partner Username	The partner administrator user name you created when configuring Zscaler.
Partner Password	The partner administrator password you created when configuring Zscaler.
Partner Key	The partner key you created when configuring your Zscaler account. Select Silver Peak from the list of partners.
Domain	The domain provisioned in Zscaler for your enterprise.

## Tunnel Settings

The **Tunnel Settings** tab helps you define the tunnels associated with Zscaler and Silver Peak EdgeConnect.

**Note:** You can configure General, IKE, and IPsec tunnel settings. The settings are automatically generated; however, you can edit if you want to do so.

### General

Field	Definition
Mode	The tunneling protocol being used.
Auto Max BW Enabled	Check this box.
NAT	The network address translation between the Silver Peak and Zscaler networks. Select <b>None</b> .

### IKE

Field	Definition
Preshared Key	The same pre-shared key you entered when creating the VPN credential.
Authentication Algorithm	Select <b>SHA1</b> .
Encryption Algorithm	Select <b>AES128</b> .
Diffie-Hellman Group	Select <b>2</b> .
Lifetime	Enter <b>1440</b> minutes.
Dead Peer Detection	
Delay Time	Enter <b>300</b> seconds.
Retry Count	Enter <b>3</b> .
IKE Identifier	Select <b>User_FQDN</b> .
Phase 1 Mode	Select <b>Aggressive</b> .

### IPSec

Field	Definition
Authentication Algorithm	Select <b>SHA1</b> .
Encryption Algorithm	Select <b>NULL</b> .
Enable IPsec Anti-replay Window	Check this box.
Lifetime	Enter <b>480</b> minutes. Enter <b>0</b> megabytes.
Perfect Forward Secrecy Group	Select <b>disable</b> .

### Interface Labels

1. Select the **Interface Labels** tab. The **Build Tunnels Using These Interfaces** displays.
2. Drag the Interface labels you want to use into the Preferred Interface Label Order column.
3. Select **Save**.

**Build Tunnels Using These Interfaces** ×

Preferred Interface Label Order		Interface Label
INET2		MPLS1
INET1		LTE
		MPLS2

Save
Close

### Remote Endpoint Exception

You can use the **Remote Endpoint Exception** if you want to override the automatically selected ZEN pair for specific sites. You have the option to add this exception to one or more sites within your network.

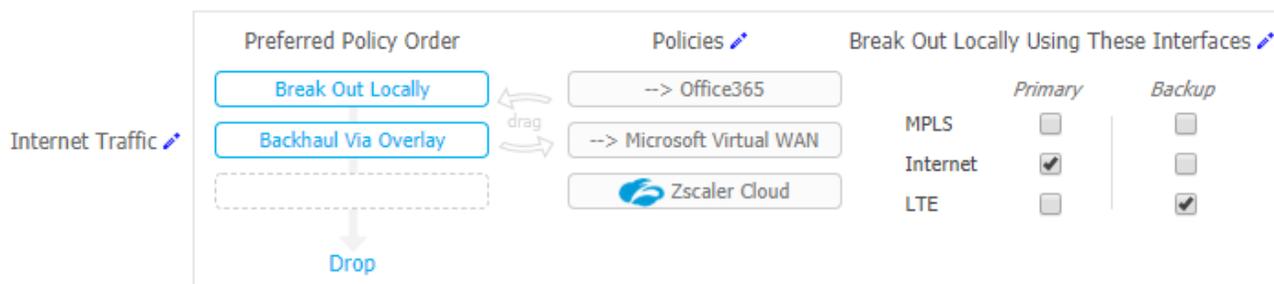
1. Select the **Remote Endpoint Exception** tab.
2. Enter the appliance name, the interface label, and the Primary and Secondary IP addresses. Orchestrator will build tunnels to those ZENS.

Field	Definition
Appliance	The appliance for which we override Zscaler ZENS.
Interface Label	The interface label from where tunnels are built.
Primary IP	The IP address of the primary Zscaler ZEN.
Secondary IP	The IP address of the secondary Zscaler ZEN.

## Enabling Zscaler

Lastly, you need to enable the Zscaler service.

1. Go to the **Business Intent Overlay** tab in Orchestrator.
2. Select the overlay that breaks out traffic to Zscaler.
3. Drag **Zscaler Cloud** from the **Policies** column to the **Preferred Policy Order** column.



## Verification

You can first verify Zscaler has been deployed in the **BIO** (Business Intent Overlay) tab. Once the Zscaler Internet Access is configured and the Zscaler policy is applied successfully in the BIO, deployment will begin automatically. Go to the Zscaler Internet Access tab to verify deployment was successful.

Zscaler Internet Access Tunnels

Subscription Tunnel Settings Interface Labels Remote Endpoint Exception

4 Rows				Search
Appliance	Interface Label	VPN Credentials and Location Status	Zscaler ZENs	
ECV-A	INET2	Deployed	Discovered: 165.225.72.39, 199.168.148.132	
ECV-B	INET2	Deployed	Discovered: 165.225.72.39, 199.168.148.132	
ECV-C	INET1	Deployed	Discovered: 165.225.72.39, 199.168.148.132	
ECV-D	INET1	Deployed	Discovered: 165.225.72.39, 199.168.148.132	

You can also verify your Zscaler tunnels have been successfully deployed in the **Tunnels** tab. Zscaler tunnels should be listed in the **Passthrough Tunnel** column with a green status of **up - active**.

## Tunnels ?

Overlay Underlay **Passthrough** Status All ▼

138/800 Rows											Search	third
Edit	Appliance	Passthrough Tunnel	Admin Status	Status	Local IP	Remote IP	Mode	NAT	Peer/Service	Max BW Kbps	Advanced Options	
	EAST1-Virginia-AWS	ThirdParty_Zscaler_INETB...	up	up - active	10.50.22.199	165.225.48.10	IPSec	none	Zscaler_INETB_Primary	1000000(Auto)		
	EAST1-Virginia-AWS	ThirdParty_Zscaler_INETA...	up	up - active	10.50.20.134	165.225.48.10	IPSec	none	Zscaler_INETA_Primary	1000000(Auto)		
	EAST1-Virginia-AWS	ThirdParty_Zscaler_INETB...	up	up - active	10.50.22.199	165.225.0.165	IPSec	none	Zscaler_INETB_Backup	1000000(Auto)		
	EAST1-Virginia-AWS	ThirdParty_Zscaler_INETA...	up	up - active	10.50.20.134	165.225.0.165	IPSec	none	Zscaler_INETA_Backup	1000000(Auto)		
	EMEA2-Paris-Devaux...	ThirdParty_Zscaler_INETA...	up	up - active	51.15.159.48	165.225.76.42	IPSec	none	Zscaler_INETA_Primary	10000(Auto)		
	EMEA2-Paris-Devaux...	ThirdParty_Zscaler_INETA...	up	up - active	51.15.159.48	165.225.88.39	IPSec	none	Zscaler_INETA_Backup	10000(Auto)		
	EMEA1-Amsterdam-A...	ThirdParty_Zscaler_INETA...	up	up - active	172.29.100.4	165.225.28.14	IPSec	none	Zscaler_INETA_Primary	1000000(Auto)		
	EMEA1-Amsterdam-A...	ThirdParty_Zscaler_INETA...	up	up - active	172.29.100.4	165.225.88.39	IPSec	none	Zscaler_INETA_Backup	1000000(Auto)		
	EAST2-Virginia-AWS	ThirdParty_Zscaler_INETB...	up	up - active	10.50.20.124	165.225.0.165	IPSec	none	Zscaler_INETB_Backup	1000000(Auto)		
	EAST2-Virginia-AWS	ThirdParty_Zscaler_INETA...	up	up - active	10.50.22.198	165.225.48.10	IPSec	none	Zscaler_INETA_Primary	1000000(Auto)		
	EAST2-Virginia-AWS	ThirdParty_Zscaler_INETA...	up	up - active	10.50.22.198	165.225.0.165	IPSec	none	Zscaler_INETA_Backup	1000000(Auto)		
	EAST2-Virginia-AWS	ThirdParty_Zscaler_INETB...	up	up - active	10.50.20.124	165.225.48.10	IPSec	none	Zscaler_INETB_Primary	1000000(Auto)		
	APJ1-Singapore-Azure	ThirdParty_Zscaler_INETA...	up	up - active	10.5.0.4	165.225.116.24	IPSec	none	Zscaler_INETA_Backup	2000000(Auto)		
	APJ1-Singapore-Azure	ThirdParty_Zscaler_INETB...	up	up - active	10.5.1.4	165.225.112.24	IPSec	none	Zscaler_INETB_Primary	2000000(Auto)		
	APJ1-Singapore-Azure	ThirdParty_Zscaler_INETB...	up	up - active	10.5.1.4	165.225.116.24	IPSec	none	Zscaler_INETB_Backup	2000000(Auto)		
	APJ1-Singapore-Azure	ThirdParty_Zscaler_INETA...	up	up - active	10.5.0.4	165.225.112.24	IPSec	none	Zscaler_INETA_Primary	2000000(Auto)		
	EAST3-Charleston-Go...	ThirdParty_Zscaler_INETA...	up	up - active	10.180.1.3	104.129.206.161	IPSec	none	Zscaler_INETA_Backup	1000000(Auto)		
	EAST3-Charleston-Go...	ThirdParty_Zscaler_INETA...	up	up - active	10.180.1.3	165.225.48.10	IPSec	none	Zscaler_INETA_Primary	1000000(Auto)		
	APJ2-Singapore-Azure	ThirdParty_Zscaler_INETB...	up	up - active	10.5.1.5	165.225.112.24	IPSec	none	Zscaler_INETB_Primary	2000000(Auto)		
	APJ2-Singapore-Azure	ThirdParty_Zscaler_INETA...	up	up - active	10.5.0.5	165.225.116.24	IPSec	none	Zscaler_INETA_Backup	2000000(Auto)		

Please note the following:

- Zscaler is applied to all your EdgeConnect appliance's associated overlays that have the Zscaler policy enabled.
- Only IPsec mode is supported for Zscaler and the bandwidth limit is 200Mbps per EdgeConnect appliance.

# NAT

*Configuration > Networking > NAT*

NAT allows for multiple sites with overlapping IP addresses to connect to a single SD-WAN fabric. You can configure S-NAT (Source Network Address Translation), D-NAT (Destination Network Address Translation), destination TCP, and UDP port translation rules to LAN to SD-WAN fabric traffic in the ingress and egress direction. The following address translation options are supported:

- 1:1 source and destination IP address translation
- 1:1 subnet to subnet source and destination IP address translation
- Many to one IP source address translation
- NAT pools for translated source IP address

You can view both NAT Rules and NAT Pools within your network by selecting **NAT Rule** or **NAT Pools** at the top of the page. You can also export a CSV file of your branch NAT traffic. Select the **Edit** icon to add rules to your NAT and NAT Pools.

## NAT Rules and Pools

You can add NAT rules by completing all the values in the table shown below. Each NAT rule has a directional field or value. Outbound rules are applied to the traffic flows initiated from the LAN, destined to the SD-WAN fabric. Inbound rules are applied to the traffic flows initiated from the SD-WAN fabric destined to the LAN. Return traffic for a given flow does not require an additional rule. The destination IP address must be configured for each rule.

**NOTE** You must disable advertisements of local, static routes on the LAN side at the site so the routes are completely unique. Additionally, you must configure static routes for NAT pools and advertise them to the SD-WAN fabric by enabling **Advertise to Silver Peak Peers**.

Complete the following steps to add a rule to your NAT:

1. Select **Add Rule**.
2. Complete the following values in the table by selecting any of the columns.

Field	Definition
<b>Priority</b>	The order in which the rules are executed: the lower the priority, the higher the chance your NAT rule will be applied.
<b>LAN Interface</b>	The name of the LAN interface the NAT rule is using. This is configurable for an outbound NAT rule only.
<b>Direction</b>	Select the direction the traffic is going: <ul style="list-style-type: none"> <li>■ Outbound (LAN to Fabric)</li> <li>■ Inbound (Fabric to LAN)</li> </ul>
<b>Protocol</b>	The type of protocol being used for each NAT.
<b>Source</b>	The original source IP address of the IP packet.
<b>Destination</b>	The address of the LAN/WAN interface where the traffic is going to.
<b>Translated Source</b>	The translated source IP address when the NAT rule is applied.
<b>Translated Destination</b>	The translated destination IP address when the NAT rule is applied.
<b>Enabled</b>	Check this box to enable your customized NAT rule. Direction can be both inbound or outbound.
<b>Comment</b>	Any comment you want to add pertaining to your NAT rule.
<b>Criteria</b>	<b>Match:</b> LAN interface, direction, source, destination <b>Set:</b> Translated source, translated destination

### NAT Pools

You also have the option to configure a NAT pool. Complete the following steps to create a NAT pool:

1. Select the **Edit** icon in the NAT tab. The **NAT** window opens.
2. Select the **NAT Pools** icon. The **NAT Pools** window opens.
3. Select **Add**.
4. Select the columns in the table, starting with **Name**, to enter information regarding your Pool.

Field	Definition
<b>Name</b>	The name of your pool.
<b>Direction</b>	Whether the traffic is outbound or inbound.
<b>Subnet</b>	The IP address of the subnet.
<b>Translate Ports</b>	Enable source port address translation if the NAT pool is too small to accommodate multiple, flows simultaneously with 1:1 IP address translation.

# Policy Configuration Tabs

These topics describe the pages related to managing access lists and policies.

## DNS Proxy Policies

*Configuration > Policies > DNS Proxy Policies*

The DNS (Domain Name Server) Proxy stores public IP addresses with their associated domain name. Server A is primarily used as a private DNS to backhaul traffic and Server B is used to match all other domains that are not included under Server A. Server B is also used for public (cloud services) to breakout traffic. See the table below for the field descriptions in this tab.

Field	Definition
Appliance Name	The name of the appliance associated with DNS proxy.
DNS Proxy Enabled	Whether the DNS Proxy is enabled. Select <b>True</b> or <b>False</b> .
Interface	The name of the interface associated with the DNS proxy.
Server A Addresses	The IP addresses of Server A.
Server A Domains	The domain addresses of Server A.
Server A Caching	Whether you configured the server to be cached.
Server B Addresses	The IP addresses of Server B.
Server B Domains	The domain addresses of Server B.
Server B Caching	Whether you configured the server to be cached.

### DNS Proxy Policies

Complete the following steps to configure and define your DNS Proxy policies.

**NOTE** This feature is only configurable if you have loopback interfaces configured.

1. Choose if you want to enable the DNS Proxy by selecting **ON** or **OFF**.
2. Select the name of the loopback interface or the LAN-side label associated with your DNS proxy.
3. Enter the IP addresses for Server A in the **Server A Addresses** field.
4. Choose if you want Caching to be **ON** or **OFF**. If selected, the domain name to the IP address mapping is cached. By default, caching is **ON**.
5. Enter the domain names of the Server A for the above IP addresses.
6. Enter Server B IP addresses in the **Server B Addresses** field. Server B will be used if there are no matches to the Server A domains.

**NOTE** You can **Clear DNS Cache**. This will erase the domain name to the IP address mapping you had cached for both Server A and B.

## Route Policies Tab

*Configuration > [Policies] 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 Appliance Manager or CLI), and entries that result from applying an Orchestrator Route Policies template, or applying Business Intent Overlays (if you're deploying an SD-WAN).

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 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 select the **Edit** icon to manage Routing policies directly for a particular appliance.

If you're deploying an SD-WAN network and setting up Internet breakout from the branch, you must create manual route policy entries for sanctioned SaaS applications or Guest WiFi.

## 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.

## Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.
- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## 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**.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## QoS Policies Tab

*Configuration > [Policies] 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 Appliance Manager or CLI), and entries that result from applying an Orchestrator QoS Policy template or Business Intent Overlay.

Use the **Shaper** to define, prioritize, and name traffic classes. Think of it as, 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
					Traffic Class	LAN QoS	WAN QoS	
	Albuquerque	map1 (active)	20000	ACL RealTime	3 - RealTime	trust-lan	be	RealTime overlay
	Albuquerque	map1 (active)	20001	ACL Interactive	2 - Interactive	trust-lan	be	Interactive overlay
	Albuquerque	map1 (active)	20002	ACL AnyTraffic	1 - default	trust-lan	be	DefaultOverlay overlay
	Albuquerque	map1 (active)	65535	Protocol ip	1 - default	trust-lan	trust-lan	
	Barcelona	map1 (active)	20000	ACL RealTime	3 - RealTime	trust-lan	be	RealTime overlay
	Barcelona	map1 (active)	20001	ACL Interactive	2 - Interactive	trust-lan	be	Interactive overlay
	Barcelona	map1 (active)	20002	ACL AnyTraffic	1 - default	trust-lan	be	DefaultOverlay overlay
	Barcelona	map1 (active)	65535	Protocol ip	1 - default	trust-lan	trust-lan	
	Boston	map1 (active)	20000	ACL RealTime	3 - RealTime	trust-lan	be	RealTime overlay
	Boston	map1 (active)	20001	ACL Interactive	2 - Interactive	trust-lan	be	Interactive overlay
	Boston	map1 (active)	20002	ACL AnyTraffic	1 - default	trust-lan	be	DefaultOverlay overlay
	Boston	map1 (active)	65535	Protocol ip	1 - default	trust-lan	trust-lan	
	Chennai	map1 (active)	20000	ACL RealTime	3 - RealTime	trust-lan	be	RealTime overlay
	Chennai	map1 (active)	20001	ACL Interactive	2 - Interactive	trust-lan	be	Interactive overlay
	Chennai	map1 (active)	20002	ACL AnyTraffic	1 - default	trust-lan	be	DefaultOverlay overlay
	Chennai	map1 (active)	65535	Protocol ip	1 - default	trust-lan	trust-lan	
	Chicago	map1 (active)	20000	ACL RealTime	3 - RealTime	trust-lan	be	RealTime overlay

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

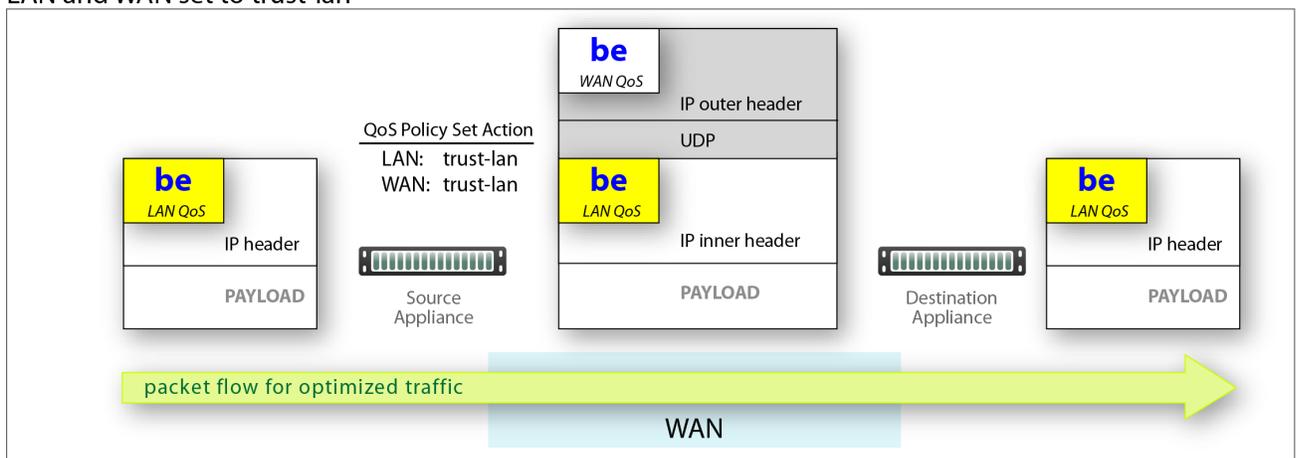
## Handling and Marking DSCP Packets

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

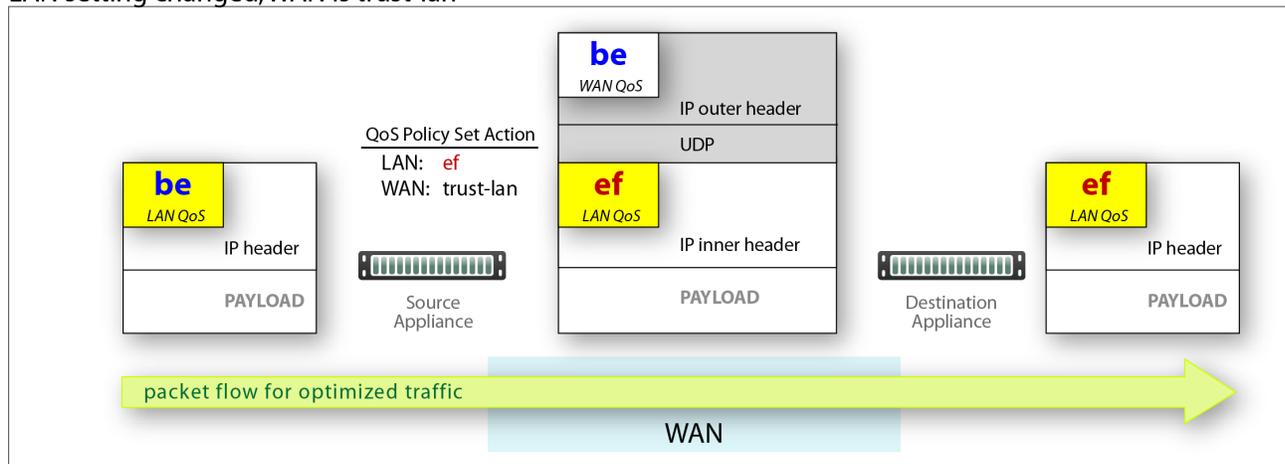
### Applying DSCP Markings to Optimized (Tunnelized) Traffic

- 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.

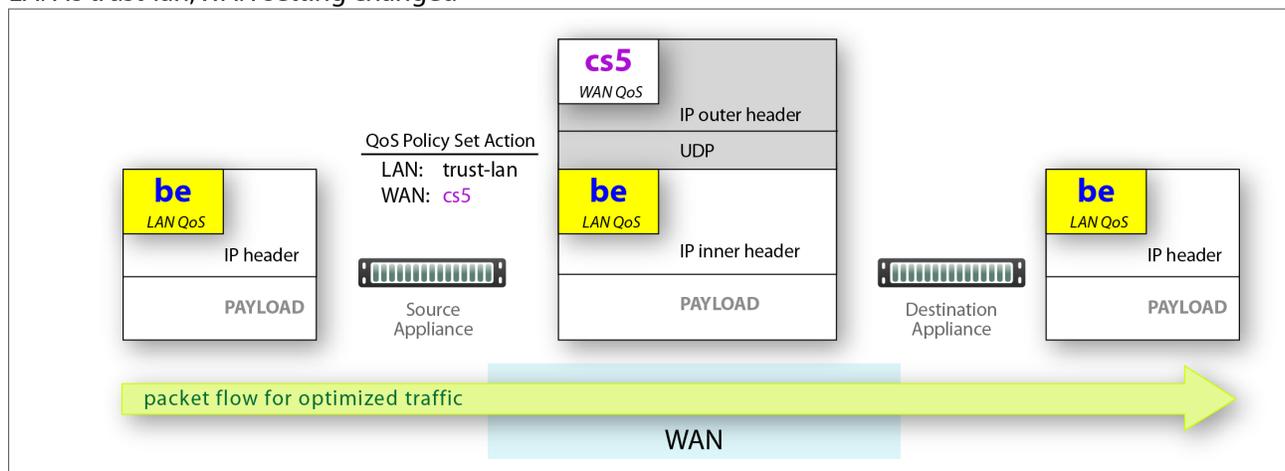
#### LAN and WAN set to trust-lan



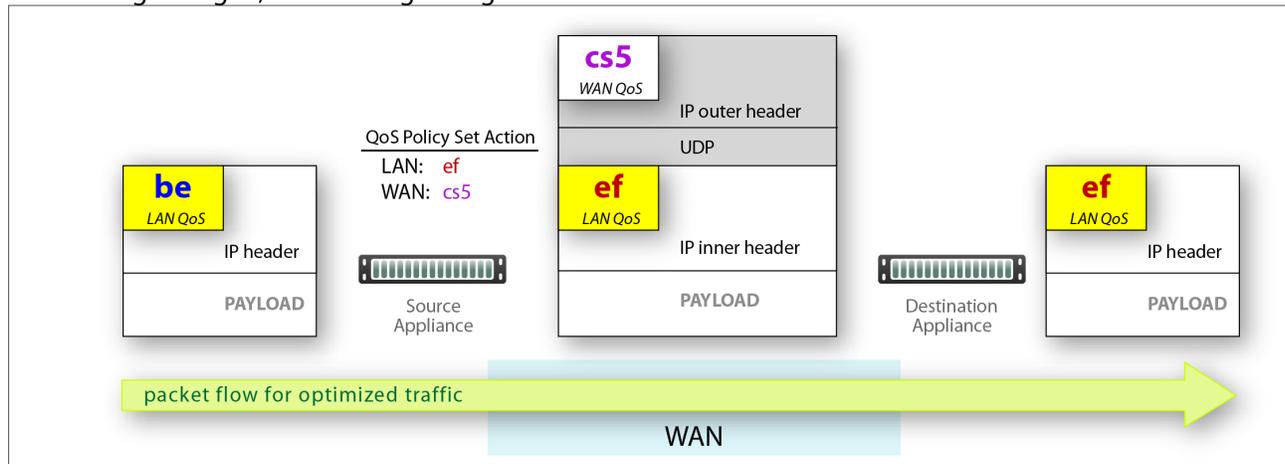
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



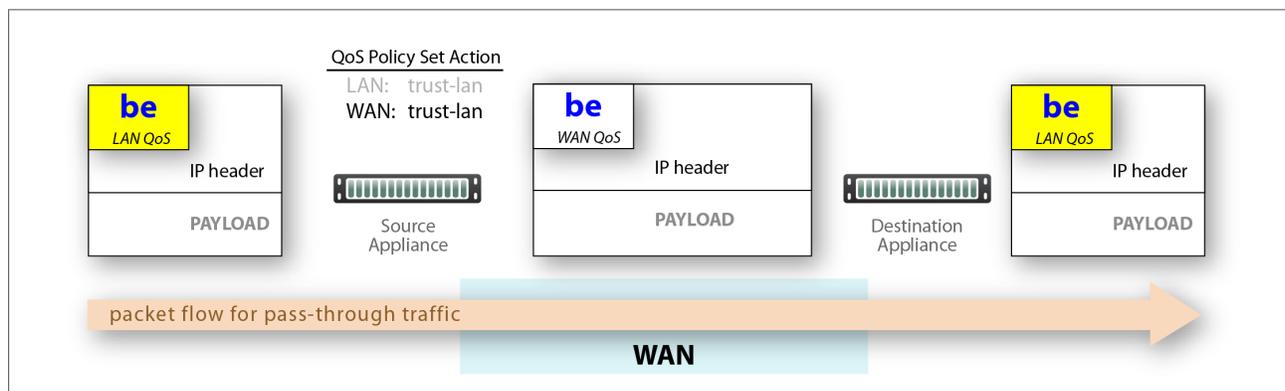
LAN setting changed, WAN setting changed



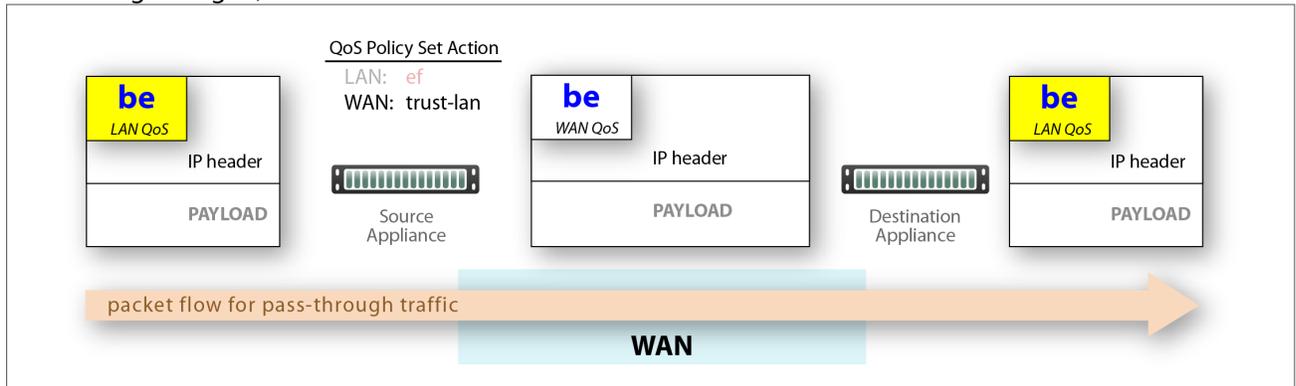
Applying DSCP Markings to Pass-through Traffic

- 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 marking.
  - When the packet reaches the remote appliance, it retains the modified QoS setting as it travels to its destination.

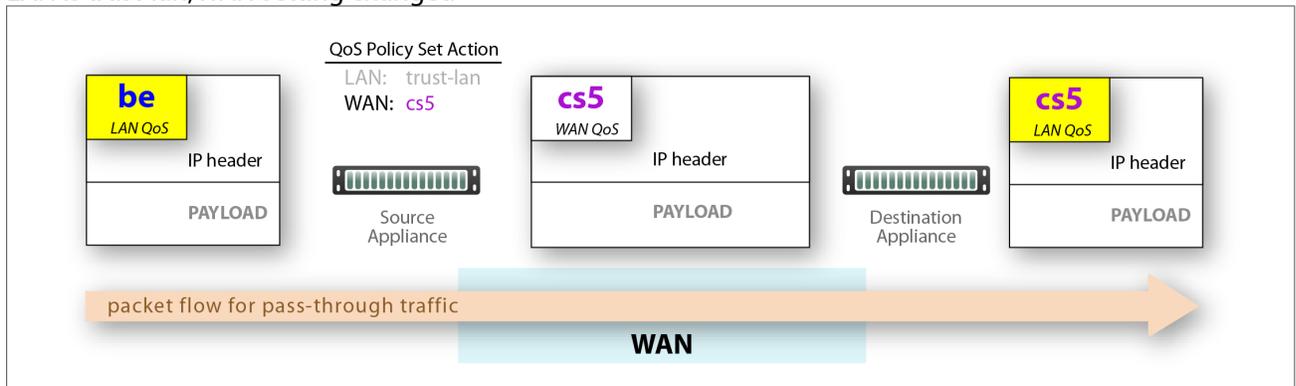
LAN and WAN set to trust-lan



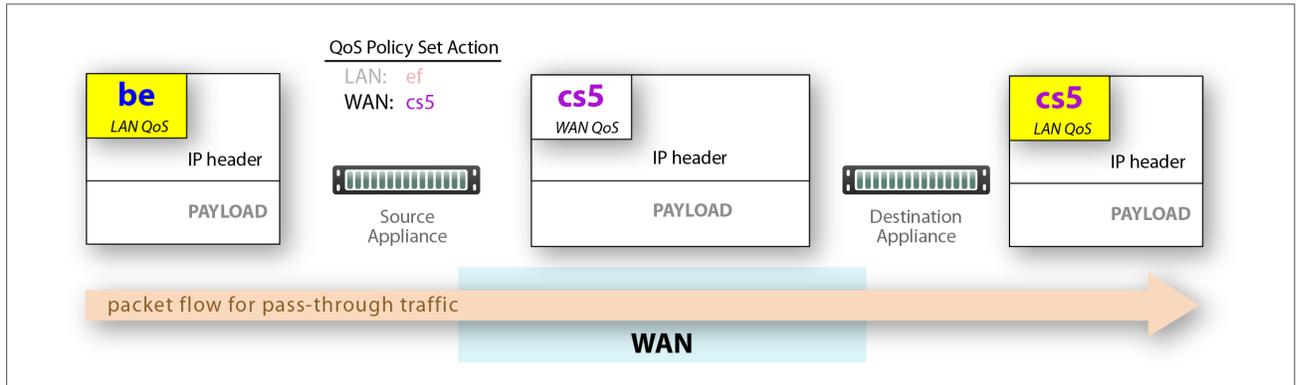
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



## LAN setting changed, WAN setting changed



## 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.

## Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.

- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## 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.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## Schedule QoS Map Activation

*Configuration > [Policies] 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

California

Map ▲	Schedule	Re-Classify	Description
map1	Every day at 6:00 starting 13-Jan-17 20:48 GMT <span style="float: right;"><a href="#">Edit</a></span>	<input checked="" type="checkbox"/>	primary map
map2	Every day at 20:00 starting 13-Jan-17 20:48 GMT <span style="float: right;"><a href="#">Edit</a></span>	<input checked="" type="checkbox"/>	evening map

Before using this option, verify the following:

- The desired Template Group has the QoS maps you need.
- 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 > [Software & Setup > Setup] Timezone for Scheduled Jobs**.

## Optimization Policies Tab

*Configuration > [Policies] Optimization Policies*

The **Optimization Policies** tab displays the Optimization policy entries that exist on the appliances. This includes the appliance-based defaults, entries applied manually (via the Appliance Manager or CLI), and entries that result from applying an Orchestrator Optimization Policy template or Business Intent Overlay.

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

Optimization Policies											
Manage Optimization Policies with Templates											
<span>Export</span> <input checked="" type="checkbox"/> Display active policies											
Optimization Policies											
49 Rows											
Search											
Edit	Appliance N...	Map	Priority	Match Criteria	Set Actions						Comment
					Network M...	IP Header ...	Payload Co...	TCP Accel	TCP Accel D...	Protocol Ac...	
✓	Chennai	map1 (active)	10000	Protocol tcp, Destination Port 139	balanced	Yes	Yes	Yes		cifs	
✓	Chennai	map1 (active)	10010	Protocol tcp, Destination Port 445	balanced	Yes	Yes	Yes		cifs	
✓	Chennai	map1 (active)	10020	Protocol tcp, Destination Port 443	balanced	Yes	Yes	Yes		ssl	
✓	Chennai	map1 (active)	10021	Protocol tcp, Source Port 443	balanced	Yes	Yes	Yes		ssl	
✓	Chennai	map1 (active)	10030	Protocol tcp, Destination Port 2598	balanced	Yes	Yes	Yes		citrix	
✓	Chennai	map1 (active)	10040	Protocol tcp, Destination Port 1494	balanced	Yes	Yes	Yes		citrix	
✓	Chennai	map1 (active)	10050	Protocol tcp, Destination Port 860	balanced	Yes	Yes	Yes		iscsi	
✓	Chennai	map1 (active)	10060	Protocol tcp, Destination Port 3260	balanced	Yes	Yes	Yes		iscsi	
✓	Chennai	map1 (active)	10070	Protocol tcp, Destination Port 9100	balanced	Yes	Yes	Yes		none	
✓	Chennai	map1 (active)	10071	Protocol tcp, Source Port 9100	balanced	Yes	Yes	Yes		none	
✓	Chennai	map1 (active)	65535	Match Everything	balanced	Yes	Yes	Yes		none	
✓	Chicago	map1 (active)	65535	Match Everything	balanced	Yes	Yes	Yes		none	
✓	London	map1 (active)	10000	Protocol tcp, Destination Port 139	balanced	Yes	Yes	Yes		cifs	
✓	London	map1 (active)	10010	Protocol tcp, Destination Port 445	balanced	Yes	Yes	Yes		cifs	
✓	London	map1 (active)	10020	Protocol tcp, Destination Port 443	balanced	Yes	Yes	Yes		ssl	
✓	London	map1 (active)	10021	Protocol tcp, Source Port 443	balanced	Yes	Yes	Yes		ssl	
✓	London	map1 (active)	10030	Protocol tcp, Destination Port 2598	balanced	Yes	Yes	Yes		citrix	
✓	London	map1 (active)	10040	Protocol tcp, Destination Port 1494	balanced	Yes	Yes	Yes		citrix	

### 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.

## Match Criteria

- These are universal across all policy maps – **Route, QoS, Optimization, NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.
- The available parameters are **Application, Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain, Geo Location, Interface, Protocol, DSCP, IP/Subnet, Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## 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.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).

- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## Set Actions

Set Action	Definition
<b>Network Memory</b>	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.
	<p><b>Maximize Reduction</b>      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.</p>
	<p><b>Minimize Latency</b>      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.</p>
	<p><b>Balanced</b>      Is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types.</p>
	<p><b>Disabled</b>      Turns off Network Memory.</p>
<b>IP Header Compression</b>	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.

Set Action	Definition
<b>Payload Compression</b>	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.
<b>TCP Acceleration</b>	Uses techniques such as selective acknowledgements, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links. For more information, see <a href="#">TCP Acceleration Options</a> .
<b>Protocol Acceleration</b>	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 <b>client</b> ) 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.

### TCP Accel Options ✕

**IMPORTANT:** Changing these settings can affect service. Consult the documentation before editing default values.

Adjust MSS to Tunnel MTU	<input checked="" type="checkbox"/>	
Preserve Packet Boundaries	<input checked="" type="checkbox"/>	
Enable Silver Peak TCP SYN option exchange	<input checked="" type="checkbox"/>	
Route Policy Override	<input checked="" type="checkbox"/>	
Auto Reset Flows	<input type="checkbox"/>	
IP Black Listing	<input type="checkbox"/>	
End to End FIN handling	<input checked="" type="checkbox"/>	
WAN Window Scale	<input type="text" value="8"/>	(1..14)
Slow LAN Defense	<input type="text" value="9"/>	(0..12, 0=off)
WAN Congestion Control	<input type="text" value="optimized"/>	
<b>Per-Flow Buffer</b>		
Max LAN to WAN Buffer	<input type="text" value="64000"/>	KB (64..1000000)
Max WAN to LAN Buffer	<input type="text" value="64000"/>	KB (64..1000000)
Slow LAN Window Penalty	<input type="text" value="0"/>	(0..254, 0=off)
LAN Side Window Scale Factor Clamp	<input type="text" value="0"/>	(0..14, 0=off)
Persist timer Timeout	<input type="text" value="0"/>	Sec (0..64000, 0=off)
<b>Keep Alive Timer</b>		
Probe Interval	<input type="text" value="30"/>	Sec (1..64000)
Probe Count	<input type="text" value="8"/>	(1..254)
First Timeout (Idle)	<input type="text" value="600"/>	Sec (1..64000)



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

Option	Description
Adjust MSS to Tunnel MTU	<p>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 <math>TCP\ MSS = Tunnel\ MTU - Tunnel\ Packet\ Overhead</math>.</p> <p>This feature is enabled by default so that the <b>maximum value</b> 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.</p>
Auto Reset Flows	<p><b>NOTE:</b> 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). The internal criteria can also include:</p> <ul style="list-style-type: none"> <li>■ Resetting all TCP accelerated flows on a Tunnel Down event.</li> <li>■ Resetting <ul style="list-style-type: none"> <li>● TCP acceleration is enabled</li> <li>● 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).</li> </ul> </li> </ul>
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>
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 <b>ON</b> (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 <b>OFF</b> (Default TCP), no such synchronization happens and the two LAN segments at the ends gracefully shutdown, independently.</p>
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>
Keep Alive Timer	<p>Allows us to change the Keep Alive timer for the TCP connections.</p> <ul style="list-style-type: none"> <li>■ <b>Probe Interval</b> - Time interval in seconds between two consecutive Keep Alive Probes</li> <li>■ <b>Probe Count</b> - Maximum number of Keep Alive probes to send</li> <li>■ <b>First Timeout (Idle)</b> - Time interval until the first Keep Alive timeout</li> </ul>

Option	Description
<b>LAN Side Window Scale Factor Clamp</b>	This setting allows the appliance to present an artificially lowered Window Scale Factor (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.
<b>Per-Flow Buffer</b>	<b>(Max LAN to WAN Buffer and Max WAN to LAN Buffer)</b> This setting clamps the maximum buffer space that can be allocated to a flow, in each direction.
<b>Persist timer Timeout</b>	Allows the TCP to terminate connections that are in Persist timeout stage after the configured number of seconds.
<b>Preserve Packet Boundaries</b>	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.
<b>Route Policy Override</b>	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. 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.
<b>Slow LAN Defense</b>	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. This feature is enabled by default. The number relates indirectly to the amount of time the system waits before resetting such slow flows.
<b>Slow LAN Window Penalty</b>	This setting ( <b>OFF</b> 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.
<b>WAN Congestion Control</b>	Selects the internal Congestion Control parameter: <ul style="list-style-type: none"> <li>■ <b>Optimized</b> - This is the default setting. This mode offers optimized performance in almost all scenarios.</li> <li>■ <b>Standard</b> - In some unique cases it may be necessary to downgrade to Standard performance to better interoperate with other flows on the WAN link.</li> <li>■ <b>Aggressive</b> - Provides aggressive performance and should be used with caution. Recommended mostly for Data Replication scenarios.</li> </ul>
<b>WAN Window Scale</b>	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.

## NAT Policies Tab

Configuration > Policies > NAT Policies

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

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

NAT Policies		NAT All Inbound			NAT All Outbound		
Edit	Appliance Name	Enable	NAT IP	Fallback	Enable	NAT IP	Fallback
✓	Chennai	No	auto	No	No	auto	No
✓	Chicago	No	auto	No	No	auto	No
✓	London	No	auto	No	No	auto	No
✓	Los-Angeles	No	auto	No	No	auto	No
✓	Miami	No	auto	No	No	auto	No
✓	Minneapolis	No	auto	No	No	auto	No
✓	Mumbai	No	auto	No	No	auto	No
✓	Munich	No	auto	No	No	auto	No
✓	Portland	No	auto	No	No	auto	No

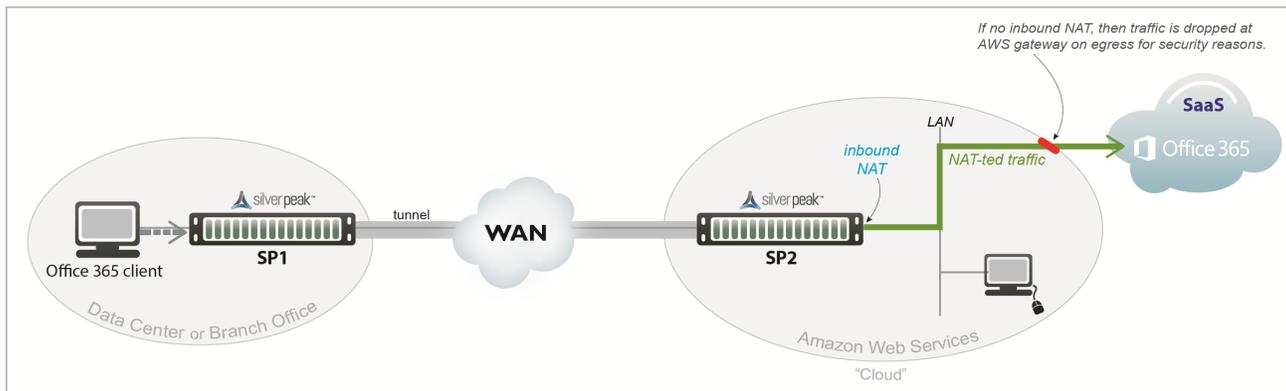
2. The **Advanced** view displays all the NAT map rules.

Edit	Appliance N...	Map	Priority	Match Criteria	Set Actions				Comment
					NAT Type	NAT Direct...	NAT IP	Fallback	
✓	Chennai	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Chicago	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	London	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Los-Angeles	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Miami	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Minneapolis	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Mumbai	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Munich	map1 (active)	65535	Match Everything	no-nat	none	auto	No	
✓	Portland	map1 (active)	65535	Match Everything	no-nat	none	auto	No	

Two use cases illustrate the need for NAT:

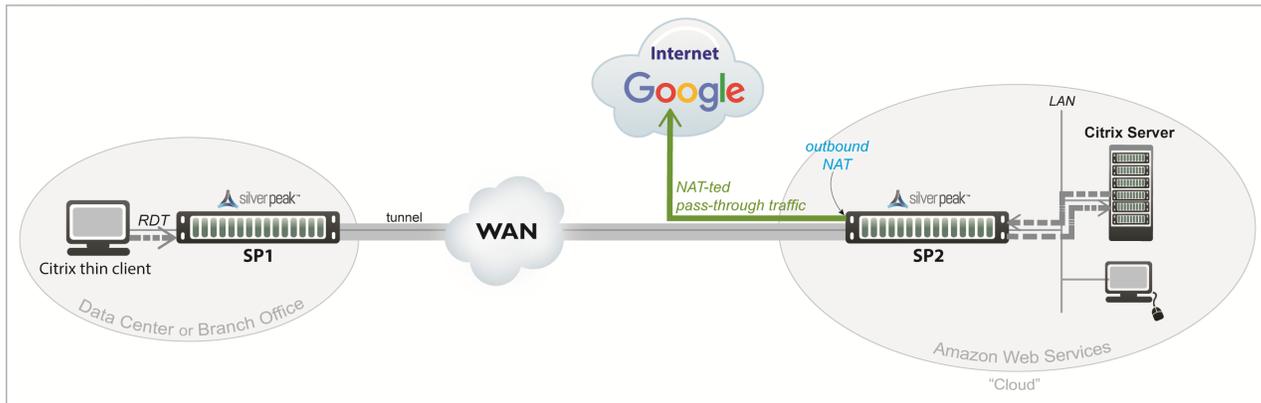
1. **Inbound NAT.** The appliance automatically creates a source NAT (Network Address Translation) map when retrieving subnet information from the Silver Peak Cloud portal. This ensures that traffic destined to SaaS servers has a return path to the appliance from which that traffic originated.

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**:

### Match Criteria

- These are universal across all policy maps – **Route, QoS, Optimization, NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.
- The available parameters are **Application, Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain, Geo Location, Interface, Protocol, DSCP, IP/Subnet, Port, and Traffic Behavior**.

- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

### 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.

### Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

### Set Actions

Set Action	Option	Definition
NAT Type	<b>no-nat</b>	Is the <i>default</i> . No IP addresses are changed.
	<b>source-nat</b>	Changes the source address and the source port in the IP header of a packet.
NAT Direction	<b>inbound</b>	NAT is on the LAN interface.
	<b>outbound</b>	NAT is on the WAN interface.
	<b>none</b>	The only option if the NAT Type is <b>no-nat</b> .
NAT IP	<b>auto</b>	Select if you want to NAT <b>all</b> traffic. The appliance then picks the first available NAT IP/Port.

Set Action	Option	Definition
	<b>tunnel</b>	Select if you only want to NAT <b>tunnel</b> traffic. Applicable only for inbound NAT, as outbound doesn't support NAT on tunnel traffic.
	<b>[IP address]</b>	Select if you want to make NAT use this IP address during address translation.
<b>Fallback</b>		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.

### Merge / Replace

At the top of the page, choose

**Merge** to use the values in the template, but keep any values set on the appliance as is (producing a mix of template and appliance rules),

*-OR-*

**Replace** (recommended) to replace all values with those in the template.

## Inbound Port Forwarding

*Configuration > Policies > Inbound Port Forwarding*

Inbound port forwarding allows traffic from the WAN to reach computers or services within a private LAN when you have a stateful firewall. It helps define and manage inbound traffic, remap a destination IP address and port number to an internal host, and create policies to manage branch devices from the WAN. Use this tab to define the desired inbound traffic.

Inbound Port forwarding is available in two modes when you add or edit a rule, depending if the translate mode is enabled or disabled.

The first operating mode for inbound port forwarding is when translate mode is disabled with inbound port forwarding. The LAN-side subnet with private IP addresses is allowed access through an inbound port forwarding rule (defined by you in the following steps) and exposes any external services. This requires LAN side private addresses to be routed on the WAN side. This represents the process of DMZ (Demilitarized Zone).

**Note:** This mode it not common unless the port forwarding source is directly connected to the EdgeConnect, or if the LAN side device address is routed from the WAN side.

To establish a DMZ connection, complete the following steps:

1. Go to the **Inbound Port Forwarding** tab.
2. Select the **Edit** icon next to **Appliance Name**.
3. Select **Add Rule**.
4. Complete each field with the appropriate information.

**Inbound Port Forwarding** x

Source IP/Subnet	<input type="text" value="10.1.1.0/24"/>		<input type="checkbox"/> Translate	
Destination IP/Subnet	<input type="text" value="10.17.44.20/32"/>	→	Translated IP	<input type="text"/>
Destination Port/Range	<input type="text"/>		Translated Port/Range	<input type="text"/>
Protocol	<input type="text"/>			
Comment	<input style="width: 100%;" type="text"/>			

Field	Definition
<b>Source IP/Subnet</b>	The source of the WAN device managing the LAN device(s) specified in the destination.
<b>Destination IP/Subnet</b>	The address of the LAN device(s) managed remotely.

The second mode is when translate mode is enabled. When enabled, the EdgeConnect WAN interface performs destination NAT to reach LAN side device(s) from an external network.

Complete the following steps to enable the translate mode. This represents the process of DNAT (Destination Network Translation).

1. Go to the **Inbound Port Forwarding** tab.
2. Select the **Edit** icon.
3. Select **Add Rule**.
4. Check the translate box to enable the **Translate** mode.
5. Complete each field with the appropriate information.

Field	Definition
Source IP/Subnet	The source of the WAN device managing the LAN device(s) specified in the destination.
Destination IP/Subnet	The address of the WAN interface IP.
Destination Port/Range	The port/range of the LAN device(s) that are managed remotely.
Protocol	Select the protocol you want to apply: <b>UDP, TCP, ICMP, Any</b> . If you select <b>Any</b> , the Destination and Translated Ports have a default value of need to be between 0-65535.
Translated IP	The IP address of the LAN device accessed inside your network.
Translated Port/Range	The port/range of the LAN device accessed inside your network.

### Additional Information

- **Interface Modes**

Port forwarding is only used when you have 'stateful' or 'stateful+snat' configured on interfaces. It does not apply when you have 'Allow All' or 'Harden' configured.

- **Security Policies**

\*If 'security policies' are configured, make sure they allow the traffic specified in the port forwarding rules.

- You can also reorder the appliances associated with inbound port forwarding by selecting **Reorder** when adding a rule.

**Note:** 'Any' is only a protocol option on versions or 8.1.9.4 and later.

## Security Policies Tab

*Configuration > Policies > Security Policies*

This tab displays the Security Policies, which manage traffic between *zone-based firewalls*.

- Zones are created on the Orchestrator. A zone is applied to an **Interface**.
- By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. Users can create exception rules (Security Policies) to allow traffic between interfaces with different zones.
- Define your Security Policies by creating **templates**. You can then apply templates to Interfaces and/or Overlays.
- Selecting the **Edit** icon opens the Security Policy that has been applied. Any changes made here are **local** to that appliance.
- Selecting **Manage Security Policies with Templates** will allow you to define policies on all appliances within your network. You can use the matrix and table view to further specify your policies.
- Logging: In table view, you can specify the log level when adding and editing a rule. Select the appropriate level from the options in the list.

### Security Policies ?

Matrix View **Table View**    Merge  Replace

Add Rule

1 Rows, 1 Selected Search

From Zone	To Zone	Priority	Match Criteria	Action	Enabled	Logging <span>▼</span>	Tag	Comment
Default	Default	1000	Match Everything <span>↗</span>	deny	<input checked="" type="checkbox"/>	<div style="border: 1px solid #ccc; padding: 5px; width: 150px;">           None <span>▼</span>  <b>None</b>            Emergency            Alert            Critical            Error            Warning            Notice            Info            Debug         </div>		<span>✕</span>

### Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.

- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## Access Lists Tab

*Configuration > Policies > ACLs > Access Lists*

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

Access Lists ? Orchestrator Template Range

1523 Rows, 1 Selected Search

Edit	Appliance Name	ACLs	Priority	Match Criteria	Permit	Comment
	Albuquerque	Overlay_Default	1000	Match Everything	permit	
	Albuquerque	Overlay_BulkApps	1000	Application Box	permit	
	Albuquerque	Overlay_BulkApps	1010	Application Dropbox	permit	
	Albuquerque	Overlay_BulkApps	1020	Application Github	permit	
	Albuquerque	Overlay_BulkApps	1030	Application group File_Sharing	permit	
	Albuquerque	Overlay_BulkApps	1040	Application Ftp	permit	
	Albuquerque	Overlay_BulkApps	1050	Application Sftp	permit	
	Albuquerque	Overlay_BulkApps	1060	Application Rsync-file-synchronization	permit	
	Albuquerque	Overlay_CriticalApps	1010	Application Webex	permit	
	Albuquerque	Overlay_CriticalApps	1030	Application Salesforce	permit	
	Albuquerque	Overlay_CriticalApps	1060	Application Office365	permit	
	Albuquerque	Overlay_CriticalApps	1070	Application Office365Exchange	permit	
	Albuquerque	Overlay_CriticalApps	1090	Application Slack	permit	
	Albuquerque	Overlay_CriticalApps	1100	Application Workday	permit	
	Albuquerque	Overlay_CriticalApps	1110	Application Adp	permit	
	Albuquerque	Overlay_CriticalApps	1120	Application ServiceNow	permit	
	Albuquerque	Overlay_CriticalApps	1130	Application GoToMeeting	permit	
	Albuquerque	Overlay_CriticalApps	1150	Application Atlassian	permit	
	Albuquerque	Overlay_CriticalApps	1160	Application Marketo	permit	
	Albuquerque	Overlay_CriticalApps	1170	Application Tableau	permit	
	Albuquerque	Overlay_CriticalApps	1200	Application Citrix	permit	

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).

## Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.
- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.64-95** is not supported. The correct way to specify this range is **10.130-139\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

# Shaper Tab

Configuration > [Policies > Shaping] 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.

Shaper x

Manage Shaper settings with Templates Inbound Outbound Export ↺ 8 mins

Shaper ?

300 Rows Search

Edit	Host Na...	Inter...	Max Wan ...	Recalc on...	Traffic ID	Traffic Na...	Priority	Min BW %	Min BW Absolute ...	Min BW Actual (...)	Excess Weight...	Max BW %	Max BW Absolute ...	Max BW Actual ...	Max Wait Time ...	Rate Limit (kb...	Enable
✓	Albuque...	wan	200,000	No	9	UNUSED9	9	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Albuque...	wan	200,000	No	10	UNUSED10	10	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	1	default	1	0	0	0	250	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	2	Interactive	1	0	0	0	1000	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	3	RealTime	1	0	0	0	500	100	10,000,000	200,000	100	0	Yes
✓	Barcelona	wan	200,000	No	4	replication	1	0	0	0	100	100	10,000,000	200,000	1000	0	Yes
✓	Barcelona	wan	200,000	No	5	guest_vir...	1	0	0	0	100	100	10,000,000	200,000	1000	0	Yes
✓	Barcelona	wan	200,000	No	6	UNUSED6	6	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	7	UNUSED7	7	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	8	UNUSED8	8	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	9	UNUSED9	9	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Barcelona	wan	200,000	No	10	UNUSED10	10	0	0	0	1	100	10,000,000	200,000	500	0	Yes
✓	Boston	wan	200,000	No	1	default	1	0	0	0	250	100	10,000,000	200,000	500	0	Yes

Shaper x

Manage Shaper settings with Templates Inbound Outbound Export ↺ 8 mins

Shaper ?

30 Rows Search

Edit	Host N...	Interfa...	Max Wan ...	Recalc on...	Traffic ID	Traffic Na...	Priority	Min BW %	Min BW Absolute ...	Min BW Actual (...)	Excess Weight...	Max BW %	Max BW Absolute ...	Max BW Actual ...	Max Wait Time ...	Rate Limit (kb...	Enable
✓	Albuqu...																
✓	Barcel...																
✓	Boston																
✓	Chennai																
✓	Chicago																
✓	Dallas																
✓	Denver																
✓	Edinbu...																
✓	Frankfurt																
✓	Geneva																
✓	London																
✓	Los-An...																
✓	Mexico...																

- 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 tunneled and pass-through-shaped traffic --- shaping it as it exits to the WAN.
- To manage Shaper settings for an appliance's system-level **WAN Shaper**, access the Shaper template.
- For minimum and maximum bandwidth, you can configure traffic class values as a percentage of total available system bandwidth and as an absolute value. The appliance always provides the larger of the minimum values, and limits bandwidth to the lower of the maximum values.
- If you set **Min Bandwidth** to a value greater than **Max Bandwidth**, then **Max** overrides **Min**.

### Shaper Tab Settings

Field Name	Description
<b>Excess Weighting</b>	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 <b>Excess Weighting</b> column. Values range from 1 to 10,000.
<b>Interface Shaper</b>	Enables a separate shaper for a specific WAN interface. <ul style="list-style-type: none"> <li>■ For WAN optimization, the interface shaper can be used but is not recommended.</li> <li>■ For SD-WAN, it should never be used because overlay traffic isn't directed to an interface shaper; traffic is always shaped by the default WAN shaper.</li> </ul>
<b>Max Bandwidth %</b>	This limits the maximum bandwidth that a traffic class can use to a percentage of total available system bandwidth.
<b>Max Bandwidth Absolute (kbps)</b>	This limits the maximum bandwidth that a traffic class can use to an absolute value (kbps). You can specify a maximum absolute value to cap the bandwidth for downloads and streaming.
<b>Max Wait Time</b>	Any packets waiting longer than the specified <b>Max Wait Time</b> are dropped.
<b>Min Bandwidth %</b>	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 <b>Min Bandwidth</b> to a value greater than <b>Max Bandwidth</b> , then <b>Max</b> overrides <b>Min</b> .

Field Name	Description
Min Bandwidth Absolute (kbps)	This guarantees a specific level of service when total system bandwidth declines. This is useful for maintaining the quality of VoIP, for example.
Priority	Determines the order in which to allocate each class's minimum bandwidth - <b>1</b> is first, <b>10</b> is last.
Rate Limit (kbps)	You can set per-flow rate limit that a traffic class uses by specifying a number in the Rate Limit column. For no limit, use <b>0</b> (zero).
Recalc on IF State Changes	When an interface state changes to UP or DOWN, selecting this recalculates the total bandwidth based on the configured bandwidth of all UP interfaces. For example, when <b>wan0</b> goes down, <b>wan0</b> bandwidth is removed from the total bandwidth when recalculating.
Traffic ID	The number assigned to the traffic class.
Traffic Name	The name assigned to a traffic class, either prescriptively or by the user.

## SaaS Optimization Tab

Configuration > [Policies > Applications & SaaS] 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.

SaaS Optimization Configuration

161 Rows

Edit	Appliance Name	Application Name	Optimize	Advertise	RTT Threshold	Domains
✓	Chennai	Adobe	No	No	10 ms	adobe.com
✓	Chennai	AirWatch	No	No	10 ms	*.air-watch.com
✓	Chennai	AthenaHealth	No	No	10 ms	*.athenahealth.com, athenahealth.com
✓	Chennai	Box	No	No	10 ms	*.app.box.com, *.box.com, *.box.net, *.boxcdn.net, *.boxcloud.com
✓	Chennai	CCConc	No	No	10 ms	mycccportal.com, *.mycccportal.com
✓	Chennai	ConstantContact	No	No	10 ms	constantcontact.com
✓	Chennai	CornerstoneOnDemand	No	No	10 ms	cornerstoneondemand.com
✓	Chennai	Dropbox	No	No	10 ms	*.dl.dropboxusercontent.com, dropbox.com, *.dropbox.com
✓	Chennai	Eloqua	No	No	10 ms	eloqua.com, eloquatrainingcenter.com
✓	Chennai	GoToAssist	No	No	10 ms	gototraining.com
✓	Chennai	GoToMeeting	No	No	10 ms	gotomeeting.com
✓	Chennai	GoToTraining	No	No	10 ms	gototraining.com
✓	Chennai	GoToWebinar	No	No	10 ms	gotowebinar.com, gotoassist.com
✓	Chennai	Intuit	No	No	10 ms	intuit.com
✓	Chennai	Jobvite	No	No	10 ms	careers.jobvite.com, www.jobvite.com, hire.jobvite.com
✓	Chennai	Lithium	No	No	10 ms	lithium.com
✓	Chennai	LiveOps	No	No	10 ms	liveops.com
✓	Chennai	Marketo	No	No	10 ms	marketo.com
✓	Chennai	NetSuite	No	No	10 ms	netsuite.com
✓	Chennai	Office365	No	No	10 ms	*.officeapps.live.com, *.microsoftonline-p.net, *.microsoftonlinesupport.net, ...
✓	Chennai	OneNote	No	No	10 ms	onenote.com, *.onenote.com, *.onenote.net
✓	Chennai	Parature	No	No	10 ms	parature.com
✓	Chennai	PardotExactTarget	No	No	10 ms	pardot.com
✓	Chennai	Planner	No	No	10 ms	tasks.office.com, *.tasks.osi.office.net, controls.office.com
✓	Chennai	PlexSystems	No	No	10 ms	plex.com
✓	Chennai	Salesforce	No	No	10 ms	*.eu4.force.com, *.na3.force.com, *.salesforce.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. The parameters for determining Round Trip Times include the following:

- The **RTT Calculation Interval** specifies how frequently Orchestrator recalculates the Round Trip Time for the enabled Cloud applications.
  - The **RTT Ping Interface** specifies which interface to use to ping the enabled SaaS subnets for Round Trip Times. The **default** interface is **wan0**.
- 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

SaaS Optimization Monitoring

158 Rows

Appliance Name	Application Name	Subnet	Server IP	Advertised	RTT	RTT Threshold	Ping Method	Ping Port	Resolved Location
Chennai	No SaaS Optimization defined for this appliance.								
Chicago	Adobe	173.240.108.176...	173.240.108.191	No	Unreachable	10 ms			San Jose, United States
Chicago	Adobe	173.240.110.128...	173.240.110.140	No	Unreachable	10 ms			San Jose, United States
Chicago	Box	107.152.24.0/24	107.152.24.192	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	107.152.27.0/24	107.152.27.192	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	74.112.185.0/24	74.112.185.67	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	74.112.184.0/24	74.112.184.67	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	74.112.186.0/24	74.112.186.67	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	107.152.25.0/24	107.152.25.192	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	107.152.26.0/24	107.152.26.192	No	Unreachable	10 ms			Los Altos, United States
Chicago	Box	208.184.35.0/27	208.184.35.1	No	Unreachable	10 ms			Santa Clara, United States
Chicago	Box	208.184.35.32/27	208.184.35.33	No	Unreachable	10 ms			Santa Clara, United States
Chicago	Jobvite	54.210.23.251/32	54.210.23.251	No	Unreachable	10 ms			Ashburn, United States
Chicago	Jobvite	52.7.65.30/32	52.7.65.30	No	Unreachable	10 ms			Ashburn, United States
Chicago	Salesforce	136.146.42.0/24	136.146.42.0	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.146.0/24	96.43.146.15	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.145.0/24	96.43.145.15	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.144.0/24	96.43.144.15	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	85.222.137.0/24	85.222.137.5	No	Unreachable	10 ms			Staines, United Kingdom
Chicago	Salesforce	136.146.52.0/24	136.146.52.0	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.153.0/24	96.43.153.31	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.152.0/24	96.43.152.31	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.151.0/24	96.43.151.9	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.150.0/24	96.43.150.15	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.149.0/24	96.43.149.15	No	Unreachable	10 ms			San Francisco, United States
Chicago	Salesforce	96.43.148.0/24	96.43.148.15	No	Unreachable	10 ms			San Francisco, United States

# Application Definitions

Configuration > [Policies > Applications & SaaS] Application Definitions

This tab provides application visibility and control. You can search to see if Silver Peak has a definition for a specific application and, if so, how it's defined.

The screenshot displays the 'Application Definitions' interface. At the top, there is a search bar containing the text 'google'. Below the search bar, there are several tabs: 'Application Groups', 'Modified', 'Disabled', and 'Auto updates' (set to 'ON'). The main content area shows a table with 134 rows of application definitions. The table has columns for 'Type', 'Name', 'Notes', 'Confid.', 'Detail', and 'Edit'. The rows list various Google services and their corresponding domain names, such as 'Google', 'Elgoog', 'Contacts-goo...', 'GoogleDocs', 'GoogleDomains', 'DoubleClick', 'GoogleDrive', 'Google', 'GoogleFonts', 'Google', 'GoogleForms', 'Gg-google', and 'Ggoogle'. Below this table, there is a section for 'Hide Advanced App Definitions' with a dropdown menu. Underneath, there is a table with 140 rows of IP protocols. The table has columns for 'Protocol', 'Name', 'Notes', 'Confidence', and 'Edit'. The rows list various protocols like 'Hopopt', 'Icmp', 'Icmp', 'Ggp', 'Ipv4', 'St', 'Trp', 'Cbt', 'Egp', 'Igp', 'Bbn-rcv-man', 'Nsp-ii', 'Pup', 'Argus', and 'Emcon'.

- Orchestrator uses these eight dimensions for identifying and defining applications:
  - IP Protocol
  - UDP Port
  - TCP Port
  - Domain Name
  - Address Map - (formerly known as *IP Intelligence*). Given a range of IP addresses, the Address Map reveals what organization owns the segment, along with the country of origin.
  - DPI (Deep Packet Inspection). An expanded list of Orchestrator's legacy built-in applications.
  - Compound - Created by user from multiple criteria.

- **SaaS** - Created by user. If any components of the definition change, the user must manually update the definition.
- You can modify or disable an existing application.
- You can use any of the dimensions to define a new application.
- **Auto update** is enabled by default.

# Application Groups Tab

*Configuration > [Policies > Applications & SaaS] 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.

The screenshot shows the 'Application Groups' tab in the Silver Peak Orchestrator. On the left, there are two filter sections: 'Traffic Type' and 'Content Type'. The 'Interactive' group is selected in the 'Traffic Type' filter, showing a count of 139. The main table lists 139 rows of applications, each with a 'Groups' column and an 'Edit Group Membership' icon. The 'Interactive' group is highlighted in blue for all rows.

Application	Groups	Edit Group Membership
Adobe	Computer_and_Electronics, Interactive, SaaS, Software, Video	✎
Airs	Interactive	✎
Ammyy	Interactive	✎
Aol	Email, Interactive, Internet_and_Telecom	✎
Apple-remote-desktop	Interactive	✎
Asf-rmcp	Interactive, Network_Services	✎
Avira	Computer_and_Electronics, Computer_Security, Interactive	✎
BlueJeans	Interactive, SaaS	✎
Bluestacks	Computer_and_Electronics, Interactive, Software	✎
Brocade	Encrypted, Interactive	✎
Cddbp	Interactive	✎
Cisco	Computer_and_Electronics, Encrypted, Interactive, Network_Services, Networking, Real-Time, Vide...	✎
Cisco-aon-amc	Interactive	✎
Citadel	Interactive	✎
Citrix-ica	Citrix, File_Sharing, Interactive	✎
Codenger	Interactive	✎
Dart	Interactive, Network_Services	✎
Dcn-meas	Interactive	✎
Default-port	Interactive	✎
Dtspcd	Interactive	✎
Farming	Interactive	✎

- The **Group Name** cannot be empty.
- Group names are case-insensitive.
- A group can be empty or contain up to 128 applications.
- An application group cannot contain an application group.
- Group name followed by \* is a group defined by a user.
- You are not allowed to change the group name for groups provided by Silver Peak. You are allowed to add or delete applications within those groups.

# Threshold Crossing Alerts Tab

Configuration > [Policies > TCAs] Threshold Crossing Alerts

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

Threshold Crossing Alerts x

Manage TCAs with Templates

System Tunnel Export

Threshold Crossing Alerts ?

36 Rows Search

Edit	Appliance Name	Name	Rising				Falling			
			Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
✓	Talinn	File-system utilization	90%	85%	1	Yes	75%	75%	1	No
✓	Talinn	LAN-side receive throughput	1000000 kbps	1000000 kbps	1	No	0 kbps	0 kbps	1	No
✓	Talinn	Total number of flows	90%	85%	1	Yes	0%	0%	1	No
✓	Talinn	Total number of optimized flows	90%	85%	1	No	0%	0%	1	No
✓	Talinn	Tunnel OOP post-POC	100%	100%	1	No	0%	0%	1	No
✓	Talinn	Tunnel OOP pre-POC	100%	100%	1	No	0%	0%	1	No
✓	Talinn	Tunnel latency	1000 ms	850 ms	1	Yes	0 ms	0 ms	1	No
✓	Talinn	Tunnel loss post-FEC	100%	100%	1	No	0%	0%	1	No
✓	Talinn	Tunnel loss pre-FEC	100%	100%	1	No	0%	0%	1	No
✓	Talinn	Tunnel reduction	100%	100%	1	No	0%	0%	1	No
✓	Talinn	Tunnel utilization	100%	100%	1	No	0%	0%	1	No
✓	Talinn	WAN-side transmit throughput	1000000 kbps	1000000 kbps	1	No	0 kbps	0 kbps	1	No
✓	laine-vxa	File-system utilization	90%	85%	1	Yes	75%	75%	1	No
✓	laine-vxa	LAN-side receive throughput	1000000 kbps	1000000 kbps	1	No	0 kbps	0 kbps	1	No
✓	laine-vxa	Total number of flows	90%	85%	1	Yes	0%	0%	1	No
✓	laine-vxa	Total number of optimized flows	85%	80%	1	No	0%	0%	1	No
✓	laine-vxa	Tunnel OOP post-POC	100%	100%	1	No	0%	0%	1	No
✓	laine-vxa	Tunnel OOP pre-POC	100%	100%	1	No	0%	0%	1	No
✓	laine-vxa	Tunnel latency	1000 ms	850 ms	1	Yes	0 ms	0 ms	1	No
✓	laine-vxa	Tunnel loss post-FEC	100%	100%	1	No	0%	0%	1	No

Threshold Crossing Alerts x

Manage TCAs with Templates

System Tunnel Export

Threshold Crossing Alerts ?

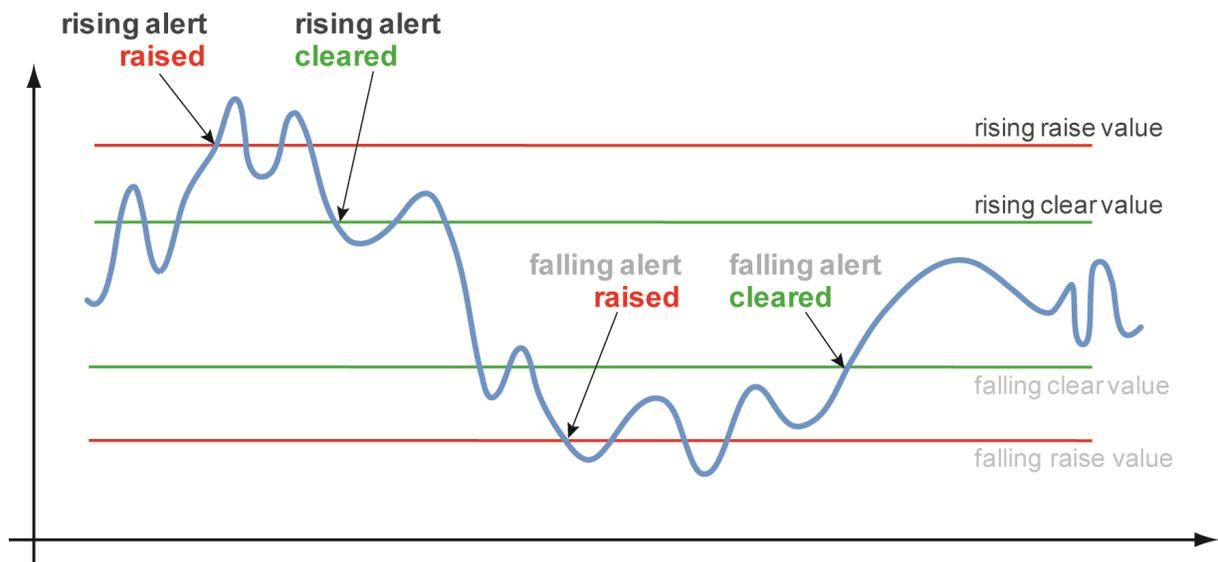
1 Rows Search

Appliance Name	Tunnel Name	TCA Name	Rising				Falling			
			Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
laine-vxa	auto_tun_10.1.153.20_to_10.1.1...	latency	20 ms	15 ms	NaN	No	undefined ms	undefined ms	NaN	No

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

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)

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

*Defaults values for Threshold Crossing Alerts*

TCA Name	Default [ON, OFF]	Default Values [Rising Raise, Rising Clear, Falling Raise, Falling Clear]	allow rising	allow falling
<b>Appliance Level</b>				
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

TCA Name	Default [ON, OFF]	Default Values [Rising Raise, Rising Clear, Falling Raise, Falling Clear]	allow rising	allow falling
File-system-utilization	ON <sup>a</sup>	95%; 85%; 0%; 0%	4	--
<b>Tunnel Level</b>				
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

---

<sup>a</sup>Cannot be disabled.

## IP SLA Tab

Configuration > [Policies > TCAs] IP SLA

6 Rows	Search								
Edit	Appliance N...	Active	State	Monitor	Down Action	Up Action	Comment	Up Stats	Down Stats
	Chennai			IP SLA managers are not configured for this appliance.					
	Mumbai			IP SLA managers are not configured for this appliance.					
	Osaka			IP SLA managers are not configured for this appliance.					
	Seoul			IP SLA managers are not configured for this appliance.					
	Singapore			IP SLA managers are not configured for this appliance.					
	Tokyo			IP SLA managers are not configured for this appliance.					

Using a polling process, **IP SLA** (Internet Protocol Service Level Agreement) tracking provides the ability to generate specific actions in the network that are completely dependent on the state of an IP interface or tunnel. The goal is to prevent black-holed traffic. For example, associated IP subnets could be removed from the subnet table, and also from subnet sharing, if the LAN-side interfaces on an appliance go down.

Four **Monitors** are available:

<b>Interface</b>	Monitors the operational status of a specific local interface.
<b>Ping</b>	Monitors the reachability of a specific IPv4 address.
<b>HTTP/HTTPS</b>	Monitors the reachability of an HTTP/HTTPS endpoint.
<b>VRRP Monitor</b>	Monitors the VRRP router state (TRUE if Master, FALSE if Backup) for a VRRP instance(s) on an interface.

Based on the Monitor chosen, the Web UI displays the appropriate fields and options.

There are eight available **Down Actions**:

<b>Remove Auto Subnet</b>	Remove from the subnet table an auto subnet for a port (including all VLAN & subinterface subnets).
<b>Increase VRRP Priority</b>	Increase the configured VRRP router priority by a delta amount.
<b>Decrease VRRP Priority</b>	Decrease the configured VRRP router priority by a delta amount.
<b>Enable Tunnel</b>	Enable a passthrough (internet breakout) tunnel Up for IP Tracking (SLA) purposes.

<b>Disable Tunnel</b>	Disable a passthrough (internet breakout) tunnel Up for IP Tracking (SLA) purposes. The tunnel can no longer be used for load balancing purposes (when load balancing traffic between multiple passthrough tunnels), although it can still be used as a last resort for traffic forwarding.
<b>Disable Subnet Sharing</b>	Disable subnet sharing of subnets to other Silver Peak peers on the appliance.
<b>Modify Subnet Metric</b>	Add a metric delta to the metric of all subnets shared with Silver Peak peers.
<b>Advertise Subnets</b>	Advertise subnets to Silver Peak peers.

There are two default **Up Actions**:

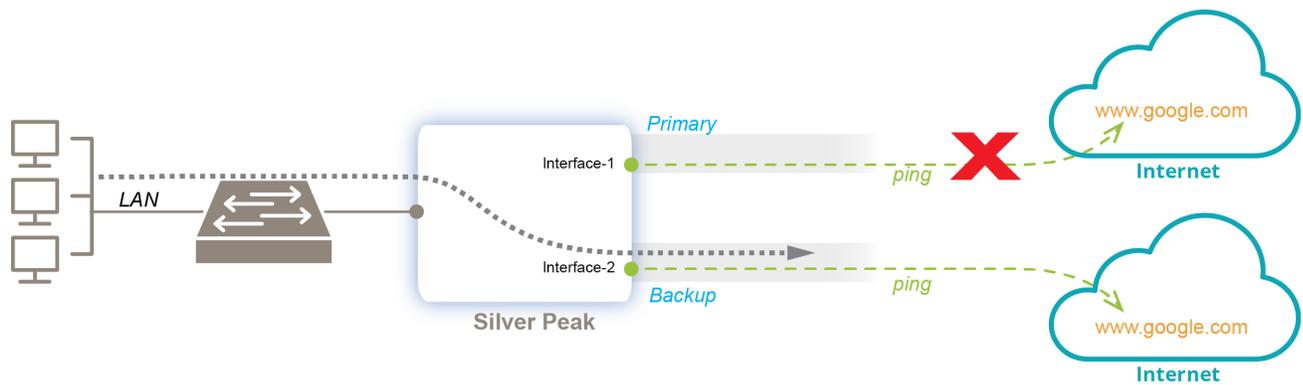
<b>Default Subnet Action</b>	This reverts whatever was the <b>Down Action</b> back to the normal state. For example: <ul style="list-style-type: none"> <li>■ If <b>Down Action = Disable Subnet Sharing</b>, the Up Action is re-enable Subnet Sharing.</li> <li>■ If <b>Down Action = Remove Auto Subnets</b>, the Up Action re-adds the auto subnet.</li> <li>■ If <b>Down Action = Modify Subnet Metric</b>, the Up Action restores subnet metrics to their original value.</li> </ul>
<b>VRRP Default</b>	Reverts the VRRP priority back to the configured value.

NOTE: If a default **Up Action** is used, it must match the **Down Action**.

## Monitor Use Cases

Following are five basic use cases.

### Example #1 - Ping via Interface



- Two passthrough tunnels configured for Internet breakout and High Availability.
- If the Primary passthrough tunnel goes down, traffic goes to Backup tunnel.
- The **IP SLA Rule** would look like this, with the same tunnel specified for the **Down** and **Up Actions**.

**IP SLA Rule** [X]

Monitor ON OFF

Monitor: Ping

Address: 8.8.8.8

Interface: Internet

Keep Alive Interval: 1 (Sec)

Up Threshold: 3 (Sec)

Down Threshold: 30 (Sec)

Interval: 30 (Sec)

**Actions**

Down Action: Disable Tunnel

Tunnel: Passthrough\_Internet\_1

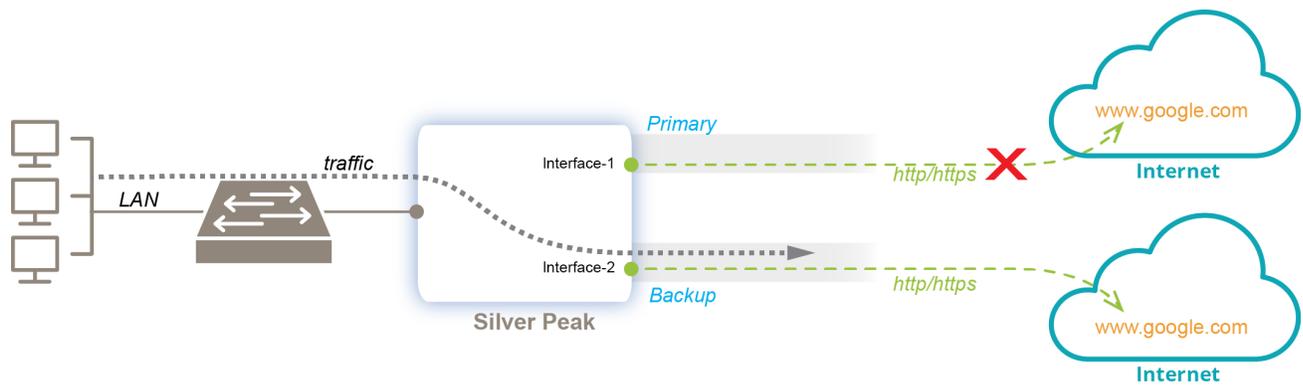
Up Action: Enable Tunnel

Tunnel: Select Tunnel

Comment:

- Select Tunnel
- Passthrough\_Internet\_Voice
- Passthrough\_Internet\_Guest\_Wifi
- Passthrough\_Internet\_Gold\_Cloud
- Passthrough\_Internet\_Email
- Passthrough\_Internet\_Video**
- Passthrough\_Internet\_Everyday\_Cloud
- Passthrough\_MPLS\_Voice
- Passthrough\_MPLS\_Guest\_Wifi
- Passthrough\_MPLS\_Gold\_Cloud
- Passthrough\_MPLS\_Email
- Passthrough\_MPLS\_Video
- Passthrough\_MPLS\_Everyday\_Cloud

## Example #2 - HTTP/HTTPS via Interface



- Two passthrough tunnels configured for Internet breakout and High Availability.
- If the Primary passthrough tunnel goes down, traffic goes to Backup tunnel.
- The **IP SLA Rule** would look like this, with the same tunnel specified for the **Down** and **Up** Actions.

**IP SLA Rule**
✕

**Monitor**

Monitor: HTTP/HTTPS

URL(s): www.google.com

Proxy Address: optional

Proxy Port: (0..65535)

User Agent: optional

HTTP Request Timeout: 60 Sec

Interface: Internet

Keep Alive Interval: 90 Sec

Mark Up after X Succeed: 2

Mark Down after X Failed: 3

Monitor Sampling Interval: 60 Sec

**Actions**

Down Action: Disable Tunnel

Tunnel: Passthrough\_Internet\_Interactive

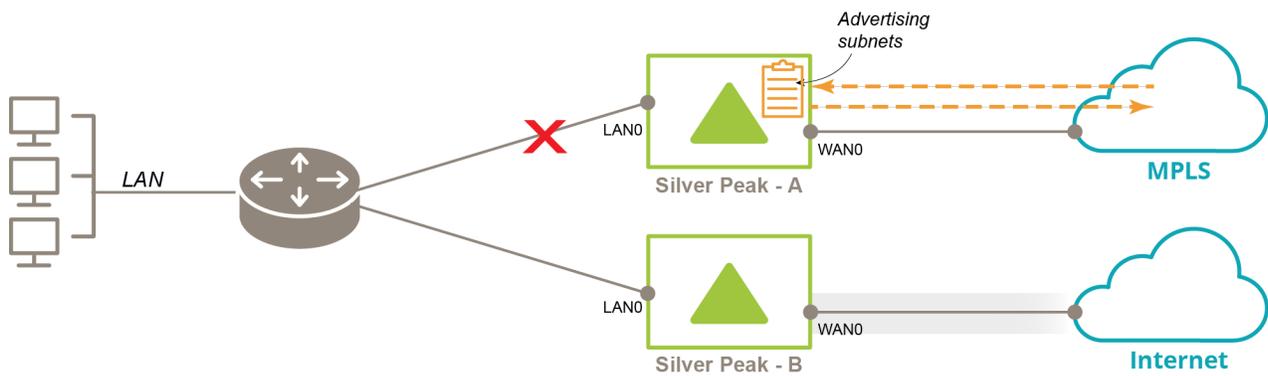
Up Action: Enable Tunnel

Tunnel: Select Tunnel

Comment:

- In the **URL(s)** field, the protocol identifier is only required when specifying HTTPS, as in **https://www.google.com**.

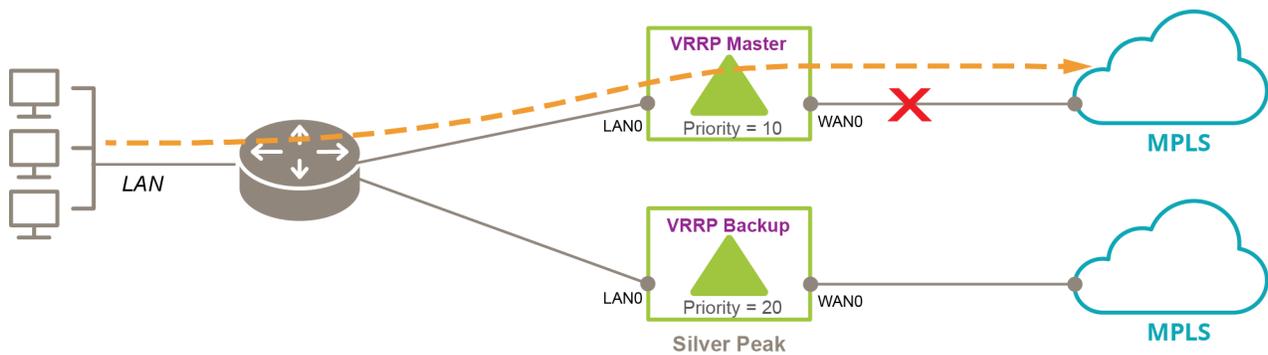
### Example #3 - Monitor Interface (LAN0)



- On *Silver Peak - A*, we want subnet advertising to be conditional on **LAN0** being up.
- Its **IP SLA Rule** would look like this, with the **Default Subnet Action** being to resume advertising subnets.



## Example #4 - Monitor Interface (WAN0) to ensure High Availability



- If **WAN0** goes down on the **VRRP Master**, we want to decrease its Priority so that traffic goes to the **VRRP Backup**.
- Its **IP SLA Rule** would look like this, with the **Default Subnet Action** being to revert to the original Priority.

**IP SLA Rule** [X]

Monitor ON OFF

Monitor

Interface

Interval  (Sec)

**Actions**

Down Action

Interface

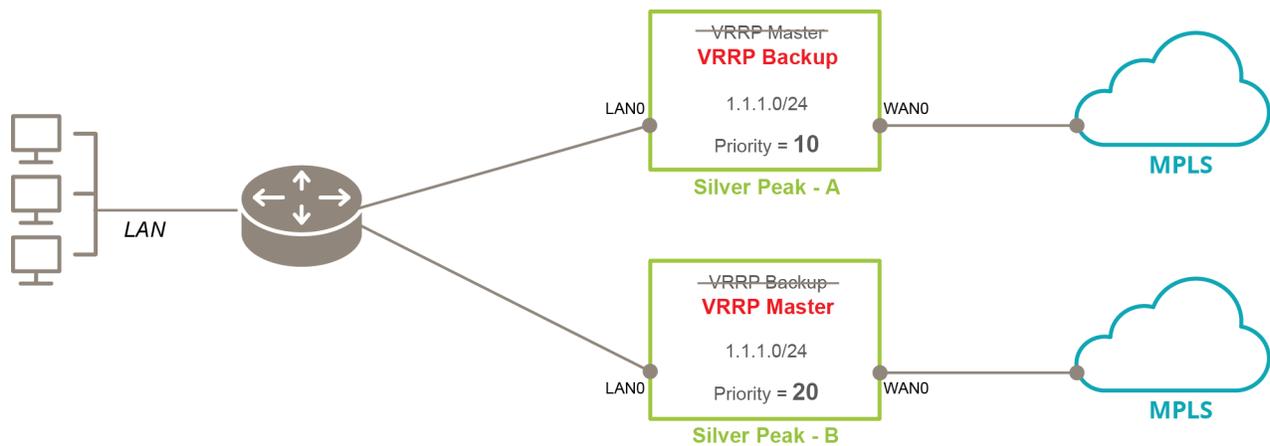
Priority

Up Action

Comment

NOTE: In this instance, the **WAN0** interface was given the label, **MPLS**, to match the service to which it connected.

### Example #5 - Monitor VRRP



- To monitor the VRRP router state, use **VRRP Monitor** and specify the interface on which the VRRP instance is configured.

In this example, it's **LAN0**.

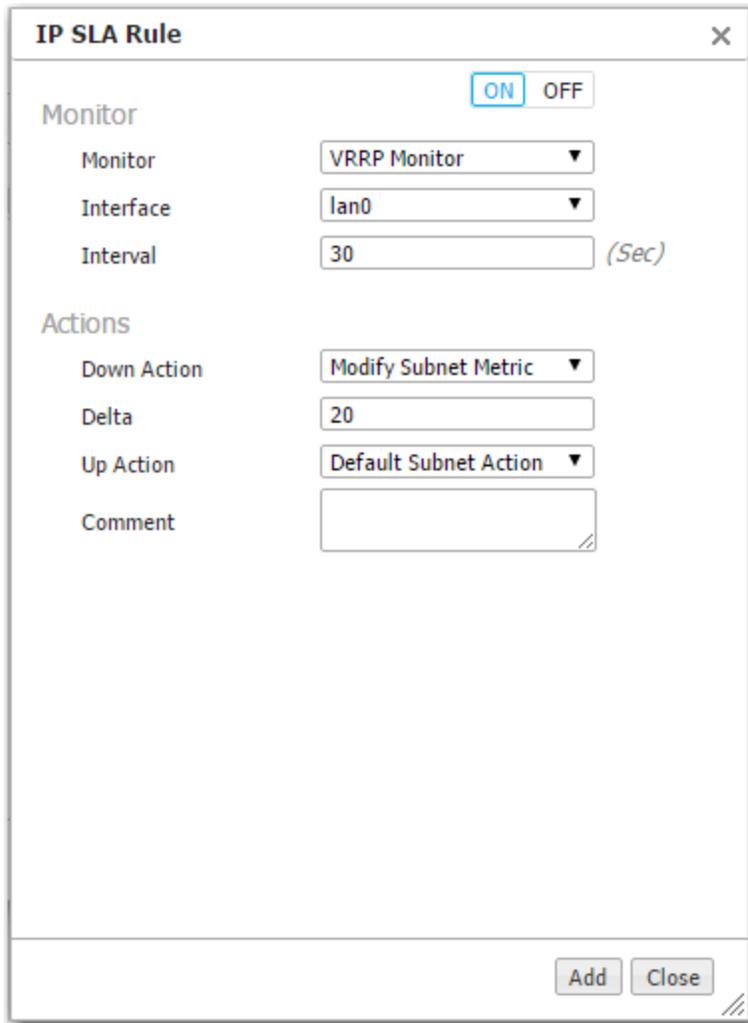
- Here, we're looking at an instance where the VRRP role changes, but priority doesn't, for whatever reason.
- Its **IP SLA Rule** would look like this, with the **Default Subnet Action** being to revert to the original Priority.

The screenshot shows the IP SLA Rule configuration window. The Monitor section is set to VRRP Monitor on the lan0 interface with a 30-second interval. The Down Action is Disable Subnet Sharing, the Up Action is Default Subnet Action, and the Comment is Monitors VRRP router state.

**NOTE:** In this instance, the **WAN0** interface was given the label, **MPLS**, to match the service to which it connected.

- Another option would be to specify **Down Action = Modify Subnet Metric**. The Web UI automatically produces another field where you can add a positive value to the current subnet

metric. **Up Action = Default Subnet Action** would return the subnet metric to its original value.



The image shows a configuration dialog box titled "IP SLA Rule" with a close button (X) in the top right corner. The dialog is divided into two main sections: "Monitor" and "Actions".

**Monitor Section:**

- At the top right of this section are two buttons: "ON" (highlighted in blue) and "OFF".
- "Monitor": A dropdown menu currently showing "VRRP Monitor".
- "Interface": A dropdown menu currently showing "lan0".
- "Interval": A text input field containing "30", followed by the label "(Sec)".

**Actions Section:**

- "Down Action": A dropdown menu currently showing "Modify Subnet Metric".
- "Delta": A text input field containing "20".
- "Up Action": A dropdown menu currently showing "Default Subnet Action".
- "Comment": A text area for entering a comment, currently empty.

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

# Configuration Templates

This section describes the templates used for assigning common **Configuration** parameters across appliances.

## 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. Some templates support a MERGE option; refer to the Help on those templates.
- To edit a template, drag (or double-click) it from the Available Templates column to the Active Templates column.
- To save the edits as a new template group, click **Save As**.
- To apply templates, click **Apply Template Groups** at the bottom of the page. This will bring you to the Apply Templates tab where you can permanently associate appliances with specific template groups.
- Associating an appliance with a template group makes Orchestrator automatically keep the templates in sync with the appliance.
- When returning to the Templates page, the Template Group field defaults to showing the last template group viewed.

# System Template

Use this page to configure system-level features.

## System [?](#)

### Optimization

- IP Id auto optimization
- TCP auto optimization
- Flows and tunnel failure

### Subnet Sharing

- Use shared subnet information
- Automatically include local LAN subnets
- Automatically include local WAN subnets
- Metric for local subnets
- Redistribute learned BGP routes to Silver Peak peers
- Allow WAN to WAN routing

### Network Memory

- Encrypt data on disk

Category	Field	Definition
Optimization	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).
	Flows and tunnel failure	<p>If there are parallel tunnels and one fails, then <i>Dynamic Path Control</i> determines where to send the flows. There are three options:</p> <p><b>fail-stick</b>      When the failed tunnel comes back up, the flows don't return to the original tunnel. They stay where they are.</p> <p><b>fail-back</b>      When the failed tunnel comes back up, the flows return to the original tunnel.</p> <p><b>disable</b>          When the original tunnel fails, the flows aren't routed to another tunnel.</p>

Category	Field	Definition
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 LAN subnets	Adds the local LAN subnet(s) to the appliance subnet information.
	Automatically include local WAN subnets	Adds the local WAN subnet(s) to the appliance subnet information.
	Metric for local subnets	Specifies 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.
	Redistribute learned BGP routes to SP peers	Enables subnet sharing of routes (subnets) learned from BGP peers.
	Allow WAN to WAN routing	Redirects inbound LAN traffic back to the WAN.
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 

## NextHop Health Check

Enable Health check Retry count  (1..255)Interval  (1..255) secondsHold down count  (1..255)

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 <b>bypass</b> flows. Or, you can choose to <b>drop</b> the packets.
NextHop Health Check	Enable Health check	Activates pinging of the next-hop router.
	Retry count	Specifies the number of ICMP echoes to send, without receiving a reply, before declaring that the link to the WAN next-hop router is down.
	Interval	Specifies the number of seconds between each ICMP echo sent.
	Hold down count	If the link has been declared down, this specifies how many successful ICMP echoes are required before declaring that the link to the next-hop router is up.

## Miscellaneous

- SSL optimization for non-IPSec tunnels
- Bridge Loop Test
- Enable IGMP snooping
- Always send pass-through traffic to original sender
- Enable default DNS lookup
- Enable HTTP/HTTPS snooping
- Quiescent tunnel keep alive time  (1..65535) seconds
- UDP flow timeout  (1..65535) seconds
- Non-accelerated TCP Flow Timeout  (1..65535) secs
- Maximum TCP MSS  (500..9000) bytes
- NAT-T keep alive time  (0..65535) seconds
- Tunnel Alarm Aggregation Threshold  Tunnel Alarms  
Raise only 1 alarm above this threshold
- Maintain end-to-end overlay mapping
- IP Directed Broadcast

---

<b>Miscellaneous</b>	<b>SSL optimization for non-IPSec tunnels</b>	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.
----------------------	---	---

<b>Bridge Loop Test</b>	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.
<b>Enable IGMP Snooping</b>	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.
<b>Always send pass-through traffic to original sender</b>	If the tunnel goes down when using WCCP and PBR, traffic that was intended for the tunnel is sent back the way it came.
<b>Enable default DNS lookup</b>	Allows the appliance to snoop the DNS requests to map domains to IP addresses. This mapping can then be used in ACLs for traffic matching.
<b>Maintain end to end Overlay Mapping</b>	Enforces the same overlay to be used end-to-end when traffic is forwarded on multiple nodes.
<b>Enable HTTP/HTTPS snooping</b>	Enables a more granular application classification of HTTP/HTTPS traffic, by inspection of the HTTP/HTTPS header, Host. This is enabled by default.
<b>IP Directed Broadcast</b>	Allows an entire network to receive data that only the target subnet initially receives.
<b>Quiescent tunnel keep alive time</b>	Specifies the rate at which to send keep alive packets after a tunnel has become idle (quiescent mode). The default is 60 seconds.
<b>UDP flow timeout</b>	Specifies how long to keep the UDP session open after traffic stops flowing. The default is 120 seconds (2 minutes).
<b>Non-accelerated TCP Flow Timeout</b>	Specifies how long to keep the TCP session open after traffic stops flowing. The default is 1800 seconds (30 minutes).
<b>Maximum TCP MSS</b>	(Maximum Segment Size). The default value is 9000 bytes. This ensures that packets are not dropped for being too large. You can adjust the value (500 to 9000) to lower a packet's MSS.
<b>NAT-T keep alive time</b>	If a device is behind a NAT, this specifies the rate at which to send keep alive packets between hosts, in order to keep the mappings in the NAT device intact.



## 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 authorization servers 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** = Remote first
  - **Second** = Local. 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 Role** = 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](#).

Templates x

Template Group ?

Default Template Group ▼

+Add -Delete

Active Templates Show All >

**General Settings**

System

**SNMP**

Policies

Shaper

Access Lists

Save Save As Cancel

*Applies to all templates in group*

Apply Template Groups

**SNMP ?**

**SNMP**

Enable SNMP

Enable SNMP Traps

Default Trap Community

**SNMP V1/V2**

Enable SNMP V1/V2

Read-Only Community

**SNMP V3**

Add

Enabled	Username	Authentication		Privacy		
		Type	Password	Type	Password	
<input type="checkbox"/>	admin	SHA1		AES-128		✕

**Trap Receivers**

Add

Host	Version	Community/Username	Enabled	

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.

## Flow Export Template

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

Templates ×

Template Group ?  
Default Template Group ▾  
+Add -Delete

Active Templates Show All >

General Settings  
System  
Flow Export

Save Save As Cancel  
*Applies to all templates in group*

Apply Template Groups

Flow Export ?

Enable Flow Exporting

Active Flow Timeout  (1..30) mins

IPFIX Template Timeout  (1..1440) mins

Traffic Type

WAN TX  
 WAN RX  
 LAN RX  
 LAN TX

Collectors

Add

Collectors	IP Address	Port	Collector Type

- 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 and IPFIX collectors.
- **Enable Flow Exporting** 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/IPFIX 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 **WAN TX**.

## DNS Proxy Policies

*Configuration > Templates & Networking > Templates*

If you select ON, Complete the following steps to configure and define your DNS Proxy policies.

**NOTE** This feature is only configurable if you have loopback interfaces configured.

1. Choose if you want the DNS Proxy enabled by selecting **ON** or **OFF**.
2. Select the name of the loopback interface or LAN-side label associated with your DNS proxy.
3. Enter the IP addresses for Server A in the **Server A Addresses** field.
4. Choose if you want Caching to be **ON** or **OFF**. If selected, the domain name to the IP address mapping is cached. By default, caching is **ON**.
5. Enter the domain names of the Server A for the above IP addresses.
6. Enter Server B IP addresses in the **Server B Addresses field**. Server B will be used if there are no matches to the Server A domains.

**NOTE** You can **Clear DNS Cache**. This will erase the domain name to the IP address mapping you had cached for both Server A and B.

## DNS Template

A **Domain Name Server** (DNS) stores the IP addresses with their associated 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.

## DHCP Failover State

EdgeConnect appliances can act as a DHCP server for clients on the LAN side. DHCP failover allows redundancy by creating failover groups when two appliances are combined in an HA configuration. DHCP failover also provides stability if one EdgeConnect dies by allowing the other EdgeConnect HA pair to take over as the DHCP server. To do so, the primary and secondary servers must be completely synchronized so each server can reply on the other if one fails.

This tab displays the DHCP failover peer states of each server for troubleshooting purposes.

### *DHCP Failover Fields*

Field Name	Description
Appliance Name	The name of the Silver Peak appliance that is part of the DHCP failover configuration.
Interface Name	The failover group name that is the same for all the tagged and untagged interfaces corresponding to one physical interface.
My State	The failover endpoint state of the selected primary appliance. The three states are: <b>Normal, Communications-Interrupted, Partner-Down, Recover, Recover-wait, Recover-done.</b>
My State Time	The date and time the selected appliance's DHCP server entered the specified state in the table.
Partner State	The failover endpoint state of the partner appliance. The three states are: <b>Normal, Communications-Interrupted, Partner-Down, Recover, Recover-wait, Recover-done.</b>
Partner State Time	The date and time the partner appliance entered the specified state in the table.
MCLT	The maximum client lead time: the maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner.

## DHCP Failover

Configure the following settings to apply to your DHCP failover servers.

1. Check the DHCP Failover box to enable the DHCP Failover feature.
2. Select whether you are configuring the failover settings for either the Primary or Secondary server.
3. Complete configuring the remaining settings in the table below.

### *DHCP Failover Fields*

Field Name	Description
<b>My IP</b>	The IP address of the LAN interface.
<b>My Port</b>	The port number of the LAN interface.
<b>Peer IP</b>	The IP address of the DHCP peer.
<b>Peer Port</b>	The port number of the DHCP peer.
<b>MLCT</b>	Optional. If selected, the default is 60 minutes. This field cannot be zero.
<b>SPLIT</b>	Optional. If selected, determines which peer (primary/secondary) should process the DHCP requests.
<b>Max Response Delay</b>	Optional. If selected, determines how many seconds the DHCP server may pass without receiving a message from its failover peer before it assumes the connection has failed.
<b>Max Unacked Updates</b>	Tells the remote DHCP server how many BNDUPD messages it can send before it receives a BNDACK from the local system.
<b>Load Balance Max Seconds</b>	Optional. Allows you to configure a cutoff after which load balancing is disabled. The cutoff is based on the number of seconds since the client sent its first DHCPDISCOVER or DHCPREQUEST message, and only works with clients that correctly implement the secs field

## 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.

Templates ×

Template Group ?  
Default Template Group ▾  
+Add -Delete

Active Templates [Show All >](#)

**General Settings**

System  
Logging  
Policies  
Shaper  
Access Lists

Save Save As Cancel  
*Applies to all templates in group*

[Apply Template Groups](#)

**Logging ?**

**Log Configuration**

Minimum severity level Notice ▾  
Start new file when log reaches 50 1-50 MB  
Keep at most log files 30 1-100

**Log Facilities Configuration**

System local1 ▾  
Audit local0 ▾  
Flow local2 ▾

**Remote Log Receivers**

Add

Remote Receiver	Minimum Severity	Facility

## Minimum Severity Levels

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

<b>EMERGENCY</b>	The system is unusable.
<b>ALERT</b>	Includes all alarms the appliance generates: <b>CRITICAL</b> , <b>MAJOR</b> , <b>MINOR</b> , and <b>WARNING</b>
<b>CRITICAL</b>	A critical event
<b>ERROR</b>	An error. This is a non-urgent failure.
<b>WARNING</b>	A warning condition. Indicates an error will occur if action is not taken.
<b>NOTICE</b>	A normal, but significant, condition. No immediate action required.
<b>INFORMATIONAL</b>	Informational. Used by Silver Peak for debugging.
<b>DEBUG</b>	Used by Silver Peak for debugging
<b>NONE</b>	If you select <b>NONE</b> , 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 a web-based configuration interface for Banner Messages. At the top left, there is a tab labeled "Templates" with a close icon. Below this, the "Template Group" section includes a dropdown menu set to "Default Template Group" and a "+Add -Delete" link. A sidebar on the left lists navigation options: "General Settings" (with sub-items "System" and "Banner Messages"), and "Policies" (with sub-items "Shaper" and "Access Lists"). The main content area is titled "Banner Messages" and contains two text input fields: "Your Login Message" and "Your Message of the Day". At the bottom left, there are "Save", "Save As", and "Cancel" buttons, with a note "Applies to all templates in group" and a link "Apply Template Groups".

## HTTPS Certificate Template

The VXOA software includes a self-signed certificate that secures the communication between the user's browser and the appliance. You also have the option to install your own custom certificate, acquired from a CA certificate authority.

The screenshot displays the 'Templates' configuration page for the 'HTTPS Certificate' template. The page is divided into two main sections: 'Template Group' and 'HTTPS Certificate'.

**Template Group:** This section includes a dropdown menu for 'Default Template Group', a '+Add -Delete' link, and a list of 'Active Templates' with a 'Show All >' link. Below this is a 'General Settings' sidebar with links for 'System', 'HTTPS Certificate' (highlighted), 'Policies', 'Shaper', and 'Access Lists'. At the bottom of this section are 'Save', 'Save As', and 'Cancel' buttons, with a note 'Applies to all templates in group' and a link 'Apply Template Groups'.

**HTTPS Certificate:** This section features two radio button options: 'Self Signed Certificate (Issuer: Silver Peak)' (which is selected) and 'Custom Certificate' (with an 'Upload and Replace' button). Below these options are labels for 'Issuer', 'Issued to', and 'Expiration'.

For a custom certificate, to use with a specific appliance:

1. Consult with your IT security team to generate a certificate signing request (CSR), and submit it to your organization's chosen SSL Certificate Authority (CA).

Examples of Certificate Authorities include GoDaddy, Verisign, Comodo, Symantec, Microsoft Entrust, GeoTrust, etc.

- For a list of what Silver Peak supports, see [Silver Peak Security Algorithms](#).
  - All certificate and key files must be in **PEM** format.
2. After the Certificate Authority provides a CA-verified certificate:
    - If your IT security team advises the use of an Intermediate CA, then use an **Intermediate Certificate File**. Otherwise, skip this file.
    - Load the **Certificate File** from the CA.
    - Upload the **Private Key File** that was generated as part of the CSR.
  3. To associate the CA verified certificate for use with Orchestrator, click **Add**.

## User Management 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.

Templates ×

Template Group ?

Default Template Group ▾

+Add -Delete

Active Templates [Show All >](#)

**General Settings**

System

User Management

Policies

Shaper

Access Lists

*Applies to all templates in group*

[Apply Template Groups](#)

**User Management** ?

**User Accounts**

User Name	Capability	Password	Confirm Password	Enabled	
admin	admin	••••	••••	Yes	
monitor	monitor	••••	••••	<input checked="" type="checkbox"/>	

**Password for CLI "Enable" privilege**

Require Password

Password

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.

## Date/Time Template

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

Templates ×

**Template Group** ?

Default Template Group ▼

[+Add](#) [-Delete](#)

Active Templates [Show All >](#)

**General Settings**

[System](#)

[Date/Time](#)

**Policies**

[Shaper](#)

[Access Lists](#)

[Save](#) [Save As](#) [Cancel](#)

*Applies to all templates in group*

[Apply Template Groups](#)

**Date / Time Setting** ?

Time Zone UTC ▼

**Manual**

*Configured when the template group is applied*

**NTP Time Synchronization**

[Add](#)

Server	Version	

- 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**.
  - a. Click **Add**.
  - b. 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.

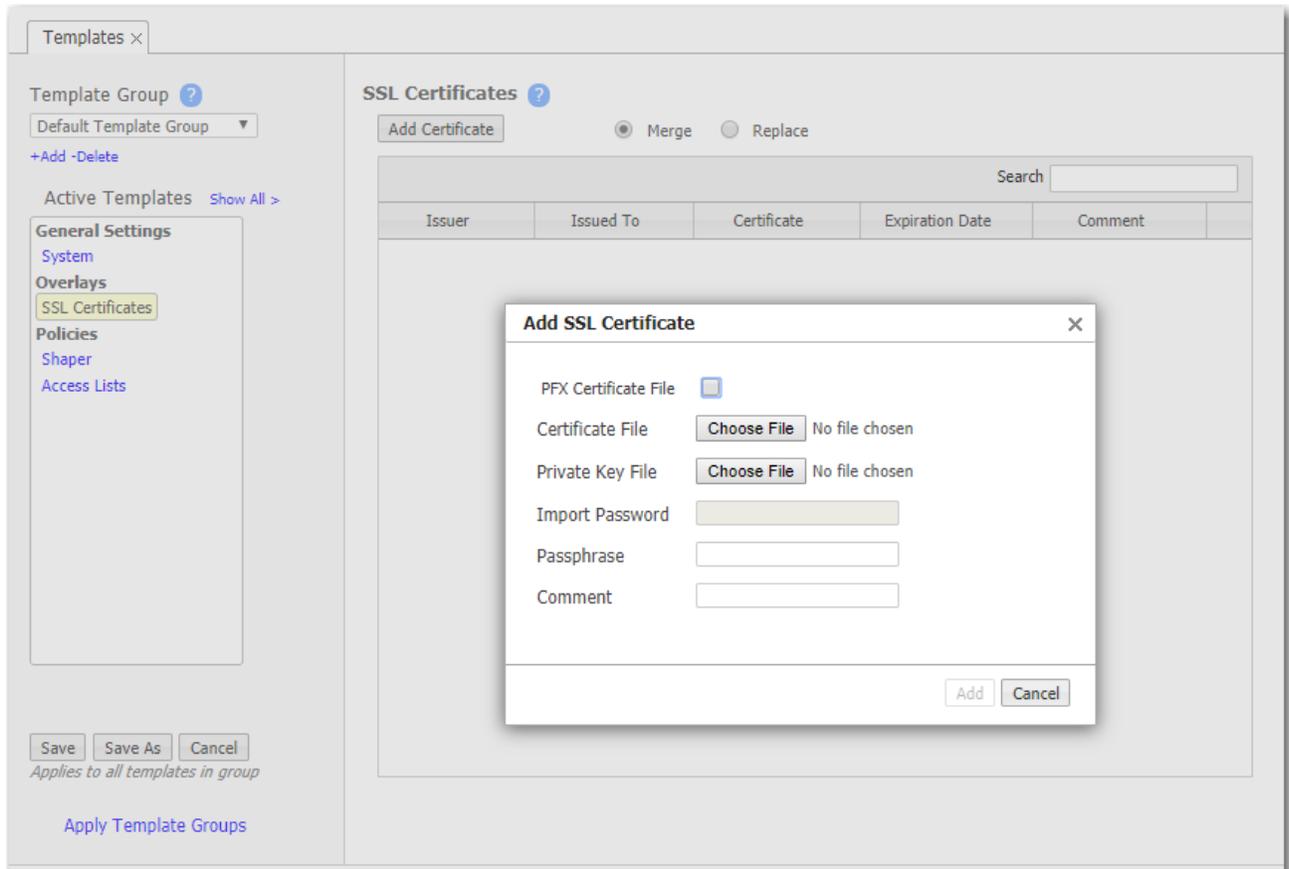
## Data Collection

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

## SSL Certificates Template

Use this page for **SSL Certificates** when the server is *part of your enterprise network* and has its own enterprise SSL certificates and key pairs.

**NOTE** To decrypt SSL for SaaS (cloud-based) services, use the **SSL for SaaS** 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.
- Before installing the certificates, you must do the following:
  - Configure the tunnels bilaterally for **IPSec** (or **IPSec\_UDP**) mode.  
To do so, access the **Configuration - Tunnels** page, select the tunnel, 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 Silver Peak security algorithms, 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.

The screenshot displays the 'SSL CA Certificates' configuration page in the Silver Peak Orchestrator. The main area features a search bar and a table with columns: Issuer, Issued To, Certificate, Expiration Date, and Comment. A modal dialog titled 'Add SSL CA Certificate' is open, containing the following fields:

- PFX Certificate File:
- Certificate File:  No file chosen
- Import Password:
- Comment:

Buttons for 'Add' and 'Cancel' are located at the bottom right of the dialog. The background interface includes a left sidebar with 'Template Group' and 'Active Templates' sections, and a bottom bar with 'Save', 'Save As', and 'Cancel' buttons.



**TIP** For a historical matrix of Silver Peak security algorithms, 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).

The screenshot displays the configuration page for the 'SSL for SaaS' template. On the left, a sidebar shows the navigation menu with 'SSL for SaaS' selected under the 'Overlays' section. The main content area is titled 'SSL Signing Authority for SaaS' and features three radio button options:

- Do Not Decrypt SSL Traffic**  
SSL traffic will be accelerated, but won't be decrypted, so traffic will not be compressed
- Decrypt - Use CA Certificate**  
For SaaS SSL decryption, the appliance will generate substitute certificates, which must be CA-signed (using the Built-In CA Certificate) or signed by a subordinate CA cert issued from an enterprise CA (Custom CA Certificate). To avoid browser warnings, import the applicable CA Certificate into the Trusted CA Store on client browsers.
- Built-In CA Certificate** *(Issuer: Silver Peak)*
- Custom CA Certificate**  
Issuer  
Issued to  
Expiration

At the bottom of the main content area, there are three buttons: 'Upload and Replace', 'Download', and 'Delete'. The left sidebar also includes a 'Template Group' dropdown set to 'Default Template Group', an 'Active Templates' section with a 'Show All >' link, and a 'General Settings' section with links for 'System', 'Overlays', 'Policies', 'Shaper', and 'Access Lists'. At the bottom of the sidebar, there are 'Save', 'Save As', and 'Cancel' buttons, with a note 'Applies to all templates in group' and a link 'Apply Template Groups'.

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 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).



**TIP** For a historical matrix of Silver Peak security algorithms, click [here](#).

---

# Tunnels Template



**NOTE** If you're deploying an SD-WAN network, the Business Intent Overlays (BIOs) govern tunnel properties. In that case, you don't need this template.

If you're not creating overlays, then 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 Groups** tab.
- To **view, edit, and delete** tunnels, use the **Tunnels** tab. The **Mode** selected determines which tabs display.

**Tunnel Settings** x

Settings for Tunnels created by Business Intent Overlays

WAN Interface Labels

- MPLS
- Internet
- LTE

**General** | IPsec

**General**

Mode: IPsec UDP

Auto Max BW Enabled:

Auto Discover MTU Enabled:

MTU: 1600 Bytes

**Packet**

Reorder Wait: 100 ms

FEC: disable

FEC Ratio: 1:10

**Tunnel Health**

Retry Count: 30

DSCP: be

**FastFail Thresholds**

Fastfail Enabled: enable  
(Use "enable" for best performance)

Latency: 0 ms

Loss: 0 %

Jitter: 0 ms

Fastfail Wait-time Base Offset: 150 ms

Fastfail RTT Multiplication Factor: 5

## Tunnels Template Settings

Field Name	Description
Admin State	Indicates whether the tunnel has been set to admin Up or Down.
Auto Discover MTU Enabled	Allows an appliance to determine the best MTU to use.
Auto Max BW Enabled	When enabled, allows the appliances to auto-negotiate the maximum tunnel bandwidth.
DSCP	Determines which DSCP marking the keep-alive messages should use.

## Tunnels Template Settings

Field Name	Description
Fastfail Thresholds	<p>When multiple tunnels are carrying data between two appliances, this feature determines how quickly to disqualify a tunnel from carrying data.</p> <p>The Fastfail connectivity detection algorithm for the wait time from receipt of last packet before declaring a <b>brownout</b> is:</p> $T_{wait} = Base + N * RTT_{avg}$ <p>where <code>Base</code> is a value in milliseconds, and <code>N</code> is the multiplier of the average Round Trip Time over the past minute.</p> <p>For example, if:</p> $Base = 200ms$ $N = 2$ <p>Then,</p> $RTT_{avg} = 50ms$ <p>The appliance declares a tunnel to be in <b>brownout</b> if it doesn't see a reply packet from the remote end within 300ms of receiving the most recent packet.</p> <p>In the Tunnel Advanced Options, <code>Base</code> is expressed as <b>Fastfail Wait-time Base Offset (ms)</b>, and <code>N</code> is expressed as <b>Fastfail RTT Multiplication Factor</b>.</p> <ul style="list-style-type: none"> <li>■ <b>Fastfail Enabled</b> - This option is triggered when a tunnel's keepalive signal doesn't receive a reply. The options are <b>disable</b>, <b>enable</b>, and <b>continuous</b>. If the disqualified tunnel subsequently receives a keepalive reply, its recovery is instantaneous. <ul style="list-style-type: none"> <li>• If set to <b>disable</b>, keepalives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost.</li> <li>• If set to <b>enable</b>, 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.</li> <li>• When set to <b>continuous</b>, keepalives are continuously sent at 10 per second. Therefore, failover occurs after one tenth of a second.</li> </ul> </li> <li>■ Thresholds for <b>Latency</b>, <b>Loss</b>, or <b>Jitter</b> are checked once every second. <ul style="list-style-type: none"> <li>• 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.</li> <li>• Receiving 3 successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.</li> </ul> </li> </ul>

*Tunnels Template Settings*

Field Name	Description
<b>FEC</b>	(Forward Error Correction) can be set to <b>enable</b> , <b>disable</b> , and <b>auto</b> .
<b>FEC Ratio</b>	Is an option when FEC is set to <b>auto</b> , that specifies the maximum ratio. The options are 1:2, 1:5, 1:10, or 1:20.
<b>IPSec Anti-replay window</b>	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.
<b>IPSec Preshared Key</b>	A shared, secret string of Unicode characters that is used for authentication of an IPSec connection between two parties.
<b>Mode</b>	Indicates whether the tunnel protocol is <b>udp</b> , <b>gre</b> , or <b>ipsec</b> .
<b>MTU (bytes)</b>	(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.
<b>Reorder Wait (ms)</b>	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.
<b>Retry Count</b>	The number of failed keep-alive messages that are allowed before the appliance brings the tunnel down.
<b>UDP destination port</b>	Used in UDP mode. Accept the default value unless the port is blocked by a firewall.
<b>UDP flows</b>	The number of flows over which to distribute tunnel data. Accept the default.

## VRRP Template

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

The screenshot shows the configuration interface for a VRRP template. On the left, there is a sidebar with a 'Templates' tab. Under 'Template Group', a dropdown menu is set to 'Default Template Group'. Below this are '+Add -Delete' and 'Active Templates Show All >' options. A list of settings categories is shown: 'General Settings' (with sub-items 'System', 'Networking', 'VRRP', 'Policies', 'Shaper', and 'Access Lists'). At the bottom of the sidebar are 'Save', 'Save As', and 'Cancel' buttons, with the text 'Applies to all templates in group' below them, and a blue link 'Apply Template Groups'.

The main configuration area is titled 'VRRP'. It contains the following settings:

- Admin: Up (dropdown)
- Advertisement Timer:  (1..255)
- Priority:  (1..254)
- Preemption:
- Authentication String:

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.

### *VRRP Template Settings*

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

## Peer Priority Template

When an appliance receives a **Subnet** with the same **Metric** from multiple remote/peer appliances, it uses the Peer Priority list as a tie-breaker.

- If a **Peer Priority** is not configured, then the appliance randomly distributes flows among multiple peers.
- The lower the number, the higher the peer's priority.

The screenshot displays the configuration interface for Peer Priority. On the left, a sidebar shows the navigation menu with categories: General Settings (System), Networking (Peer Priority), and Policies (Shaper, Access Lists). The main content area is titled "Peer Priority" and includes an "Add Peer" button, a search bar, and a table with the following structure:

Peer Name	Priority	
Type to select	0	X

At the bottom of the sidebar, there are "Save", "Save As", and "Cancel" buttons, with the text "Applies to all templates in group" below them. A blue link "Apply Template Groups" is also present.

## Admin Distance Template

This table shows the values associated with various types of **Admin Distance**. Admin Distance (AD) is the route preference value assigned to dynamic routes, static routes, and directly connected routes. When the appliance's Routes table has multiple routes to the same destination, the appliance uses the route with the lowest administrative distance.

### Admin Distance

10 Rows, 1 Selected		Search <input type="text"/>
Type	Distance ▲	
Local	1	
Subnet Shared - Static Routes	10	
Subnet Shared - BGP Remote	15	
BGP Branch (pre-8.1.9.4)	15	
Subnet Shared - OSPF Remote	15	
BGP Transit (pre-8.1.9.4)	20	
EBGP (post-8.1.9.4)	20	
BGP PE (pre-8.1.9.4)	25	
OSPF	110	
IBGP (post-8.1.9.4)	200	

Field	Description
Subnet Shared- Static Routes	A route learned from a Silver Peak peer.
Subnet Shared - OSPF Remote	A route shared from a Silver Peak peer within the same network.
Subnet Shared - BGP Branch	A route shared from a Silver Peak peer from an external network.
OSPF	A route learned from an OSPF (Open Shortest Path First) neighbor.
Local	A manually configured route, or one learned from locally connected subnets.
IBGP (post-8.1.9.4)	Internal BGP: exchanging routing information with a router inside the company-wide network after version 8.1.9.4.
EBGP(post-8.1.9.4)	External BGP: exchanging routing information with a router outside the company-wide network after version 8.1.9.4.
BGP Transit (pre-8.1.9.4)	A type of dynamic route learned from a local BGP branch-transit peer prior to version 8.1.9.4.

Field	Description
BGP PE (pre-8.1.9.4)	A type of dynamic route learned from a local BGP PE (Provider Edge) router prior to version 8.1.9.4.
BGP Branch (pre-8.1.9.4)	A type of dynamic route learned from a local BGP branch peer prior to version 8.1.9.4.

# 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 tunneled 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.
- For minimum and maximum bandwidth, you can configure traffic class values as a percentage of total available system bandwidth and as an absolute value. The appliance always provides the larger of the minimum values, and limits bandwidth to the lower of the maximum values.
- You can rename or edit any traffic class.
- To view any applied configurations, access the **Configuration > Shaper** page.

Templates ×

Template Group ?  
Default Template Group  
+Add -Delete

Active Templates Show All >

General Settings  
System  
Policies  
Shaper  
Access Lists

Save Save As Cancel  
Applies to all templates in group  
Apply Template Groups

Shaper ?

Inbound Outbound Interface Shaper Total Wan Add Interface Shaper Delete Interface Shaper  Enable Interface Shaper  Recalc on IF State Changes

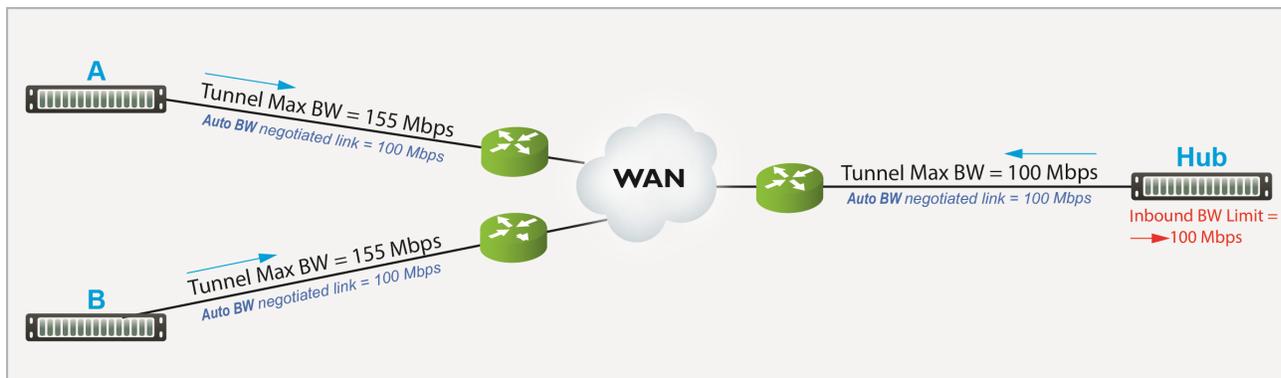
**Total Wan Traffic Classes**

ID	Traffic Name	Priority	Min Bandwidth %	Min Bandwidth Absolute (kbps)	Excess Weighting	Max Bandwidth %	Max Bandwidth Absolute (kbps)	Max Wait Time (ms)	Rate Limit (kbps)
1	Default	1	0	0	250	100	10,000,000	500	0
2	Interactive	1	0	0	1000	100	10,000,000	500	0
3	RealTime	1	0	0	500	100	10,000,000	100	0
4	Replication	1	0	0	100	100	10,000,000	1000	0
5	GuestWireless	1	0	0	100	100	10,000,000	1000	0
6	UNUSED6	6	0	0	1	100	10,000,000	500	0
7	UNUSED7	7	0	0	1	100	10,000,000	500	0
8	UNUSED8	8	0	0	1	100	10,000,000	500	0
9	UNUSED9	9	0	0	1	100	10,000,000	500	0
10	UNUSED10	10	0	0	1	100	10,000,000	500	0

## 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:

- Select **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 bandwidth the appliance can receive.

### Shaper Settings

Field Name	Description
Add Interface Shaper	Adds an interface-specific shaper for outbound or inbound traffic.
Enable Interface Shaper	Enables a separate shaper for a specific WAN interface. <ul style="list-style-type: none"> <li>■ For WAN optimization, the interface shaper can be used but is not recommended.</li> <li>■ For SD-WAN, it should never be used because overlay traffic isn't directed to an interface shaper; traffic is always shaped by the default WAN shaper.</li> </ul>

Field Name	Description
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 <b>Excess Weighting</b> column. Values range from 1 to 10,000.
Interface Shaper	The interface which is being shaped.
Max Bandwidth %	This limits the maximum bandwidth that a traffic class can use to a percentage of total available system bandwidth.
Max Bandwidth Absolute (kbps)	This limits the maximum bandwidth that a traffic class can use to an absolute value (kbps). You can specify a maximum absolute value to cap the bandwidth for downloads and streaming.
Max Wait Time	Any packets waiting longer than the specified <b>Max Wait Time</b> are dropped.
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 <b>Max Bandwidth</b> , then <b>Max</b> overrides <b>Min</b> .
Min Bandwidth Absolute (kbps)	This guarantees a specific level of service when total system bandwidth declines. This is useful for maintaining the quality of VoIP, for example.
Priority	Determines the order in which to allocate each class's minimum bandwidth - <b>1</b> is first, <b>10</b> is last.
Rate Limit (kbps)	You can set per-flow rate limit that a traffic class uses by specifying a number in the Rate Limit column. For no limit, use <b>0</b> (zero).
Recalc on IF State Changes	When an interface state changes to UP or DOWN, selecting this recalculates the total bandwidth based on the configured bandwidth of all UP interfaces. For example, when <b>wan0</b> goes down, <b>wan0</b> bandwidth is removed from the total bandwidth when recalculating.
Traffic Name	The name assigned to a traffic class, either prescriptively or by the user.

## 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**.

The screenshot shows the 'QoS Policies' configuration page. On the left, a sidebar lists navigation options: 'General Settings', 'System', 'Policies', 'Shaper', 'Access Lists', and 'QoS Policies'. The main content area is titled 'QoS Policies' and includes a dropdown for 'map1', a checked 'Active Map' checkbox, and buttons for 'Add Map', 'Delete Map', and 'Rename Map'. Below this are 'Add Rule', 'Merge', and 'Replace' options. A table displays the policy rules:

Priority	Match Criteria	Set Actions			Comment
		Traffic Class	LAN QoS	WAN QoS	
1000	Match Everything	1 - Default	trust-lan	trust-lan	X

### Priority

- With this template, you can create rules with priority from **1000 - 9999**, inclusive. When you apply the template to an appliance, 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.

## Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.
- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## 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.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.

- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

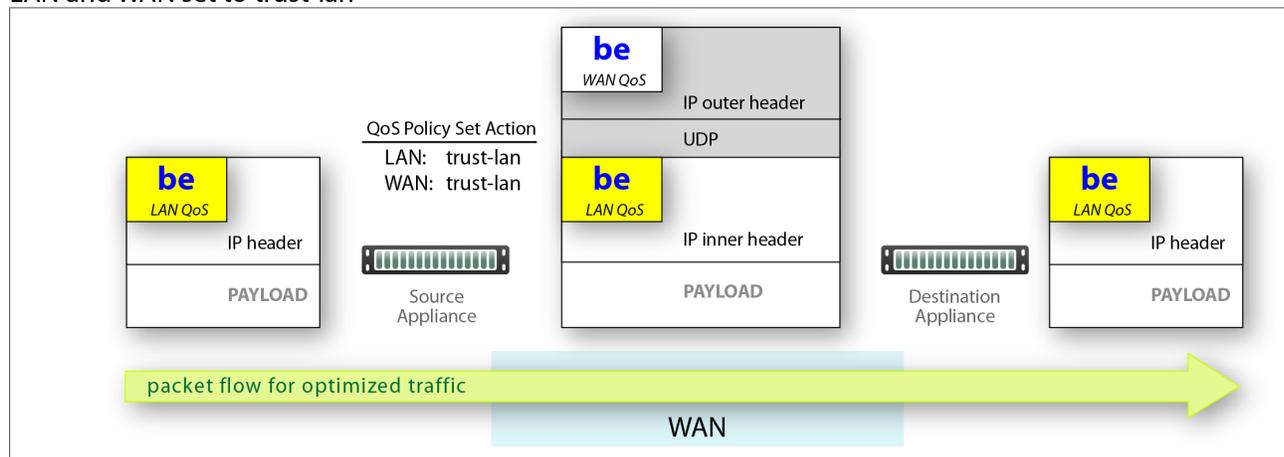
## Handling and Marking DSCP Packets

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

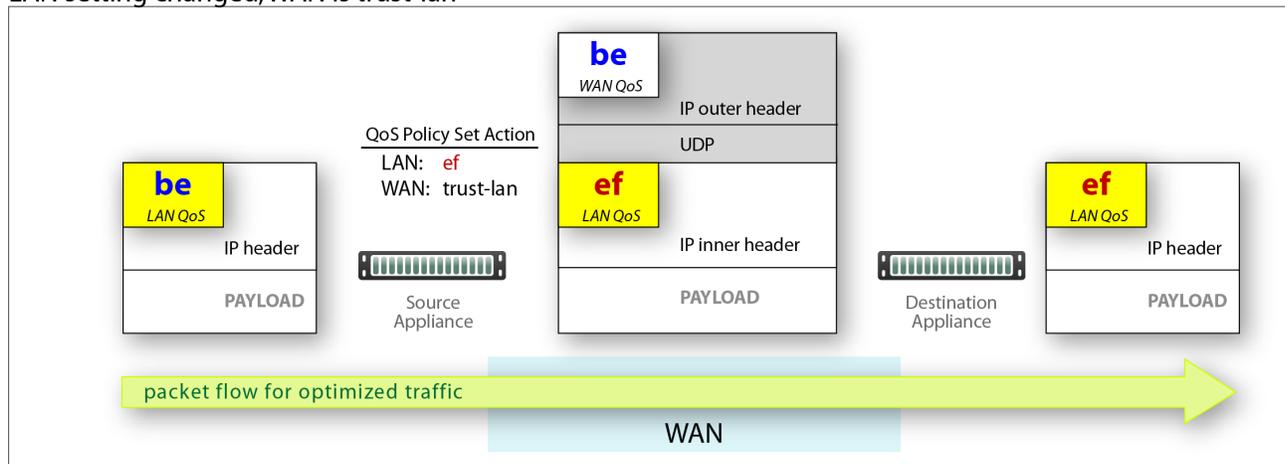
### Applying DSCP Markings to Optimized (Tunnelized) Traffic

- 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.

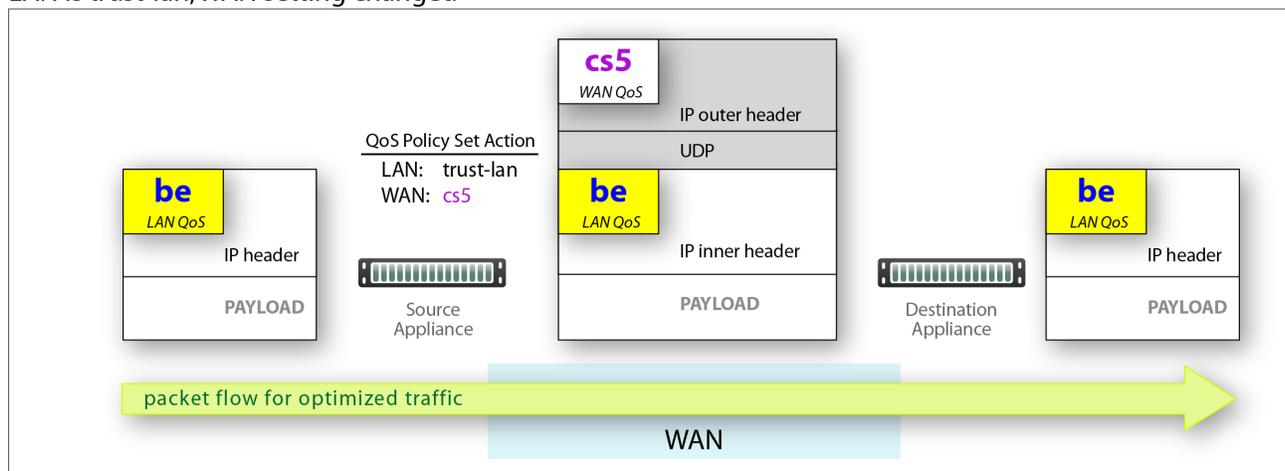
#### LAN and WAN set to trust-lan



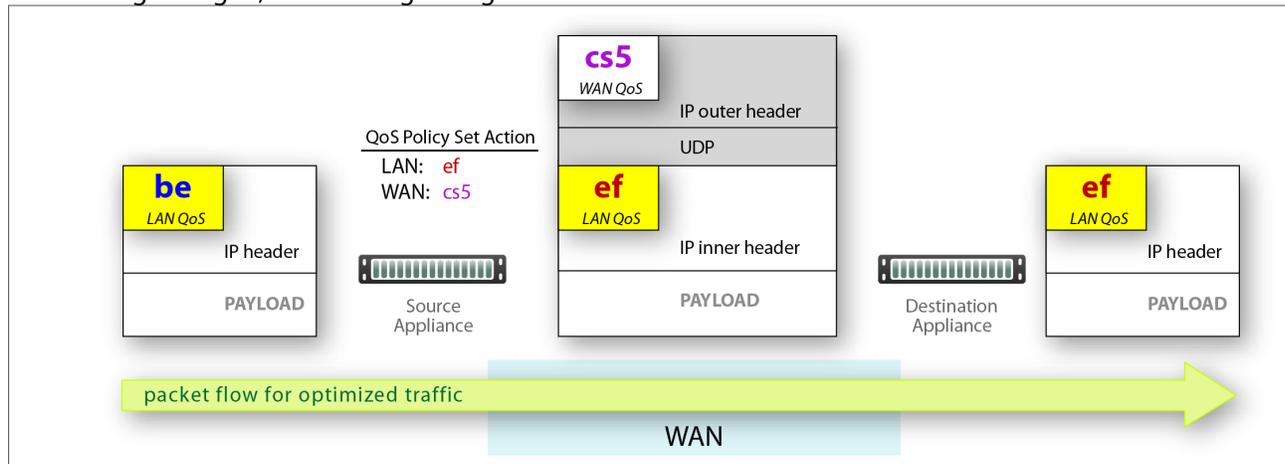
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



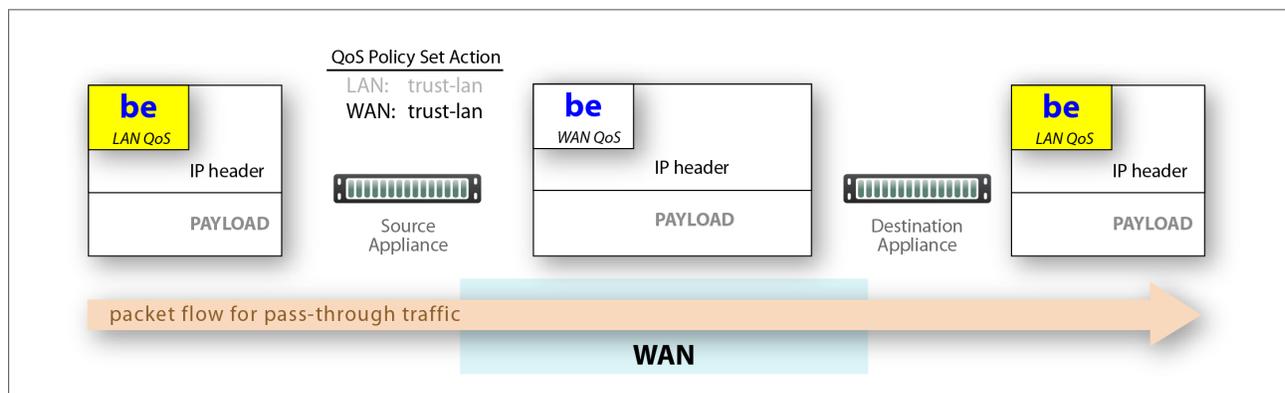
LAN setting changed, WAN setting changed



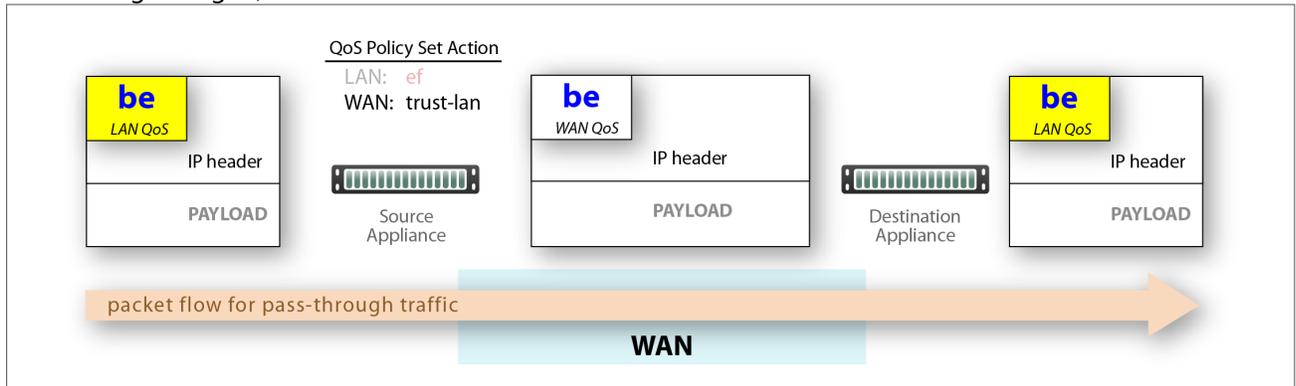
Applying DSCP Markings to Pass-through Traffic

- 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 marking.
  - When the packet reaches the remote appliance, it retains the modified QoS setting as it travels to its destination.

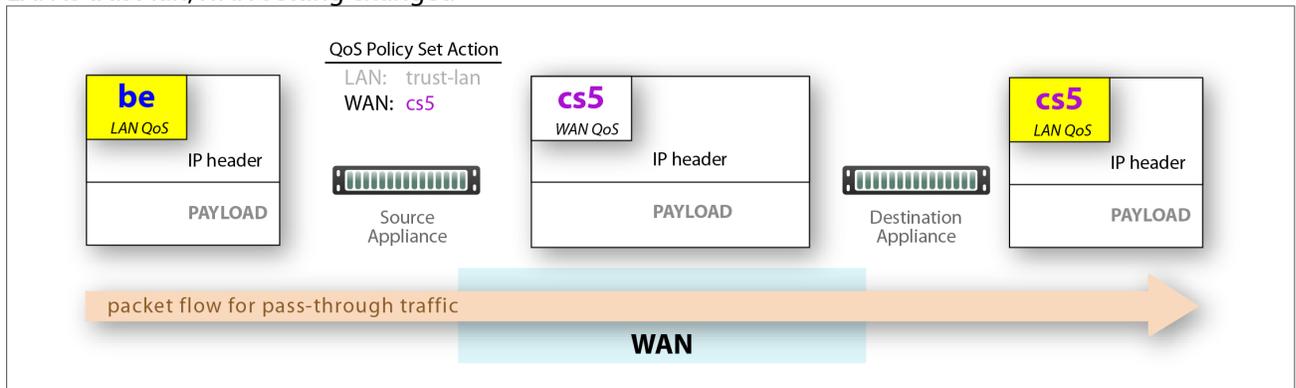
LAN and WAN set to trust-lan



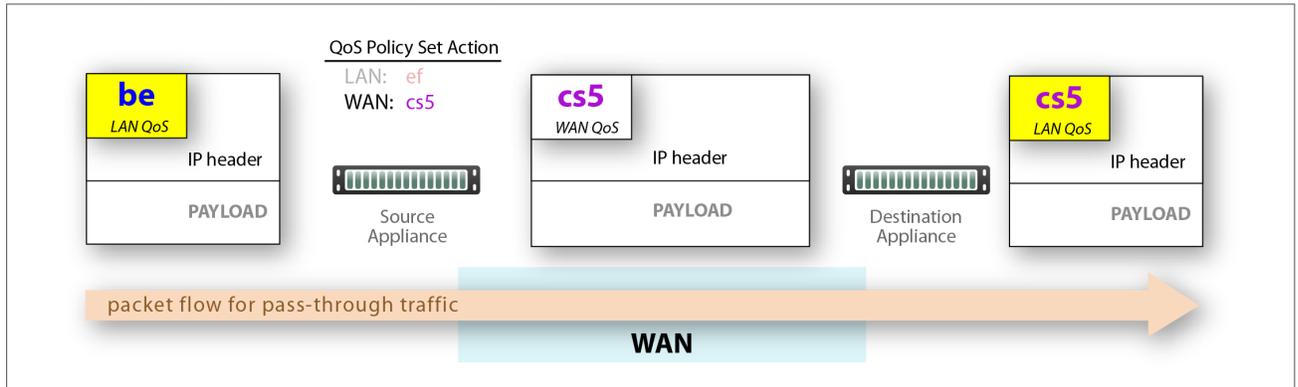
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed

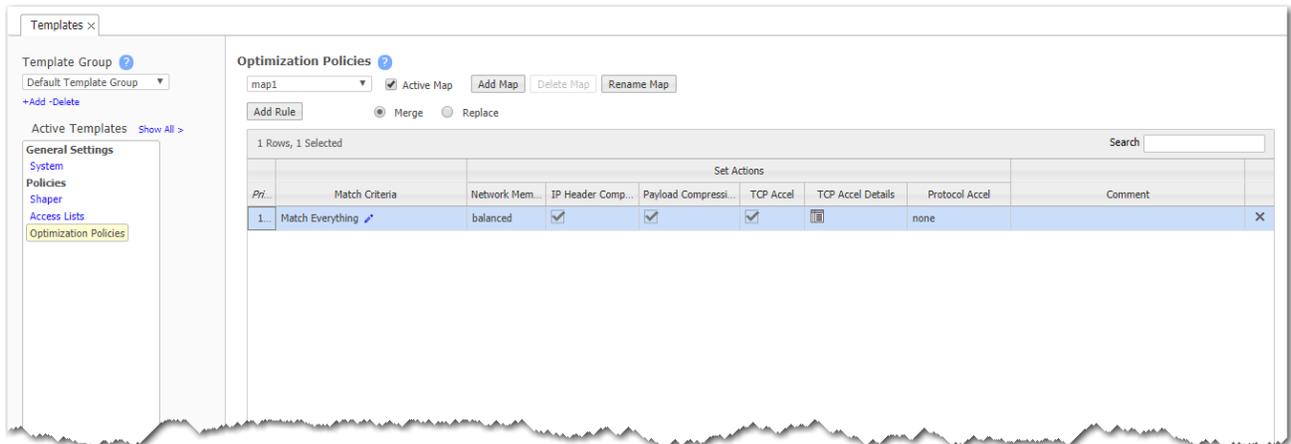


LAN setting changed, WAN setting changed



## 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, 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.

### Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.

- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## 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.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## Set Actions Fields

Set Action	Definition
<b>Network Memory</b>	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.
	<p><b>Maximize Reduction</b>      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.</p>
	<p><b>Minimize Latency</b>      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.</p>
	<p><b>Balanced</b>      Is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types.</p>
	<p><b>Disabled</b>      Turns off Network Memory.</p>
<b>IP Header Compression</b>	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.
<b>Payload Compression</b>	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.
<b>TCP Acceleration</b>	Uses techniques such as selective acknowledgements, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links. For more information, see <a href="#">TCP Acceleration Options</a> .
<b>Protocol Acceleration</b>	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 <b>client</b> ) determines the state of the protocol-specific optimization.

## Route Policies Template



If you've deployed an SD-WAN network by using Business Intent Overlays (BIO), then Orchestrator uses BIOs to automatically create the necessary Route Policies.

If you're creating a conventional WAN optimization network, then there may be occasions when you need to directly configure Route Policies. Then, the following applies.

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

The screenshot shows the 'Route Policies' configuration page in the Orchestrator. The left sidebar contains a 'Templates' section with 'Route Policies' selected. The main panel shows the 'Route Policies' configuration for a map named 'map1'. The 'Active Map' checkbox is checked. Below the map name are buttons for 'Add Map', 'Delete Map', and 'Rename Map'. There are also buttons for 'Add Rule', 'Merge', and 'Replace'. The main area displays a table with 4 rows of policy rules. The table has columns for Priority, Match Criteria, Destination, Path, Fallback, and Comment. Each row has a delete icon (X) in the rightmost column.

Priority	Match Criteria	Set Actions			Comment	
		Destination	Path	Fallback		
1000	Protocol ip, DSCP ef	auto optimized	low-latency	pass-through		X
1010	Protocol ip, Application datadomain	auto optimized	load balance	pass-through		X
1020	Protocol ip, Dest IP/Subnet 10.10.11.56/32	auto optimized	low-loss	pass-through		X
1030	Protocol ip	auto optimized	load balance	pass-through		X

## 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, 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.

## Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.
- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

## 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.

## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## Set Actions Fields

The Route Policy template's SET actions determines where to direct traffic and what the fallback is when a tunnel is down.

## Where the appliance directs traffic

- In the **Destination** field, you specify how to characterize the flow. The options are a specific overlay, **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
- 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

- The **Fallback** can be **pass-through** [shaped], **pass-through-unshaped**, or **dropped**.
- When configuring the Route Policy for an **individual** appliance, the **continue** option is available if a specific tunnel is named in the **Destination** 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](#).

## NAT Policies Template

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

The screenshot shows the 'NAT Policies' configuration page. On the left, there is a sidebar with 'Template Group' set to 'Default Template Group' and a list of active templates including 'General Settings', 'System', 'Policies', 'Shaper', 'Access Lists', and 'NAT Policies'. The main content area is titled 'NAT Policies' and includes options for 'NAT All Inbound' and 'NAT All Outbound', each with checkboxes for 'NAT IP auto' and 'Fallback'. Under 'Advanced Settings', there is a dropdown for 'map1', a checked 'Active Map' checkbox, and buttons for 'Add Map', 'Delete Map', and 'Rename Map'. Below this is an 'Add Rule' button and radio buttons for 'Merge' and 'Replace'. A table with 1 row is displayed, showing a policy with priority 1000, match criteria 'Protocol ip', NAT Type 'no-nat', NAT Direction 'none', NAT IP 'auto', and a 'Fallback' checkbox. The table has a search bar and a close button (X) in the top right corner.

Priority	Match Criteria	Set Actions				Comment	
		NAT Type	NAT Direction	NAT IP	Fallback		
1000	Protocol ip	no-nat	none	auto	<input type="checkbox"/>		X

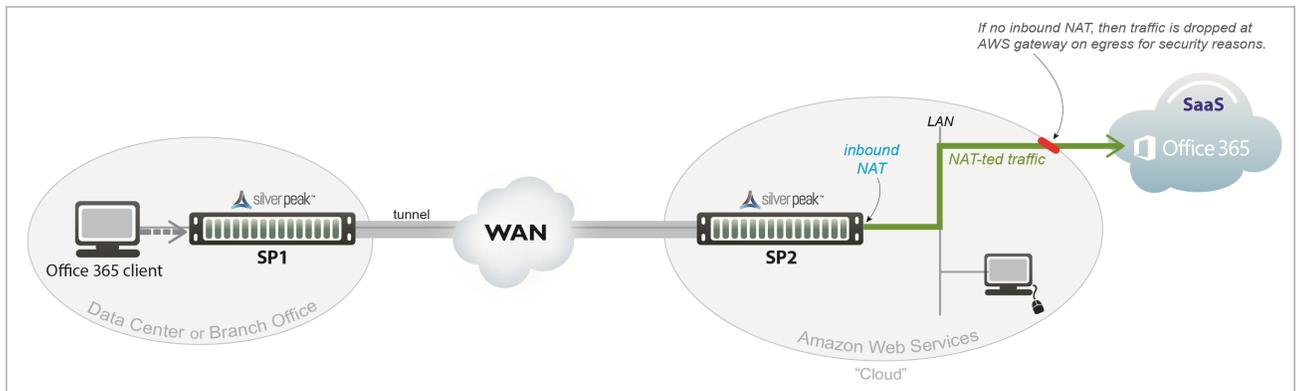
### When to NAT

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

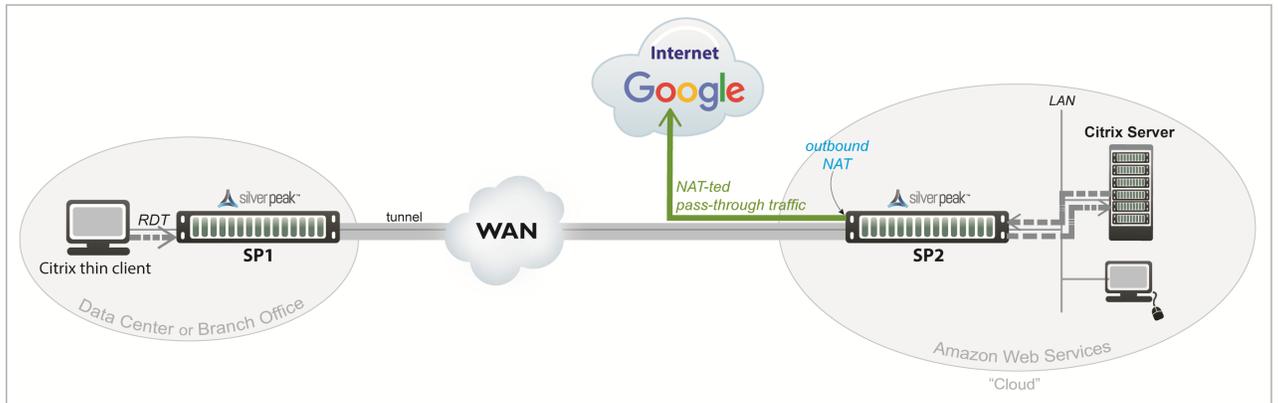
that traffic originated.

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**:

### Match Criteria

- These are universal across all policy maps – **Route**, **QoS**, **Optimization**, **NAT** (Network Address Translation), and **Security**.
- If you expect to use the same match criteria in different maps, you can create an **ACL** (Access Control List), which is a named, reusable set of rules. For efficiency, create them in **Configuration > Templates > Access Lists**, and apply them across appliances.

- The available parameters are **Application**, **Address Map** (for sorting by country, IP address owner, or SaaS application), **Domain**, **Geo Location**, **Interface**, **Protocol**, **DSCP**, **IP/Subnet**, **Port**, and **Traffic Behavior**.
- To specify different criteria for inbound versus outbound traffic, select the **Source:Dest** checkbox.

### 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.

### Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137.\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139.\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

### Set Actions

Set Action	Option	Definition
NAT Type	no-nat	Is the <i>default</i> . 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.

Set Action	Option	Definition
	<b>outbound</b>	NAT is on the WAN interface.
	<b>none</b>	The only option if the NAT Type is <b>no-nat</b> .
<b>NAT IP</b>	<b>auto</b>	Select if you want to NAT <b>all</b> traffic. The appliance then picks the first available NAT IP/Port.
	<b>tunnel</b>	Select if you only want to NAT <b>tunnel</b> traffic. Applicable only for inbound NAT, as outbound doesn't support NAT on tunnel traffic.
	<b>[IP address]</b>	Select if you want to make NAT use this IP address during address translation.
<b>Fallback</b>		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.

### Merge / Replace

At the top of the page, choose

**Merge** to use the values in the template, but keep any values set on the appliance as is (producing a mix of template and appliance rules),

*-OR-*

**Replace** (recommended) to replace all values with those in the template.

## Threshold Crossing Alerts Template

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

Templates ×

Template Group ?

Default Template Group ▾

+Add -Delete

Active Templates [Show All >](#)

**General Settings**

- System
- Policies
- Shaper
- Access Lists
- Threshold Crossing Alerts

Applies to all templates in group

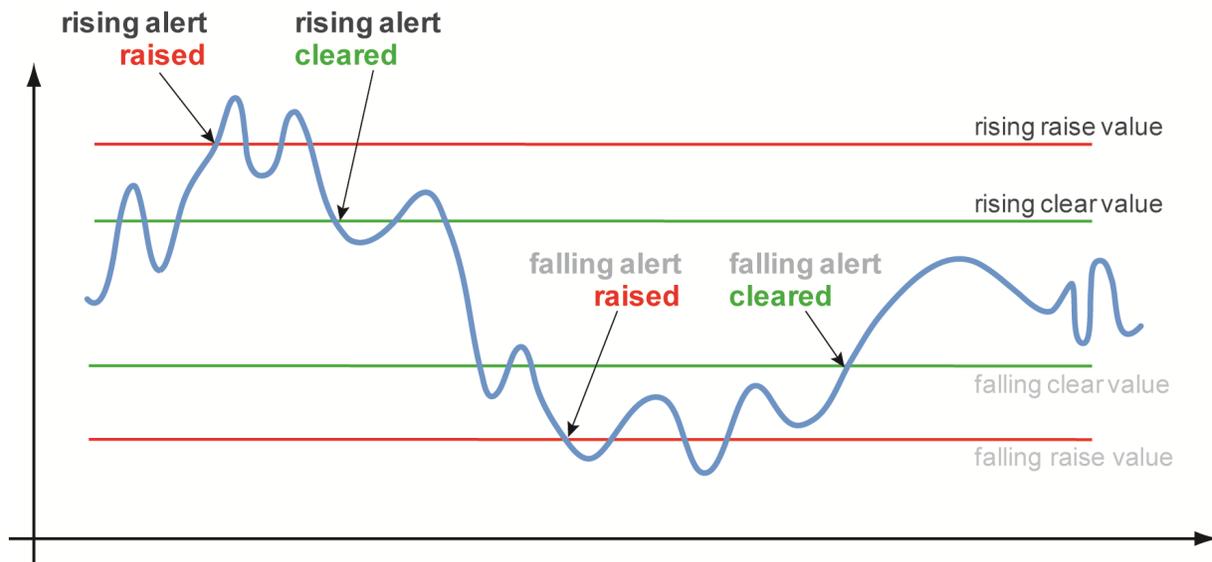
Apply Template Groups

### Threshold Crossing Alerts ?

12 Rows Search

Name ▲	Rising				Falling			
	Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
File-system utilization	90%	85%	5	<input checked="" type="checkbox"/>	75%	75%	5	<input type="checkbox"/>
LAN-side receive throughput	1000000 kbps	1000000 kbps	5	<input type="checkbox"/>	0 kbps	0 kbps	5	<input type="checkbox"/>
Total number of flows	90%	85%	5	<input checked="" type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Total number of optimized flows	90%	85%	5	<input type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Tunnel OOP post-POC	100%	100%	5	<input type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Tunnel OOP pre-POC	100%	100%	5	<input checked="" type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Tunnel latency	1000 ms	850 ms	5	<input checked="" type="checkbox"/>	0 ms	0 ms	5	<input type="checkbox"/>
Tunnel loss post-FEC	100%	100%	5	<input type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Tunnel loss pre-FEC	100%	100%	5	<input type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Tunnel reduction	100%	100%	5	<input type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
Tunnel utilization	100%	100%	5	<input type="checkbox"/>	0%	0%	5	<input type="checkbox"/>
WAN-side transmit throughput	1000000 kbps	1000000 kbps	5	<input type="checkbox"/>	0 kbps	0 kbps	5	<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

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

- **TCA's based on an end-of-minute count:**
  - Total number of flows
  - Total number of optimized flows
- **TCA's 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)

## TCA Metrics

**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:

### *Metrics for Threshold Crossing Alerts*

TCA Name	Unit	Metric
<b>Appliance Level</b>		
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
<b>Tunnel Level</b>		

TCA Name	Unit	Metric
Tunnel latency	msec	Second-sampled maximum latency during the minute
Tunnel loss pre-FEC	1/10 <sup>th</sup> %	Minute average
Tunnel loss post-FEC	1/10 <sup>th</sup> %	Minute average
Tunnel OOP pre-POC	1/10 <sup>th</sup> %	Minute average
Tunnel OOP post-POC	1/10 <sup>th</sup> %	Minute average
Tunnel utilization	% of configured bandwidth	Minute average
Tunnel reduction	%	Minute average

## 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, there's a 'Template Group' dropdown set to 'trial' and a list of policies including 'SaaS Optimization'. The main configuration area includes:

- Enable SaaS Optimization:**
- RTT Calculation Interval:** 1440 (1..1440) minutes
- RTT Ping Interface:** default

Below the settings is a table with 52 rows of SaaS applications. The table has columns: Application Name, Opti..., Adver..., RTT Threshold, Domains, and SaaS ...

Application Name	Opti...	Adver...	RTT Threshold	Domains	SaaS ...
Adobe	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	adobe.com	1
AirWatch	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.air-watch.com	31
AthenaHealth	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.athenahealth.com, athenahealth.com	34
BlueJeans	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.bluejeans.com, *.bjn.vc	69
Box	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.box.com, *.app.box.com, *.boxcloud.com, *.box.net, *.boxcdn.net	2
CCConc	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.myccportal.com, myccportal.com	30
ConstantContact	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	constantcontact.com	3
CornerstoneOnDemand	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	cornerstoneondemand.com	4
Dropbox	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	dropbox.com, *.dropbox.com, *.dl.dropboxusercontent.com	5
Dynamics	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.dynamics.com, *.microsoft.com, dynamics.com, microsoft.com	75
Eloqua	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	eloquatrainingcenter.com, eloqua.com	6
GoToAssist	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com	7
GoToMeeting	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotomeeting.com	8
GoToTraining	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com	9
GoToWebinar	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotoassist.com, gotowebinar.com	10
Intuit	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	intuit.com	11
Jobvite	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	www.jobvite.com, hire.jobvite.com, careers.jobvite.com	12
Lithium	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	lithium.com	13

**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 tunneled.
  - 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

The **RTT Calculation Interval** specifies how frequently Orchestrator recalculates the Round Trip Time for the enabled Cloud applications.

The **RTT Ping Interface** specifies which interface to use to ping the enabled SaaS subnets for Round Trip Times. The **default** interface is **wan0**.

## 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.

## Security Policies Template

Use this page to set up security policies, also known as *zone-based firewalls*.

Zones are created on the Orchestrator and applied to an **Interface**.

By default, traffic is allowed between interfaces labeled with the same zone. Any traffic between interfaces with different zones is dropped. Users can create exception rules (Security Policies) to allow traffic between interfaces with different zones.

When you create an interface, it is assigned **Default** zone.

If you create a new zone and assign that to an interface, all traffic between that interface and rest of the interfaces (which are still in the **Default** zone) are dropped. This implies that zone creation and assignment to interfaces should be performed during a planned network maintenance.

You can also assign a zone label to an **Overlay**. On a brand new system, all overlays are assigned the **Default** zone.

Traffic between an Interface and an Overlay follows the same rules as traffic between Interfaces or two Overlays; traffic is allowed between zones with the same label, and any traffic between different zones is dropped. Users can create Security Policies to allow traffic between different zones.

### Implicit Drop Logging

Implicit Drop Logging allows you to configure implicit zone-based firewall drop logging levels. Implicit zone-based firewall drop is for inter-zone traffic by default. For example, if all the zone\_x to zone\_y traffic is the default **Deny All** (all the red cells from matrix), the traffic will be dropped by the zone-based firewall engine.

Select one of the following levels for the Implicit Drop Logging from the list: **None**, **Emergency**, **Alert**, **Critical**, **Error**, **Warning**, **Notice**, **Info**, or **Debug**.

**NOTE** The default logging level is **Alert**.

### Template

Complete the following steps to create a Security Policies Template:

1. Create zone names in **Configuration > Overlays > Firewall Zones**.
2. Create security policies to define exceptions.

To edit or add a rule, select the desired square in the matrix, and when the Edit Rules pop-up appears, make the desired changes.

3. Select the edit icon in the Match Criteria column and the Match Criteria pop-up appears. Make the desired changes.

You can select **More Options** to customize your rules. Check the box next to the specific match criteria and select your desired changes from the list.

4. Select **Save**.

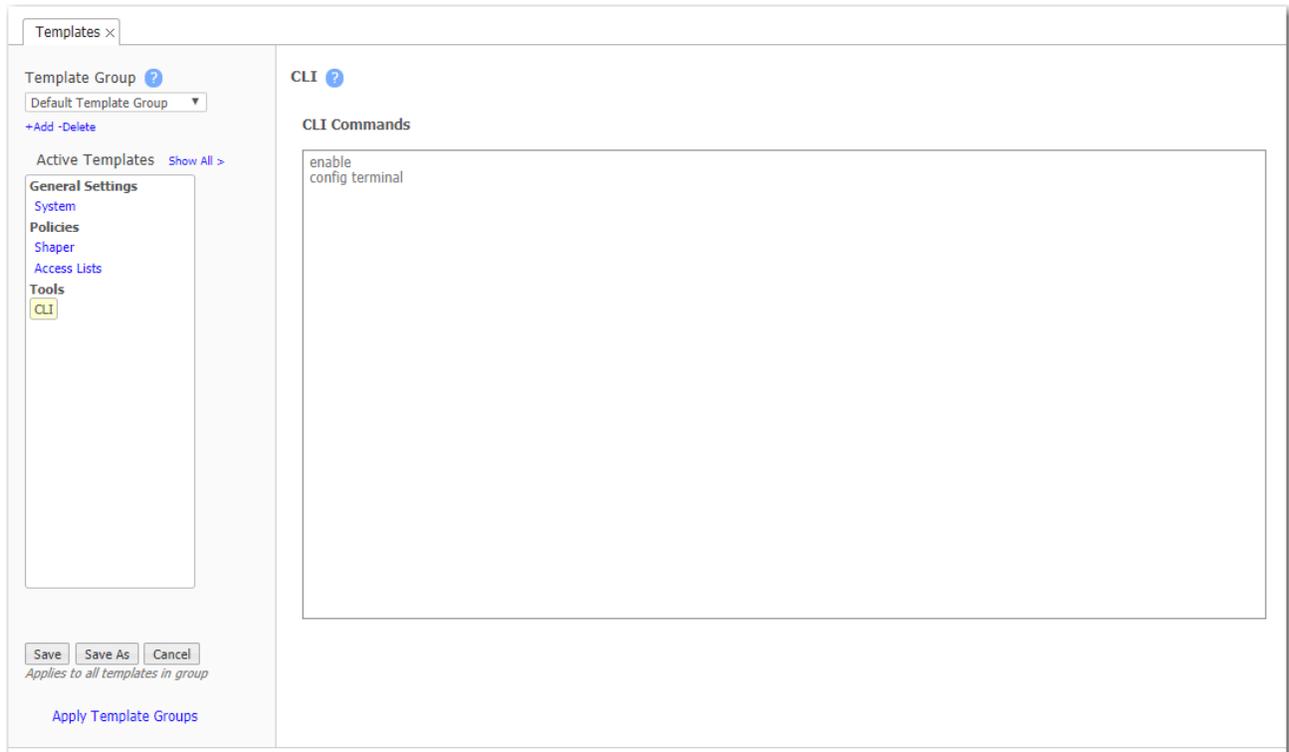
## Wildcard-based Prefix Matching

- When using a range or a wildcard, the IPv4 address must be specified in the 4-octet format, separated by the dot notation. For example, **A.B.C.D**.
- Range is specified using a dash. For example, **128-129**.
- Wildcard is specified as an asterisk (\*).
- Range and Wildcard can both be used in the same address, but an octet can only contain one or the other. For example, **10.136-137\*.64-95**.
- A wildcard can only be used to define an entire octet. For example, **10.13\*.\*.64-95** is not supported. The correct way to specify this range is **10.130-139\*.64-94**.
- The same rules apply to IPv6 addressing.
- CIDR notation and (Range or Wildcard) are mutually exclusive in the same address. For example, use either **192.168.0.0/24** or **192.168.0.1-127**.
- These prefix-matching rules only apply to the following policies: Router, QoS, Optimization, NAT, Security, and ACLs.

## CLI Template

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

Enter each CLI command on a new line.



## Session Management Template

Use this page to configure access to the web server.

The screenshot shows the 'Session Management' configuration page within a 'Templates' window. The left sidebar contains a navigation menu with categories: General Settings, System, Policies, Shaper, Access Lists, Tools, and Session Management (highlighted). The main content area is titled 'Session Management' and includes the following settings:

- Auto Logout:** A text input field containing '15', with a note '(0-60 minutes, 0 indicates no timeout)'.
- Max Session:** A text input field containing '10', with a note '(5-50)'.
- Web Protocol:** Three radio buttons labeled 'HTTP', 'HTTPS', and 'Both'. The 'Both' option is selected.

At the bottom of the left sidebar, there are three buttons: 'Save', 'Save As', and 'Cancel'. Below these buttons is the text 'Applies to all templates in group' and a link 'Apply Template Groups'.

- **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 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.

## Apply Template Groups

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

The screenshot shows the 'Apply Template Groups' page. On the left, there is a 'Template Apply Order' section with a list of template groups: Default Template Group, Demo, North America, NTP, Security1, and trial. Each group has two checkboxes, one under an 'Add' button and one under a 'Remove' button. Below this list are 'Apply' and 'Cancel' buttons. On the right, there is a table with 30 rows. The table has columns for Hostname, Present, and Changes. The 'Present' column shows 'NTP,Security1' for all listed appliances.

Hostname	Present	Changes
Chennai	NTP,Security1	
Mumbai	NTP,Security1	
Osaka	NTP,Security1	
Seoul	NTP,Security1	
Singapore	NTP,Security1	
Tokyo	NTP,Security1	
Barcelona	NTP,Security1	
Edinburgh	NTP,Security1	
Frankfurt	NTP,Security1	
Geneva	NTP,Security1	

- If multiple template groups are applied to an appliance, the order in which they're applied determines which template 'wins'. Templates applied later (lower on the apply order list) will overwrite any conflicting templates applied earlier.
- Drag templates up or down to reorder the list.
- Orchestrator automatically applies any changed templates to the associated appliances.

# Monitoring Status and Performance

These topics focus on reports related to performance, traffic, and appliance status.

Also helpful in monitoring, and the [Threshold Crossing Alerts Tab](#) are addressed in other chapters.

# Dashboard

*Monitoring > [Summary] Dashboard*

The **Dashboard** is a customizable collection of widgets for monitoring your network. Customizations persist for each user account.

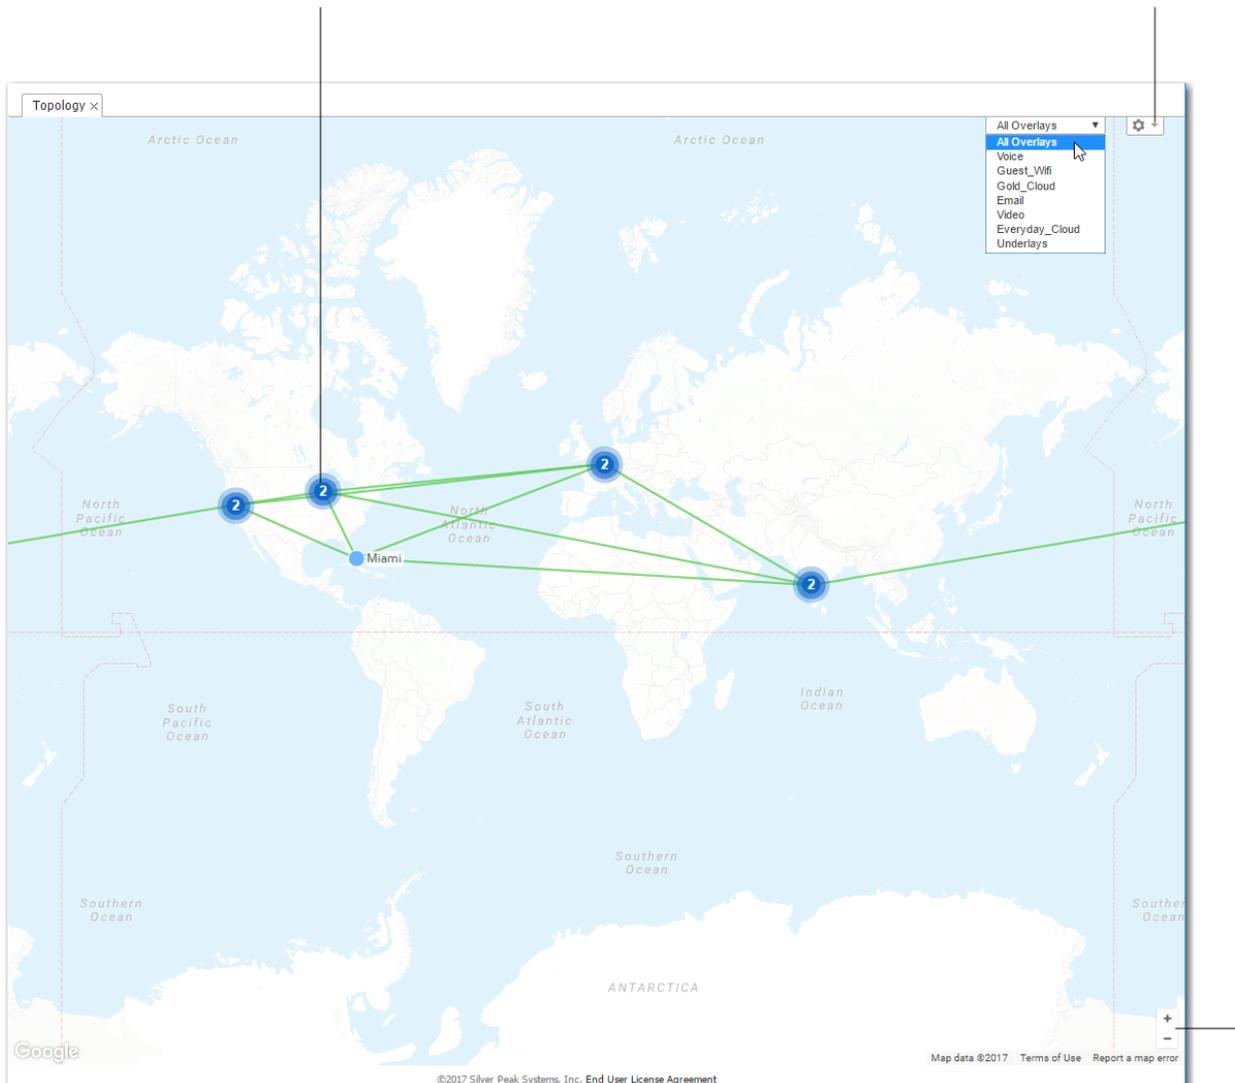
- Select which widgets to show or hide in **Settings** [  ].
- To move widgets, drag them by title.
- To access more detail in its corresponding tab, click a widget's title.
- To filter on **Src** or **Dest**, alone, click to select (illuminate) the desired filter.
- The **Appliance Summary** displays an inventory count by appliance model. It also displays license type, availability, and usage.
- The **Health Map** widget reflects the selections and settings configured on the Health Map tab.

## Topology Settings & Legend

*Monitoring > [Summary] Topology*

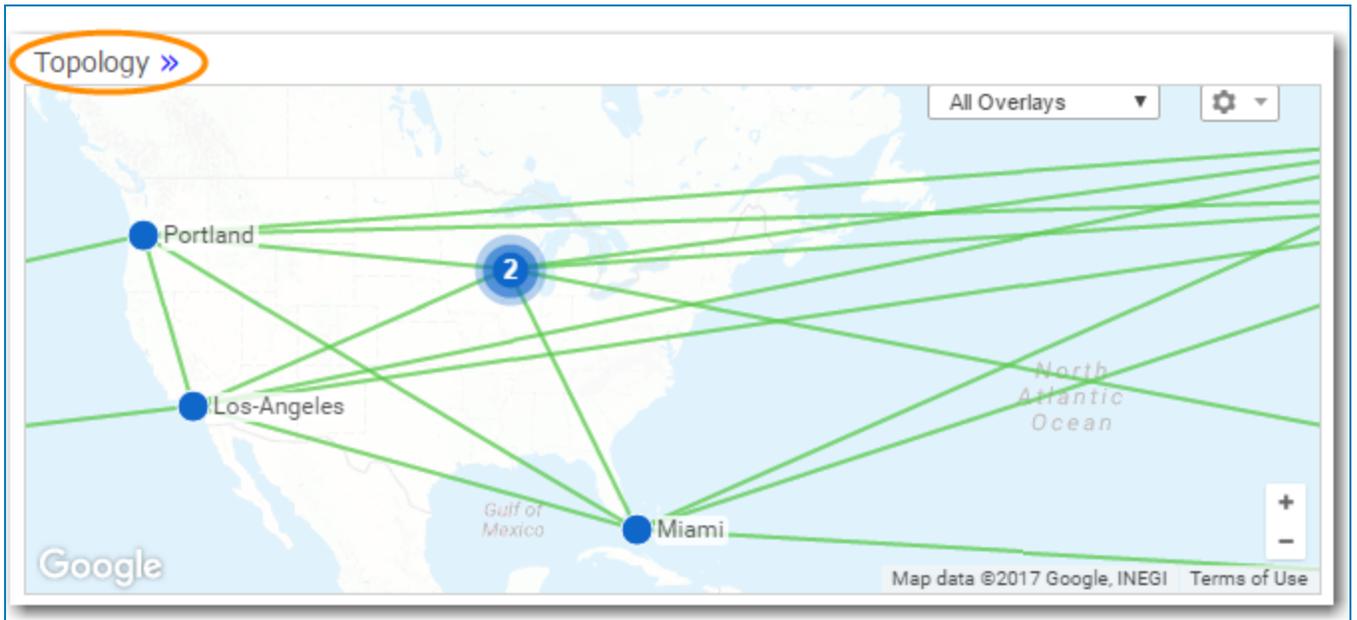
The **Topology** tab provides a visual summary of your Silver Peak network.

When configuring a software-defined WAN (**SD-WAN**), you can view **All Overlays**, individual **Business Intent Overlays** (BIOs), or the single and bonded **Underlay** tunnels that support them.

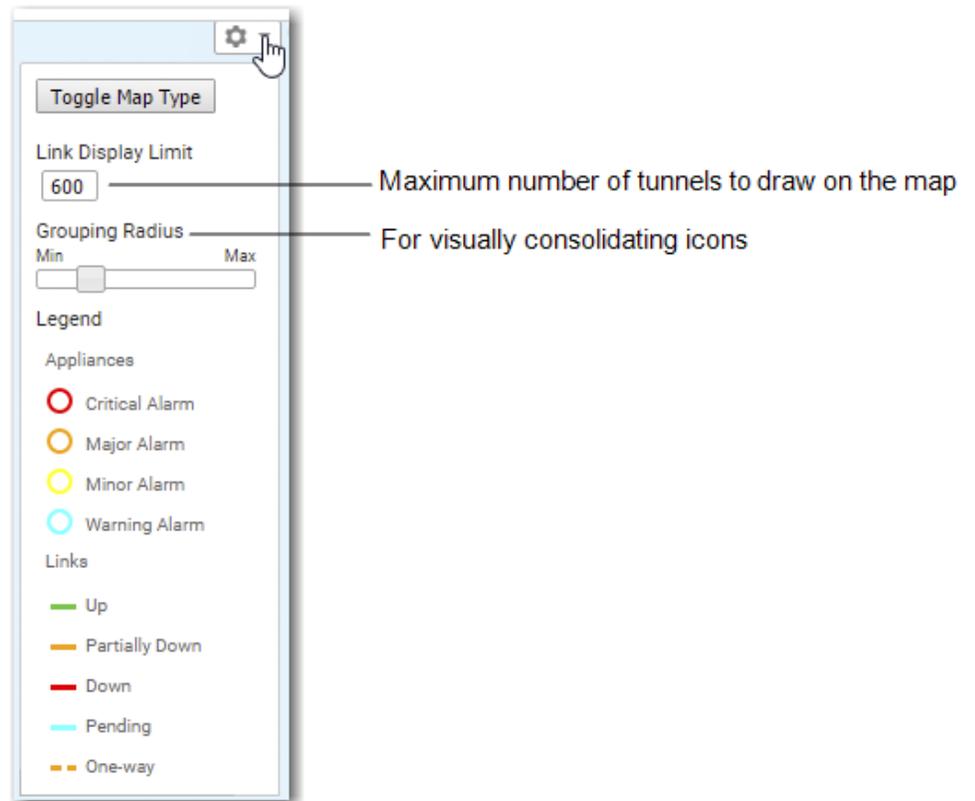


You can access it under **Monitoring** in the menu bar, or by clicking the widget title on the **Dashboard** tab.

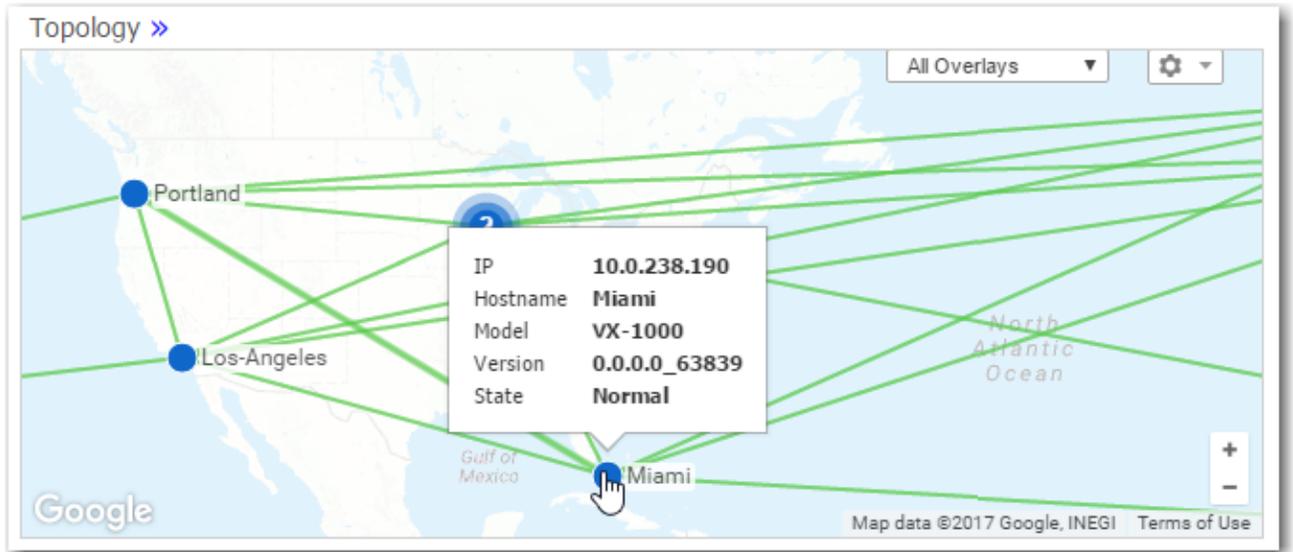
Topology widget on Dashboard tab



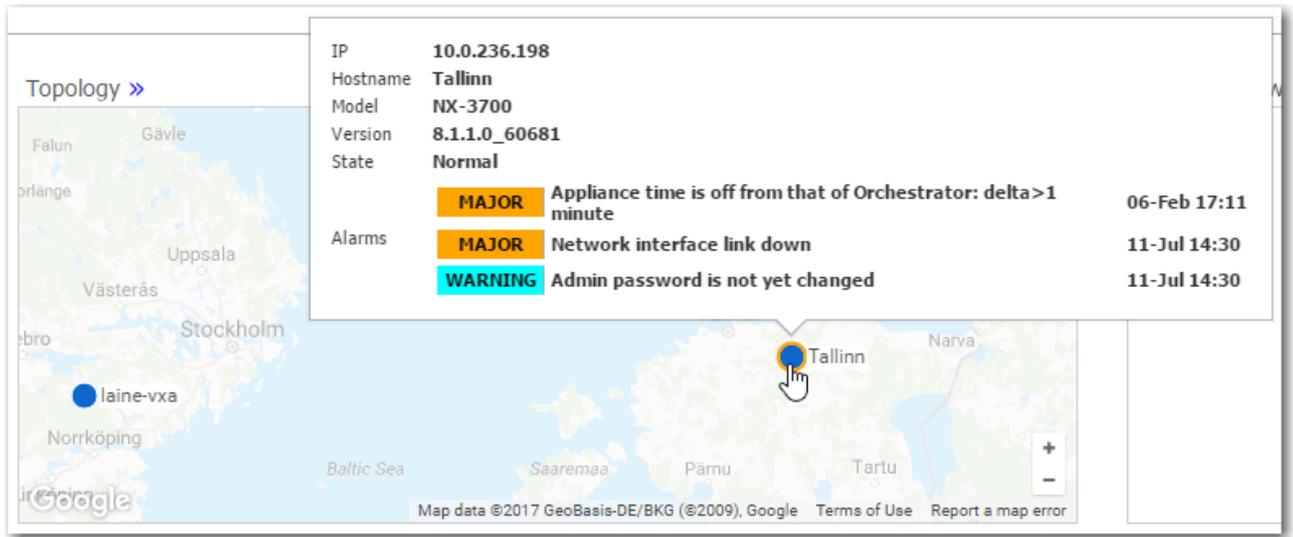
- The Legend details the appliances' management and operational states.



- The **Topology** map can dynamically geolocate an appliance when you enter a location [City, State, Country] in an appliance Configuration Wizard, or when you modify the appliance by right-clicking to access its contextual menu.
- The map tile renders to support variable detail at different zoom levels.
- You can use icon grouping to visually consolidate adjacent appliances. The status bubbles up, and you can configure relative grouping distance in the map's legend. The grouping is also a function of how far you zoom in or out.
- Rolling over an individual appliance's icon displays basic system information.

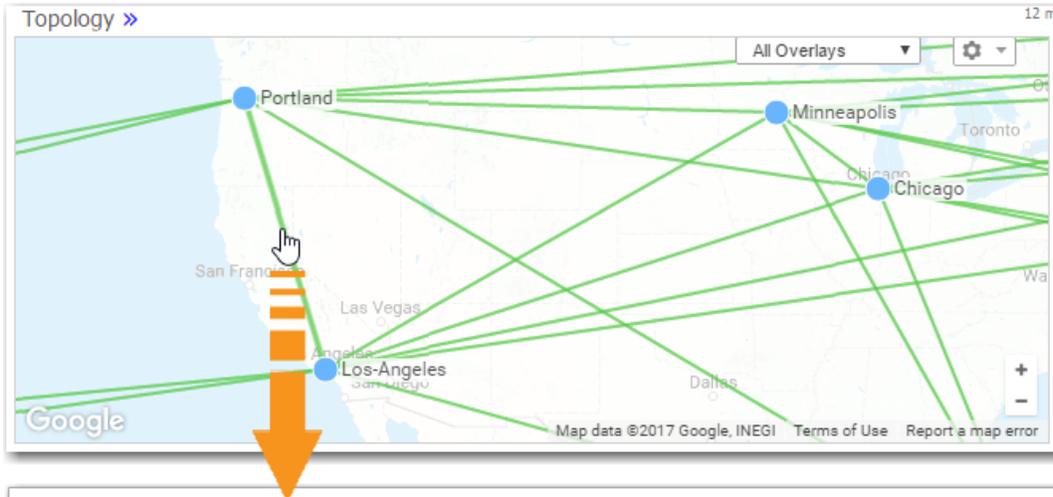


When the icon is encircled by a ring, indicating an alarm, those also display.



## Viewing Tunnels in the Topology Map

Clicking on a tunnel opens a table with access to information about that link.



**Tunnels** X

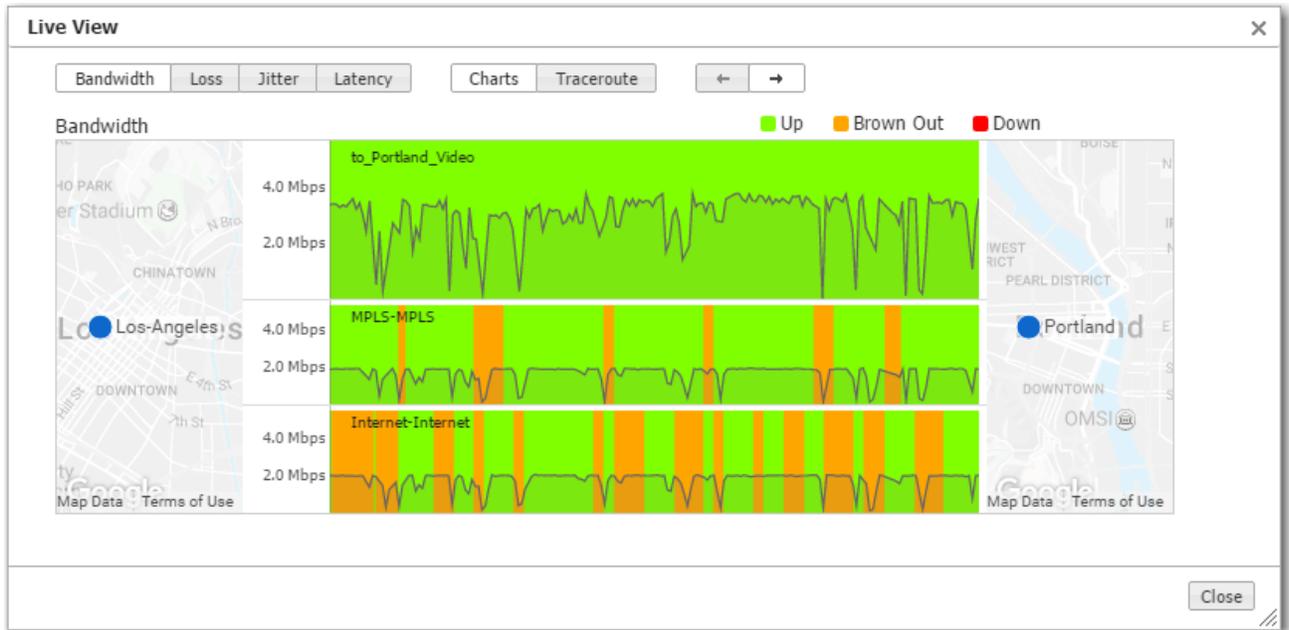
12/36 Rows, 1 Selected Search

Type	Local Appliance	Remote Appliance	Name	Status	Live View	Historical Charts
+ Overlay: Gold_Cloud (6 tunnels)						
+ Overlay: Everyday_Cloud (6 tunnels)						
+ Overlay: Voice (6 tunnels)						
- Overlay: Video (6 tunnels)						
overlay	Los-Angeles	Portland	to_Portland_Video	up - active		
underlay	Los-Angeles	Portland	to_Portland_MPLS-MPLS	up - active		
underlay	Los-Angeles	Portland	to_Portland_Internet-Internet	up - active		
underlay	Portland	Los-Angeles	to_Los-Angeles_Internet-Internet	up - active		
underlay	Portland	Los-Angeles	to_Los-Angeles_MPLS-MPLS	up - active		
overlay	Portland	Los-Angeles	to_Los-Angeles_Video	up - active		
+ Overlay: Email (6 tunnels)						
+ Overlay: Guest_Wifi (6 tunnels)						

Close

## Live View

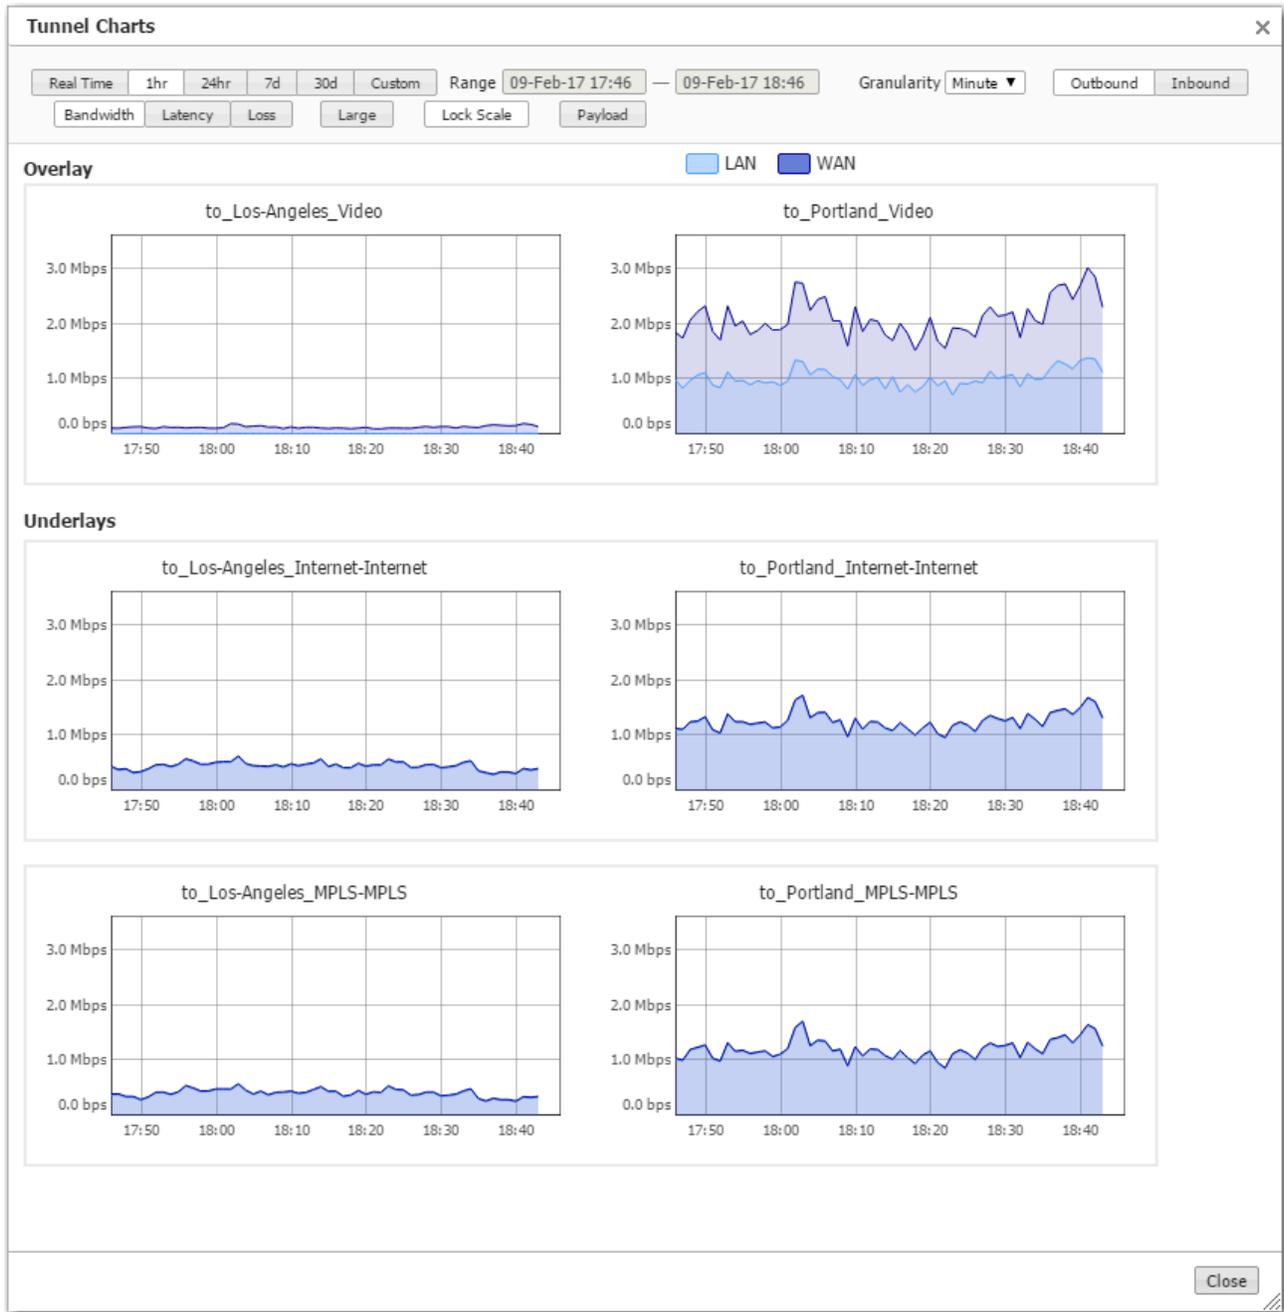
From the table, you can access the link's **Live View** graph.



In real-time, LiveView shows how Silver Peak creates synergy to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and delivering consistent application performance while both underlays are in persistent brown-out state.

### Historical Charts

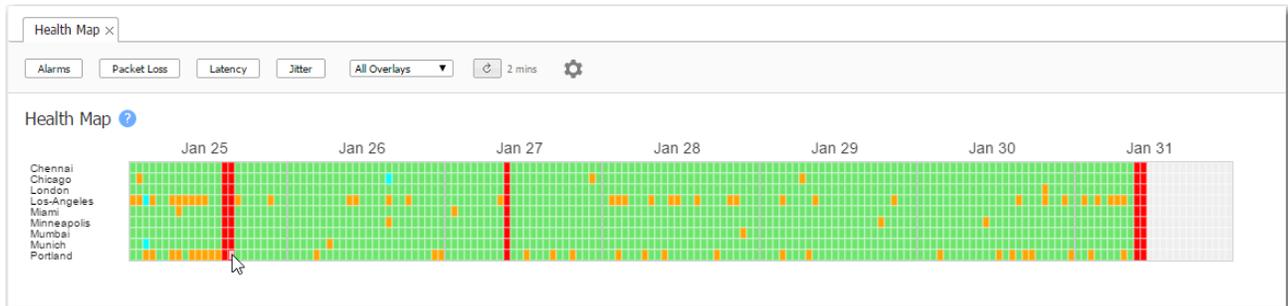
These charts let you selectively view the tunnel's components and behavior.



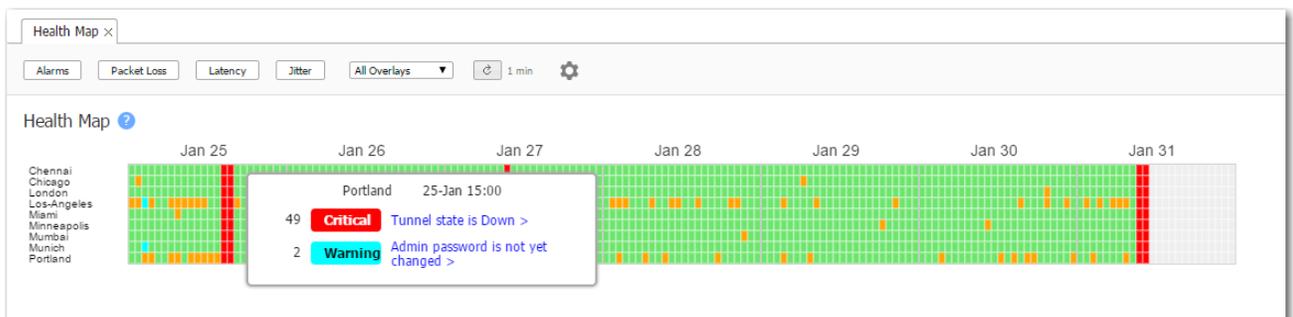
# Health Map

Monitoring > [Summary] Health Map

The **Health Map** 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*. You can configure two thresholds for each filter. 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.
  - **Green** = normal operation
  - **Orange** = marginal
  - **Red** = needs immediate attention
  - **Aqua** = warning (an alarm level)
  - **Grey** = no data available
- 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.



- 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 .



## Alarms Tab

*Monitoring > Summary > Alarms*

This tab provides various details for appliance alarms in Orchestrator.

You can apply the following filters to an alarm.

- **Time:** 1h, 4hr, 1d, 7d, or **Custom**. **Custom** allows you to select specified dates in the **Range** field.
- **Alarm Emails ON and Alarm Emails Paused:** You can enable or disable if you want to receive an email if there is an alarm that is on or paused.
- **Alarm Email Recipients:** Each configured recipient can receive emails regarding either Appliance alarms or Orchestrator alarms. Select **Add Recipient** in the **Alarm Recipients** window. Select the **alarm type** and check the boxes that you want to receive emails for. Select **Save** or **Reload**.
- **Wait to Send Emails:** You can customize the amount of time you want the system to wait to send you an email notifying you of an alarm. Select this icon and enter the amount of minutes you want the system to wait in the **Wait to Send Emails** window.

### Disable Alarms

You can specify which alarms you want to disable by selecting **Customize / Disable Alarms**.

To disable alarms:

1. Select **Disable All Alarms on Specific Appliances**.
2. Enter the name of the appliance that has the alarms you want disabled.
3. Select **Disable Alarms**.
4. Select **Save**.

### Customize Alarms

Complete the following steps to customize a pre-existing alarm.

1. Select the **Edit** icon next to the selected appliance in the Alarm Information window.
2. Choose **Enable/Disable**.
3. If selecting, **Enable**, specify the **Custom Severity** by choosing from the list: **None, CRITICAL, MAJOR, MINOR, WARNING**.

If selecting **Disable**, the following message will display: \*You are about to disable this alarm. Select **Save**.

- **Export:** You can export a CSV file of your alarms.
- **Additional 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.

**NOTE** Orchestrator keeps alarms for 90 days.

## Alarm Severity

Alarms have one of four severity levels: **None**, **Critical**, **Major**, **Minor**, or **Warning**. Only Critical and Major alarms are service-affecting.

- **None:** no level of severity has been applied to the alarm.
- **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.

## Alarm Recipients

Complete the following to add alarm recipients to receive an email notifying you of an alarm within your network.

1. Select **Alarm Email Recipients**.
2. Select **Add Recipient**.
3. Enter the following information in the correct fields.
  - The Hostname is **Orchestrator** for Orchestrator alarms, and **<Appliance hostname>** for appliance-generated alarms.

- Groups display in a drop-down list, based on the groups configured in the navigation pane.
- By default, alarms are **HTML formatted**. However, you can choose **Plain Text** or **Both**.
- **Plain Text** alarms are emailed as pipe-separated data. Users can create a script to parse the email and read the fields.

Example:

```
Hostname|Alarm_Status|Time|Alarm_ID|Type_  
ID|Source|Severity|Description|Recommended_action
```

```
Orchestrator|1|1526341365000|94|6815775|orchestrator|MINOR|Backup configuration not  
set|
```

```
Orchestrator|1|1526341362000|93|6815762|orchestrator|MAJOR|Orchestrator is using the  
default SMTP settings
```

- The **Alarm ID** is the auto-incremented, primary key in the database.
- **Alarm Status**: 1 - Raised | 2 - Cleared

## Additional Alarm Indications

- A cumulative (Orchestrator + appliances) alarm summary always displays at the right side of the header. Clicking it opens a top-level summary and access to the Alarms tab.
- Appliances are color-coded to indicate their severest alarms in the Topology tab and in the navigation pane.
- **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](#) and [Threshold Crossing Alerts Tab](#).

# Configuring & Distributing Custom Reports

Monitoring > [Reporting] Schedule & Run Reports

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

Schedule & Run Reports

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

*for Global Report*

Name  
Global Report ▼

New Report Delete Report

Email Recipients [email image sizes](#)

dmerwin@silver-peak.com

*(separate with commas or semicolons)*

**Appliances in Report** Use Tree Selection

All Appliances

**Data Granularity - Time Range**

Daily - 14 days

Hourly - 24 hours

**Scheduled or Single Report**

Run Scheduled Report

Every day at 0:00 starting 17-Oct-17 9:31 PDT Edit

Run Single Report with Custom Time Range

2018-05-07 16:26 - 2018-05-14 16:26

Run Now Stop

**Top Traffic Type**

50 ▼ Optimized Traffic ▼

**Application Charts**

Application Bandwidth ↗

Application Pie Charts ↗

Application Trends ↗

Filter App

**Tunnel Charts**

All Overlays ▼

Health Map ↗

Flow Counts ↗

Packet Counts ↗

Loss ↗

Loss Trends ↗

Out-Of-Order Packets ↗

Out-Of-Order Packet Trends ↗

Latency ↗

Latency Trends ↗

Jitter ↗

Jitter Trends ↗

Tunnel Bandwidth ↗

Tunnel Bandwidth Trends ↗

Tunnel Bandwidth Pie Charts ↗

DRC Trends ↗

Filter Tunnel

**Appliance Charts**

Top Talkers ↗

Top Domains ↗

Top Ports ↗

Top Countries ↗

Bandwidth Cost Savings ↗

Appliance Bandwidth ↗

Bandwidth Utilization ↗

Appliance Bandwidth Trends ↗

Appliance Max Bandwidth ↗

Appliance Flow Counts ↗

Appliance Flow Trends ↗

Appliance Packet Counts ↗

DSCP Bandwidth ↗

Traffic Class Bandwidth ↗

DSCP Pie Charts ↗

Traffic Class Pie Charts ↗

DSCP Trends be ▼ ↗

QoS Trends 1 ▼ ↗

Interface Bandwidth Trends ↗

Interface Summary ↗

Save Cancel

- On schedule or on demand, Orchestrator can generate Daily, Hourly, and/or Minute Reports containing user-selected charts.
- The **Email Recipients** field is pre-populated from the settings in the Orchestrator Configuration Wizard.
- The Orchestrator server sends **reports via email**.
  - To send a test email and/or to configure another SMTP server instead, go to **Orchestrator > [Setup] SMTP Server Settings**.
  - If a test email doesn't arrive within minutes, check your firewall.
- **Global Report** - By default, Orchestrator emails this preconfigured subset of charts every day. Clicking on a chart's image opens the associated tab in the browser.
- To access all reports residing on the Orchestrator server, click **View Reports**.
- Default range of reports: **Daily** = 14 days, **Hourly** = 24 hours. Increasing the scope uses additional memory.



**TIP** To specify the timezone for scheduled jobs and reports, go to **Orchestrator > Timezone for Scheduled Jobs**.

---

## View Reports

*Monitoring > [Reporting] View Reports*

Use this page to **view** and **download** reports in PDF form. Reports can be filtered by keywords or sorted by **name**, **size**, or **date last modified**. These reports can also be emailed depending on the configuration set in the **Schedule & Run Reports** tab.

View Reports x

View Reports ?

90 Rows Search

Report	File Size	Last Modified	Download
08.Dec.16-07.30.03-Daily-Global_Report.pdf	337 KB	07-Dec-16 23:31	
08.Dec.16-07.30.03-Minutes-Global_Report.pdf	334 KB	07-Dec-16 23:33	
08.Dec.16-07.30.03-Hourly-Global_Report.pdf	333 KB	07-Dec-16 23:33	
09.Dec.16-07.30.03-Daily-Global_Report.pdf	336 KB	08-Dec-16 23:31	
09.Dec.16-07.30.03-Minutes-Global_Report.pdf	334 KB	08-Dec-16 23:32	
09.Dec.16-07.30.03-Hourly-Global_Report.pdf	333 KB	08-Dec-16 23:32	
10.Dec.16-07.30.03-Daily-Global_Report.pdf	337 KB	09-Dec-16 23:31	
10.Dec.16-07.30.03-Minutes-Global_Report.pdf	334 KB	09-Dec-16 23:32	
10.Dec.16-07.30.03-Hourly-Global_Report.pdf	333 KB	09-Dec-16 23:33	
11.Dec.16-07.30.03-Daily-Global_Report.pdf	337 KB	10-Dec-16 23:31	
11.Dec.16-07.30.03-Minutes-Global_Report.pdf	334 KB	10-Dec-16 23:33	
11.Dec.16-07.30.03-Hourly-Global_Report.pdf	333 KB	10-Dec-16 23:34	
12.Dec.16-07.30.03-Daily-Global_Report.pdf	337 KB	11-Dec-16 23:31	
12.Dec.16-07.30.03-Minutes-Global_Report.pdf	334 KB	11-Dec-16 23:34	
12.Dec.16-07.30.03-Hourly-Global_Report.pdf	333 KB	11-Dec-16 23:34	
13.Dec.16-07.30.03-Daily-Global_Report.pdf	337 KB	12-Dec-16 23:31	
13.Dec.16-07.30.03-Minutes-Global_Report.pdf	334 KB	12-Dec-16 23:32	
13.Dec.16-07.30.03-Hourly-Global_Report.pdf	333 KB	12-Dec-16 23:36	
13.Dec.16-07.30.03-Daily-Global_Report.pdf	337 KB	13-Dec-16 23:31	

# Sample Report



## Scheduled & Historical Jobs

Monitoring > [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.

Scheduled & Historical Jobs x

Scheduled Jobs Historical Jobs Export ↻

Scheduled Jobs ⓘ

2 Rows Search

Job	Appliances	Description	Schedule	Last Run	Next Run ↕	Status	
Orchestrator Report	Silver Peak Systems	Global Report	Every day at 10:10 starting 15-Jul-15 17:53 GMT	13-Jan-17 10:10 GMT	14-Jan-17 10:10 GMT	Success - Global Report Time t...	✕
Orchestrator Backup	All appliances	Weekly Orchestrator Bac...	Every Friday at 0:30 starting 30-Jun-16 21:34 GMT	13-Jan-17 00:30 GMT	20-Jan-17 00:30 GMT	Failed - 13-Jan-17 00:30 GMT - ...	✕

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

Scheduled & Historical Jobs x

Scheduled Jobs Historical Jobs Export ↻

Historical Jobs ⓘ

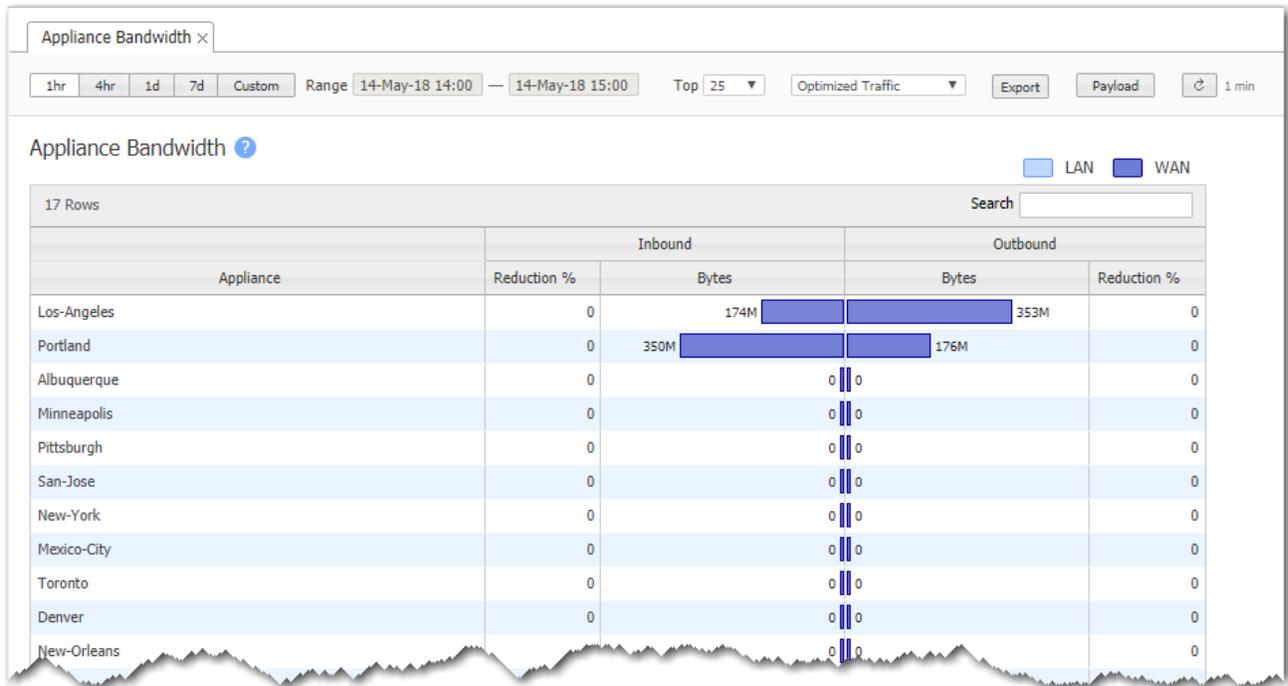
805 Rows Search

Job	Appliances	Description	Start Time ↕	Duration	Status
Orchestrator Report	Silver Peak Systems	Global Report	13-Jan-17 10:10 GMT	3m 47s	Success - Global Report Time taken(s): 227 ...
Orchestrator Backup	All appliances	Weekly Orchestrator Backup	13-Jan-17 00:30 GMT	17m 11s	Failed - 13-Jan-17 00:30 GMT - Backing up A...
Orchestrator Report	Silver Peak Systems	Global Report	12-Jan-17 10:10 GMT	4m 15s	Success - Global Report Time taken(s): 255 ...
Orchestrator Report	Silver Peak Systems	Global Report	11-Jan-17 10:10 GMT	4m 18s	Success - Global Report Time taken(s): 258 ...
Orchestrator Report	Silver Peak Systems	Global Report	10-Jan-17 10:10 GMT	4m 50s	Success - Global Report Time taken(s): 290 ...
Orchestrator Report	Silver Peak Systems	Global Report	09-Jan-17 10:10 GMT	3m 30s	Success - Global Report Time taken(s): 210 ...
Orchestrator Report	Silver Peak Systems	Global Report	08-Jan-17 10:10 GMT	3m 31s	Success - Global Report Time taken(s): 211 ...
Orchestrator Report	Silver Peak Systems	Global Report	07-Jan-17 10:10 GMT	3m 27s	Success - Global Report Time taken(s): 207 ...
Appliance Reboot	Asia,Europe,US-East,US-West		06-Jan-17 22:23 GMT	0s	Failed - Failed to run reboot/shutdown sche...
Orchestrator Report	Silver Peak Systems	Global Report	06-Jan-17 10:10 GMT	3m 32s	Success - Global Report Time taken(s): 212 ...
Orchestrator Backup	All appliances	Weekly Orchestrator Backup	06-Jan-17 00:30 GMT	8s	Failed - 06-Jan-17 00:30 GMT - Backing up A...
Orchestrator Report	Silver Peak Systems	Global Report	05-Jan-17 10:10 GMT	4m 43s	Success - Global Report Time taken(s): 283 ...
Orchestrator Report	Silver Peak Systems	Global Report	04-Jan-17 10:10 GMT	4m 7s	Success - Global Report Time taken(s): 247 ...
Orchestrator Report	Silver Peak Systems	Global Report	03-Jan-17 10:10 GMT	8m 21s	Failed - Global Report, Error: Failed to run re...
Orchestrator Report	Silver Peak Systems	Global Report	02-Jan-17 10:10 GMT	6m 26s	Success - Global Report Time taken(s): 386 ...
Orchestrator Report	Silver Peak Systems	Global Report	01-Jan-17 10:10 GMT	3m 36s	Success - Global Report Time taken(s): 216 ...
Orchestrator Report	Silver Peak Systems	Global Report	31-Dec-16 10:10 GMT	3m 33s	Success - Global Report Time taken(s): 213 ...
Orchestrator Report	Silver Peak Systems	Global Report	30-Dec-16 10:10 GMT	4m 22s	Success - Global Report Time taken(s): 262 ...
Orchestrator Backup	All appliances	Weekly Orchestrator Backup	30-Dec-16 00:30 GMT	16m 16s	Failed - 30-Dec-16 00:30 GMT - Backing up A...
Orchestrator Report	Silver Peak Systems	Global Report	29-Dec-16 10:10 GMT	3m 51s	Success - Global Report Time taken(s): 231 ...
Orchestrator Report	Silver Peak Systems	Global Report	28-Dec-16 10:10 GMT	3m 19s	Success - Global Report Time taken(s): 199 ...
Orchestrator Report	Silver Peak Systems	Global Report	24-Dec-16 10:10 GMT	3m 32s	Success - Global Report Time taken(s): 212 ...
Orchestrator Report	Silver Peak Systems	Global Report	23-Dec-16 10:10 GMT	9m 25s	Failed - Global Report, Error: Failed to run re...

# Appliance Bandwidth

Monitoring > [Bandwidth > Appliances] Summary

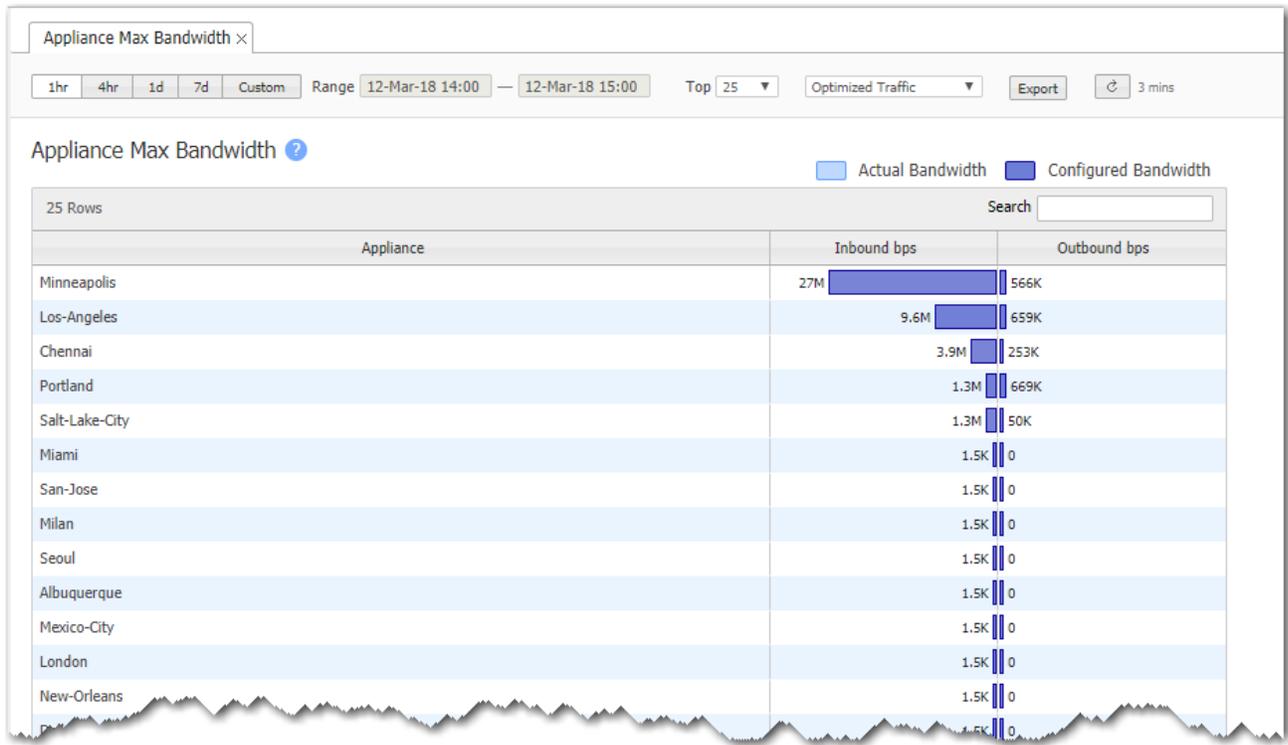
The **Appliance Bandwidth** chart lists the top appliances based on the total volume of inbound and outbound traffic before reduction. It shows how many bytes the Silver Peak appliance saved when transferring data, aggregated over a selectable time period.



# Appliance Max Bandwidth

*Monitoring > [Bandwidth > Appliances] Max*

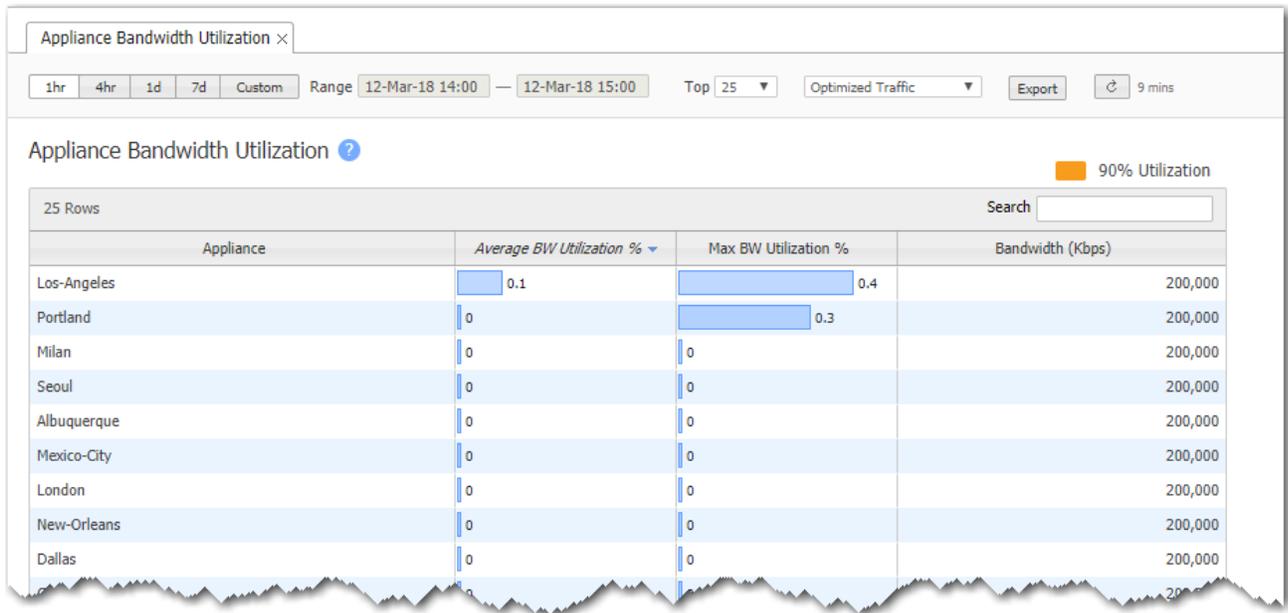
The **Appliance Max Bandwidth** chart lists the top appliances by the peak throughput (in either direction), within a selected time period. It compares the system bandwidth of the appliance to the effective bandwidth it's providing.



# Appliance Bandwidth Utilization

Monitoring > [Bandwidth > Appliances] 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.

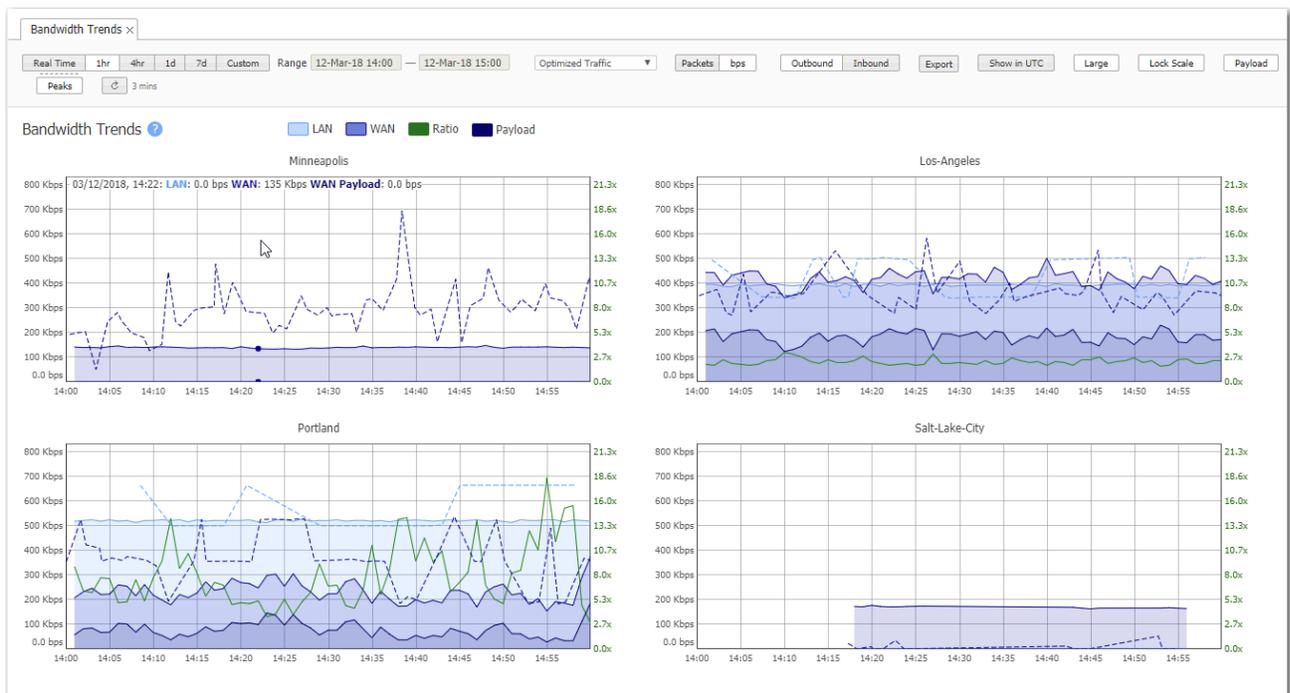


# Appliance Bandwidth Trends

Monitoring > [Bandwidth > Appliances] Trends

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

For each Business Intent Overlay, the Link Bonding Policy specified determines the bandwidth efficiency. To guarantee service quality levels, High Availability requires the most overhead, and High Efficiency requires the least. Charts display the total bandwidth used. The Payload option shows how much raw data is transmitted. At the same time, it exposes the Peaks option, which enables the viewing of peak transmissions.



# Appliance Packet Counts

Monitoring > [Bandwidth > Appliances] Packet Counts

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

Appliance Packet Counts x

1hr 4hr 1d 7d Custom Range 12-Mar-18 14:00 — 12-Mar-18 15:00 Top 25 Optimized Traffic Export Payload 2 mins

Appliance Packet Counts ?

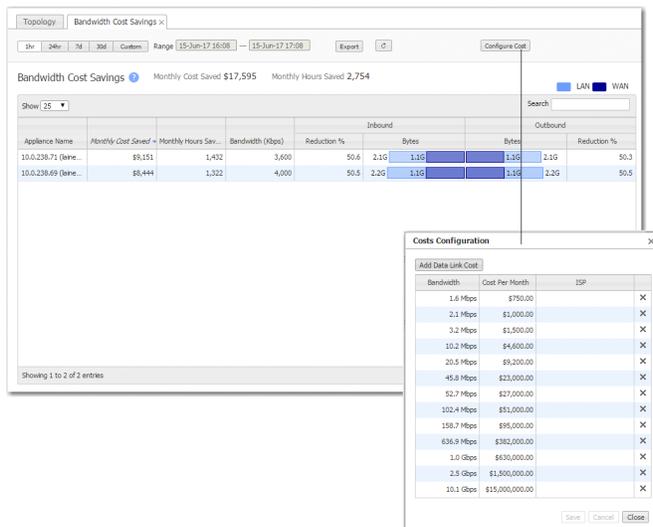
17 Rows Search

Appliance	Inbound			Outbound		
	LAN Packets	LAN Max pps	WAN Packets	LAN Packets	LAN Max pps	WAN Packets
Los-Angeles	187,741	880	582,249	279,684	808	604,592
Portland	151,491	132	443,103	284,349	92	433,854
Denver	0	2	190,119	0	0	191,507
New-York	0	2	193,013	0	0	196,445
Boston	0	2	171,849	0	0	173,435
Minneapolis	0	2,574	286,378	0	1,188	289,339
Toronto	0	2	170,280	0	0	170,761
Miami	0	2	177,138	0	0	178,810
San-Jose	0	2	201,153	0	0	202,662
New-Orleans	0	2	208,142	0	0	210,076
Dallas	0	2	186,596	0	0	189,649
San-Diego	0	2	177,138	0	0	178,810

# Appliance Bandwidth Cost Savings

Monitoring > [Bandwidth > Appliances] 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.



## Calculations

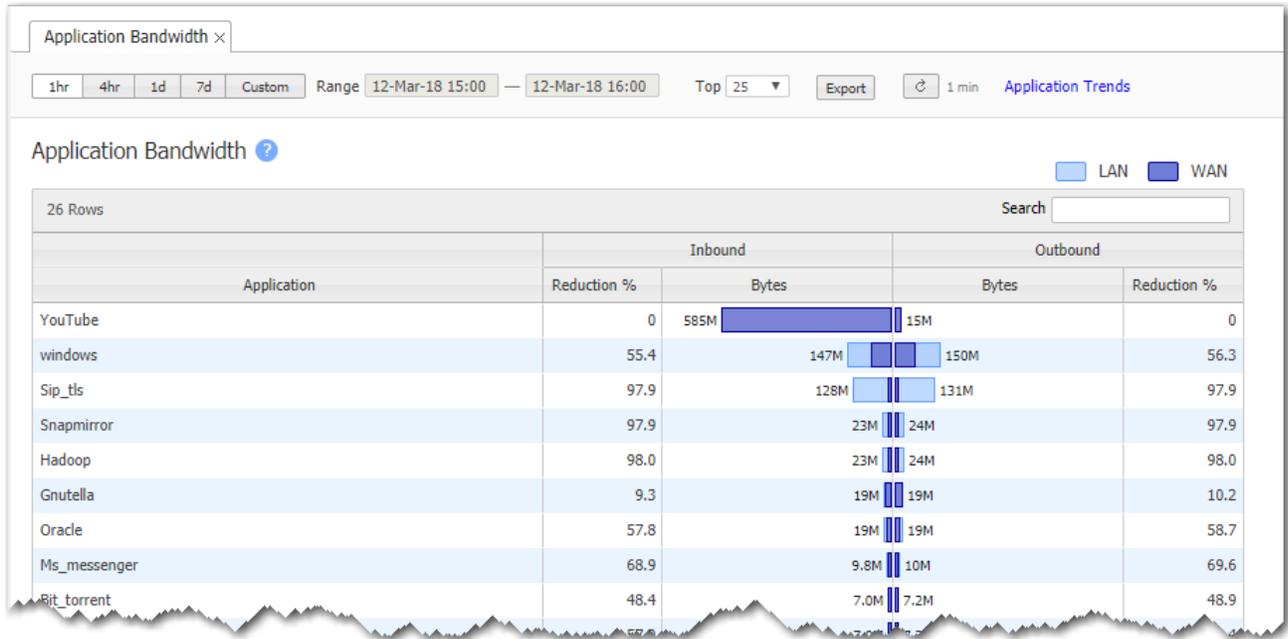
- For the monthly cost savings, it subtracts the maximum bytes you could send **without** the appliance from how many bytes the Silver Peak **actually** sent, and multiplies the difference by the data link cost.
- For the monthly time savings, it uses the data link speed to calculate how many more hours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

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

# Application Bandwidth

Monitoring > [Bandwidth > Applications] Summary

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

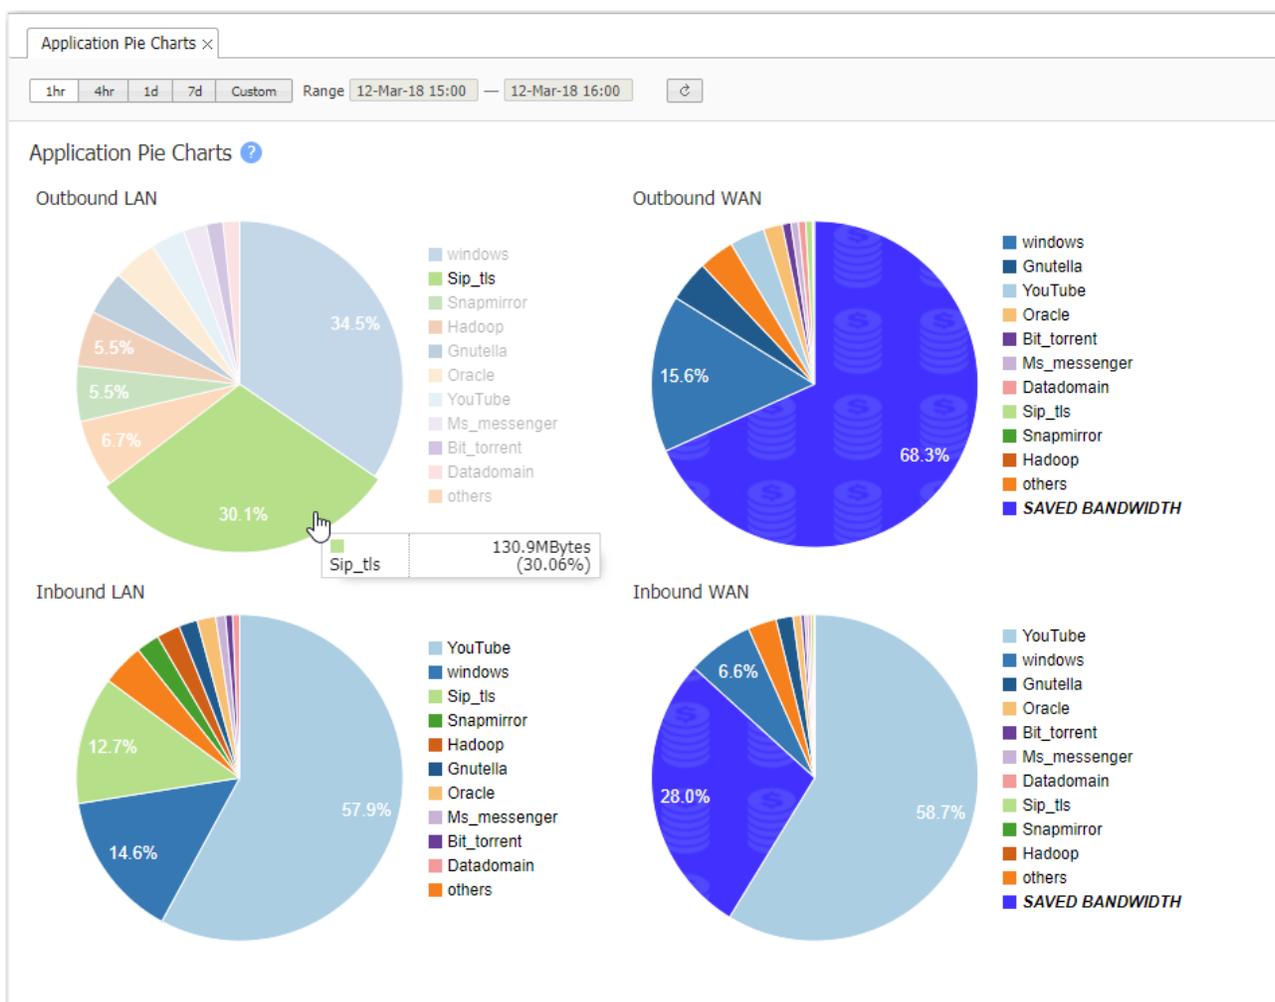


## Application Pie Charts

Monitoring > [Bandwidth > Applications] 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 > [Bandwidth > Applications] Trends

This tab shows application trends over time.



## Firewall Drops

You can use the Firewall Drops tab to see the statistics on various flows, packets, and bytes dropped or allowed by a zone-based firewall for a given time range.

*Monitoring > Bandwidth > Firewall Drops > Summary*

- You can select a range of time (in hours and days) to view the firewall drops. You can also select if you want to view in Matrix or Table view.
- Select Export to export the report to an excel spreadsheet.

Dashboard | Business Intent Overlays | Security Policies | Flows | Firewall Drops x

Security Policies | Real Time | 1hr | 4hr | 1d | 7d | Custom | Range 11-Dec-18 08:00 – 11-Dec-18 09:00 | Matrix View | Table View | Export | 11 mins

Firewall Drops ⓘ

Appliance Name	From Zone	To Zone	Flows Dropped	Flows Allowed	Packets Dropped	Packets Allowed	Bytes Dropped	Bytes Allowed	Charts	
Los-Angeles	Default	CorporateWAN	0	0	0	2.5K	0	577.5K	📊	
Portland	CorporateWAN	Default	0	0	0	1.9K	0	461.5K	📊	
Chennai			No Data Available							
Mumbai			No Data Available							

- In the charts column, you can select the chart icon.
  - In this pop-up, you can see packets, and bytes dropped or allowed by a zone-based firewall for a given time range.



## Top Talkers

Monitoring > [Bandwidth > Identifiers] Top Talkers

This tab lists the IP addresses that use the most bandwidth.

Top Talkers x

1hr 4hr 1d 7d Custom Range 12-Mar-18 15:00 — 12-Mar-18 16:00 Top 25 Export 4 mins

Top Talkers 2

LAN WAN

25 Rows Search

IPs	Domain	Top Destinations	Inbound		Outbound		Flows Started	Flows Ended
			Bytes	Bytes	Bytes	Bytes		
10.17.9.11			402M	109M	109M	411M	0	0
10.17.11.11			402M	109M	109M	411M	0	0
10.17.9.10			241M		10M		3469	3763
173.194.167.231	r2---sn-n4v7knll.googlevideo.com		234M		6.1M		155	153
10.17.15.10			185M		7.8M		1053	1123
10.17.17.10			180M		5.9M		1377	1326
173.194.166.104	r2---sn-n4v7sne7.googlevideo.com		51M		929K		13	13
173.194.167.233	r4---sn-n4v7knll.googlevideo.com		46M		792K		17	17
74.125.170.234	r4---sn-n4v7sn76.googlevideo.com		30M		718K		7	7
173.194.167.7	r2---sn-n4v7knl6.googlevideo.com		25M		390K		4	4
74.125.166.201	r3---sn-o097znld.googlevideo.com		25M		534K		6	6
74.125.166.168	r2---sn-o097znl6.googlevideo.com		21M		458K		7	7
173.194.12.92	r6---sn-o097znlr.googlevideo.com		19M		447K		4	4
173.194.12.74	r4---sn-o097znlk.googlevideo.com		16M		397K		5	5
74.125.170.216	r2---sn-n4v7sn7z.googlevideo.com		16M		241K		3	3

You can also view each IP's destinations.

10.17.15.10's Destinations				
10 Rows				
Destination	Inbound Bytes	Outbound Bytes	Flows Started	Flows Ended
r4---sn-n4v7knll.googlevideo.com (173.194.167.233)	36M	506K	11	11
r2---sn-n4v7knl6.googlevideo.com (173.194.167.7)	25M	390K	4	4
r2---sn-n4v7sne7.googlevideo.com (173.194.166.104)	25M	324K	2	2
r2---sn-n4v7sn7z.googlevideo.com (74.125.170.216)	16M	241K	3	3
r3---sn-o097znld.googlevideo.com (74.125.166.201)	12M	226K	2	2
r1---sn-n4v7sn7l.googlevideo.com (74.125.170.183)	11M	189K	2	2
r1---sn-n4v7sn7z.googlevideo.com (74.125.170.215)	11M	175K	0	1
r2---sn-o097znl6.googlevideo.com (74.125.166.168)	10M	192K	4	4
quote.cnn.com (104.68.113.65)	3.3M	3.5M	2	3
r2---sn-n4v7sn7y.googlevideo.com (74.125.170.120)	5.6M	103K	2	2

[Close](#)

## Domains

*Monitoring > [Bandwidth > Identifiers] Domains*

This tab lists the domains that use the most bandwidth.

The number of **Subdomains** selected determines how the table aggregates subdomains for display. An asterisk (\*) indicates that more subdomains would be displayed if a higher number were selected. This is not a filter, but rather a grouping convenience.

Domains ×

1hr 4hr 1d 7d Custom
Range 14-May-18 15:00 — 14-May-18 16:00
Src Dest
Subdomains 2 ▼
Top 25 ▼
Export ↻

Domains ?

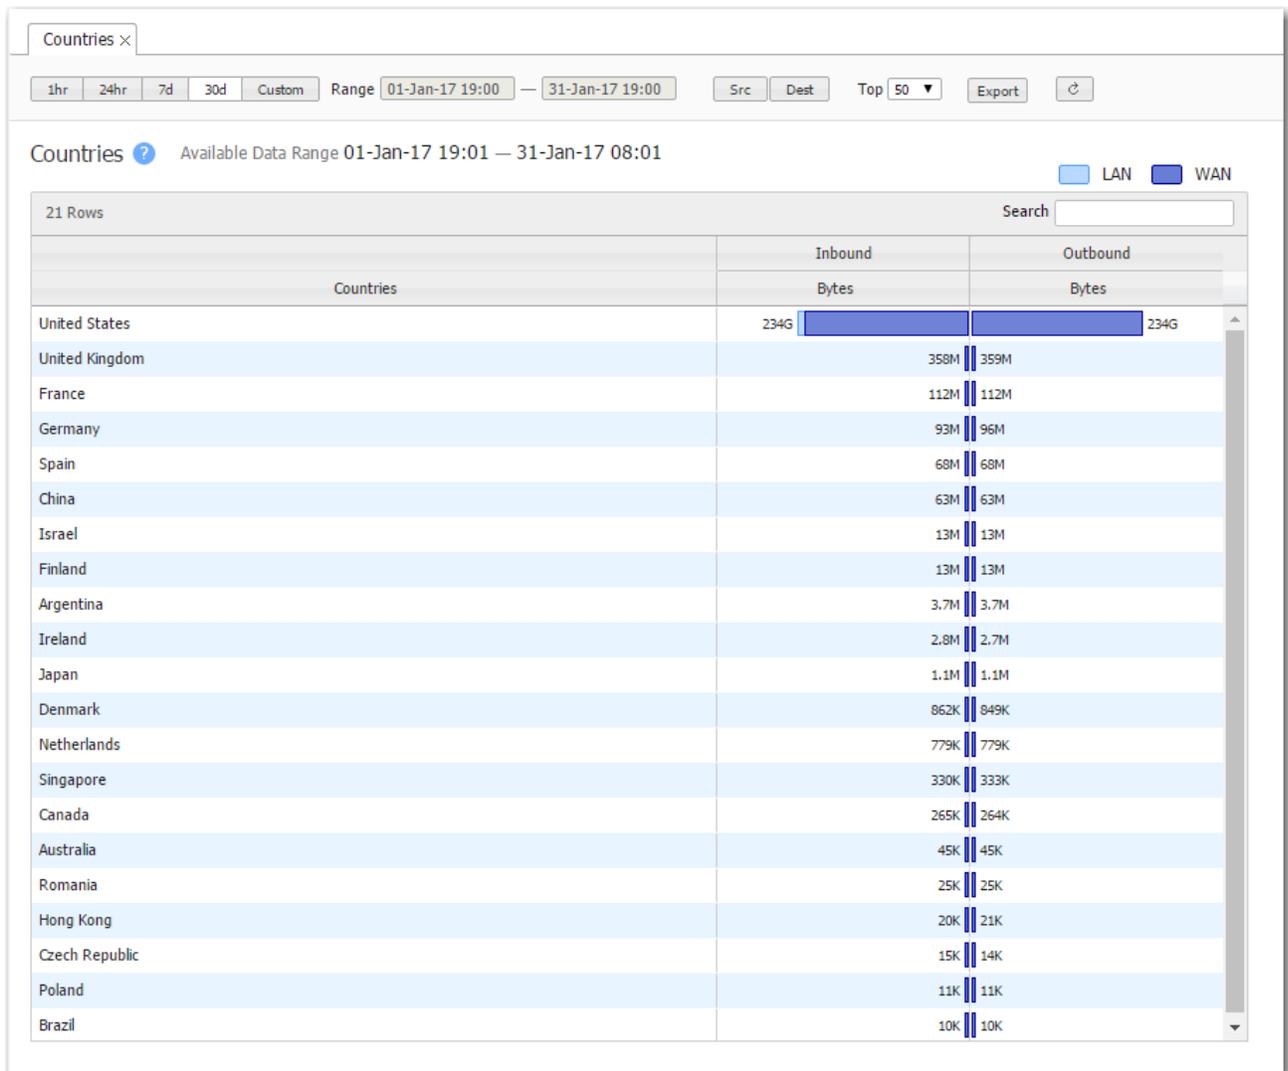
■ LAN
 ■ WAN

Domains	Inbound		Outbound	
	Reduction %	Bytes	Bytes	Reduction %
*googlevideo.com	0	145M <div style="width: 100%; height: 10px; background-color: blue;"></div>	3.3M	0
*cnbc.com	0	1.7M <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	311K	0
*youtube.com	0	173K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	197K	0
*nytimes.com	0	192K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	147K	0
*mozilla.net	0	252K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	12K	0
*nyt.com	0	97K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	31K	0
*nr-data.net	0	32K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	65K	0
*yimg.com	0	74K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	5.1K	0
*doubleclick.net	0	28K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	21K	0
*googleapis.com	0	24K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	16K	0
*googlesyndication.com	0	13K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	13K	0
*mozilla.com	0	12K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	5.5K	0
*ggpht.com	0	5.8K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	2.1K	0
*google.com	0	2.6K <div style="width: 100%; height: 10px; background-color: lightblue;"></div>	2.8K	0

# Countries

Monitoring > [Bandwidth > Identifiers] Countries

This tab lists the countries that use the most bandwidth.



# Ports

*Monitoring > [Bandwidth > Identifiers] Ports*

This tab lists the ports that use the most bandwidth.

Ports x

1hr 24hr 7d 30d Custom Range 31-Jan-17 18:00 — 31-Jan-17 19:00 Src Dest Top 50 Export

Ports ? LAN WAN

50 Rows Search

Ports	Description	Protocol	Inbound		Outbound	
			Bytes		Bytes	
1521	nCube	TCP	379M		581M	
10565		TCP	237M		479M	
5001	complex-link	TCP	236M		476M	
1863	MSNP	TCP	99M		200M	
54330		TCP	99M		200M	
39699		TCP	95M		146M	
50655		TCP	79M		161M	
48869		TCP	80M		160M	
54873		TCP	79M		159M	
57049		TCP	93M		143M	
48227		TCP	59M		119M	
45796		TCP	59M		119M	
5631	pcANYWHEREdata	TCP	57M		114M	
54150		TCP	59M		92M	
36371		TCP	56M		86M	
49713		TCP	40M		81M	
40650		TCP	39M		79M	
38547		TCP	38M		77M	
60557		TCP	38M		58M	
44535		TCP	37M		56M	
35186		TCP	29M		58M	

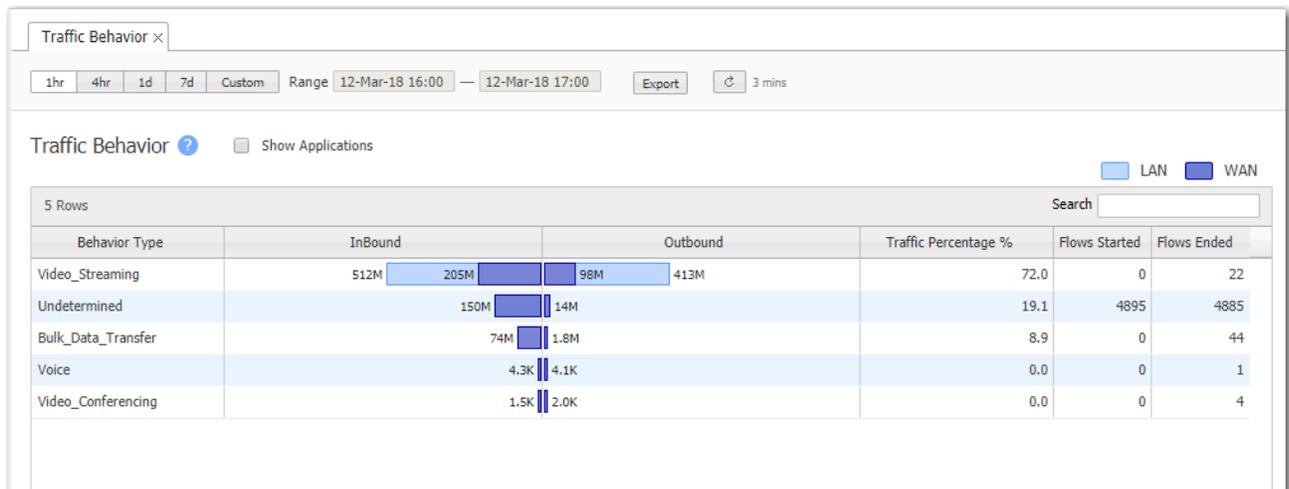
# Traffic Behavior

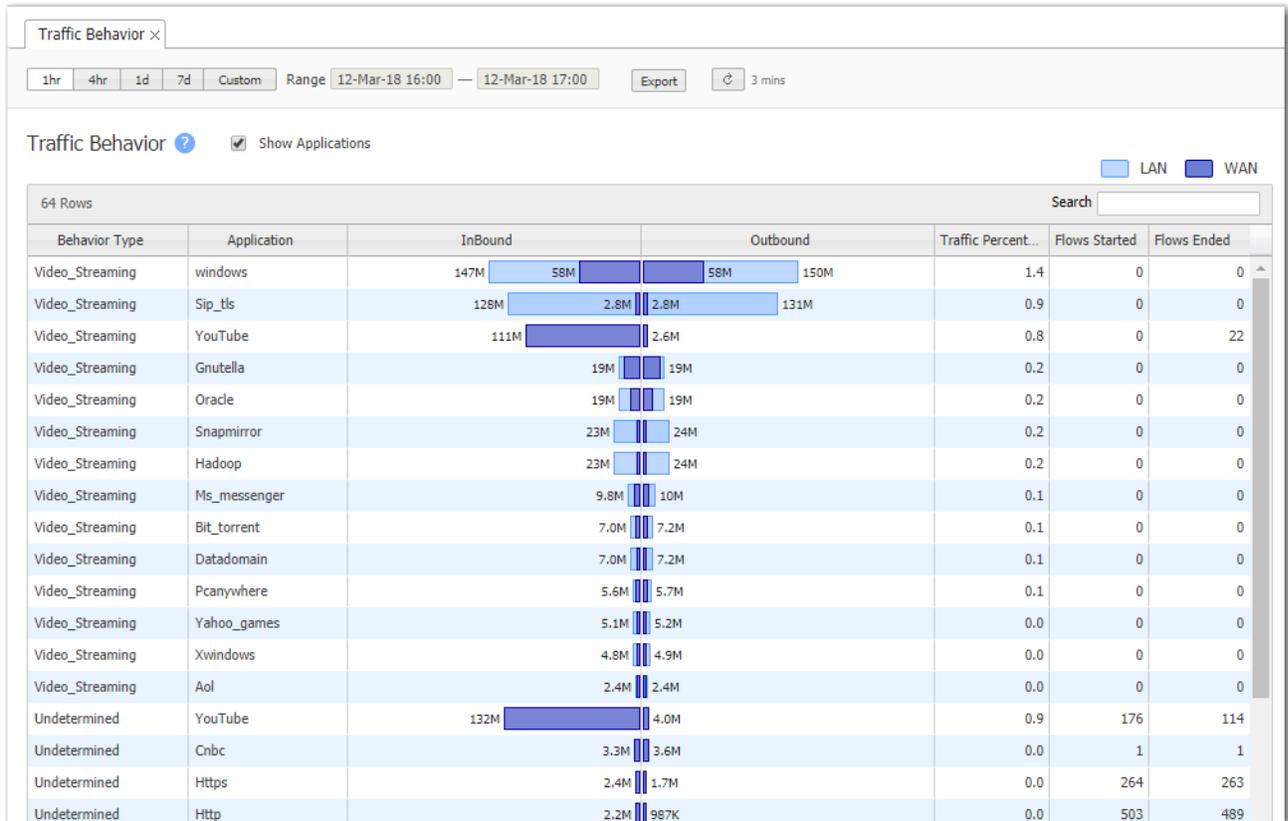
*Monitoring > [Bandwidth > Identifiers] Traffic Behavior*

The **Traffic Behavior** report identifies and categorizes traffic based on low-level characteristics of the data streams. The behavior types are:

- Voice
- Video Conferencing
- Video Streaming
- Bulk Data Transfer
- Interactive
- Undetermined

You can also specify these categories as match criteria when creating policies or ACLs (Access Control Lists).





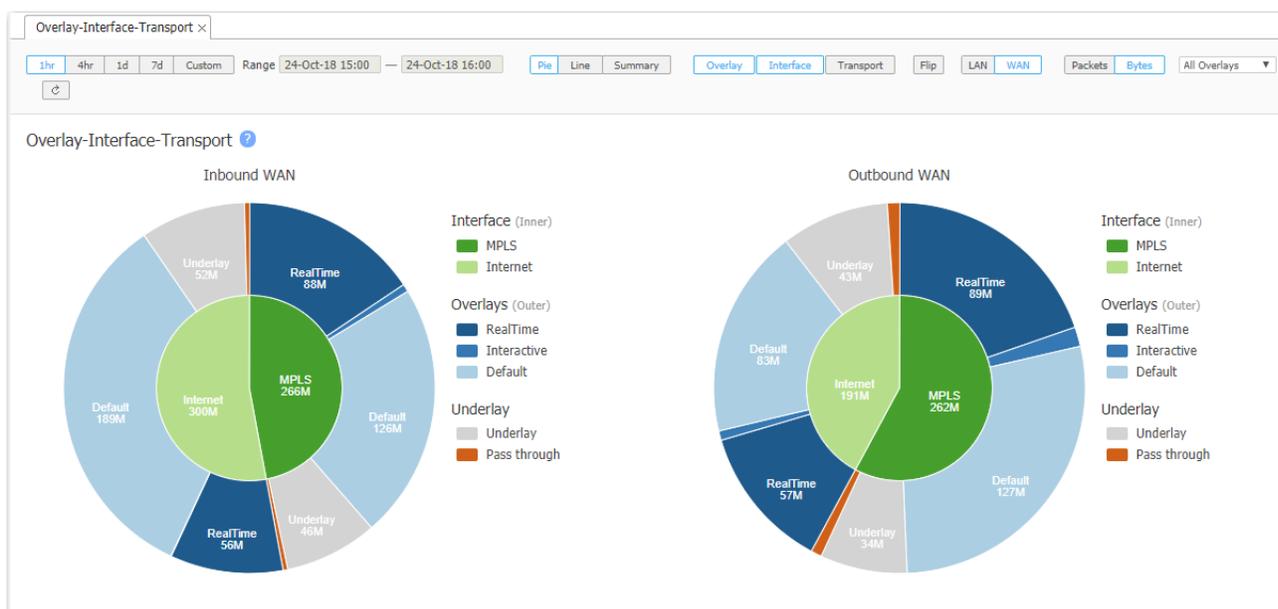
## Overlay-Interface-Transport

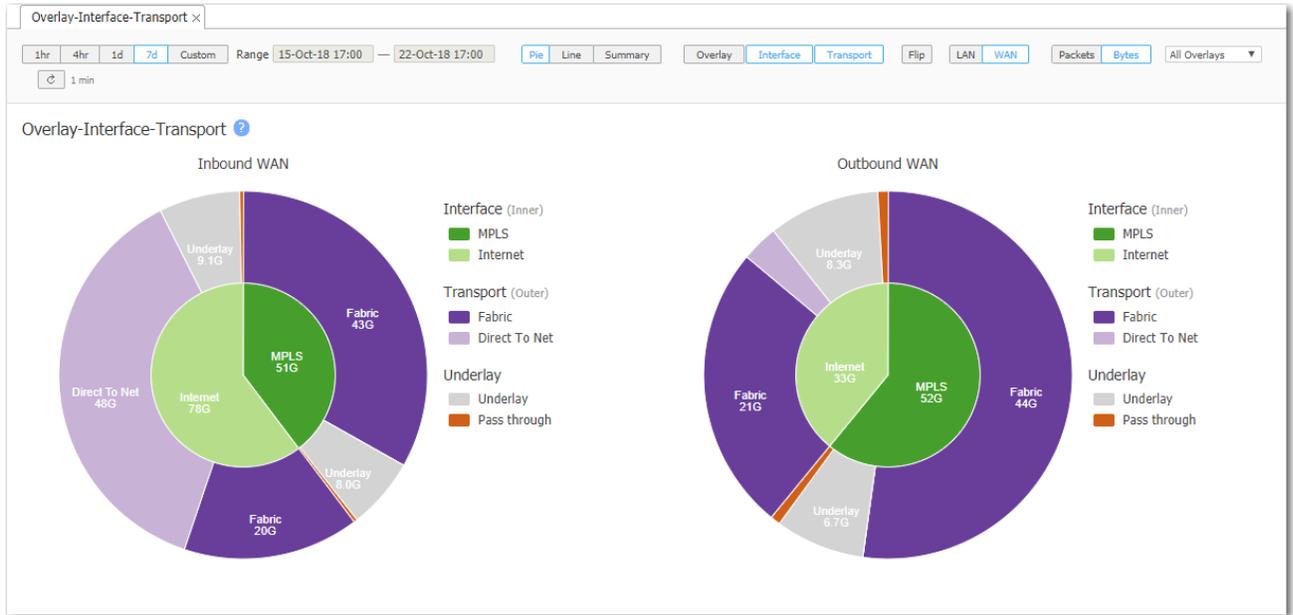
Monitoring > [Bandwidth > Overlays & Interfaces] Overlay-Interface-Transport

These charts display the distribution of traffic across three dimensions—overlays, interfaces, and transport. You can view each option individually, or in relation to another.

For instance, for a given interface, you can see how the overlay traffic is distributed.

You can also view how much traffic is transported from one Silver Peak appliance to another on the SD-WAN fabric (Overlays), versus how much is broken out locally, direct to the internet. The Underlay legend displays non-overlay traffic.

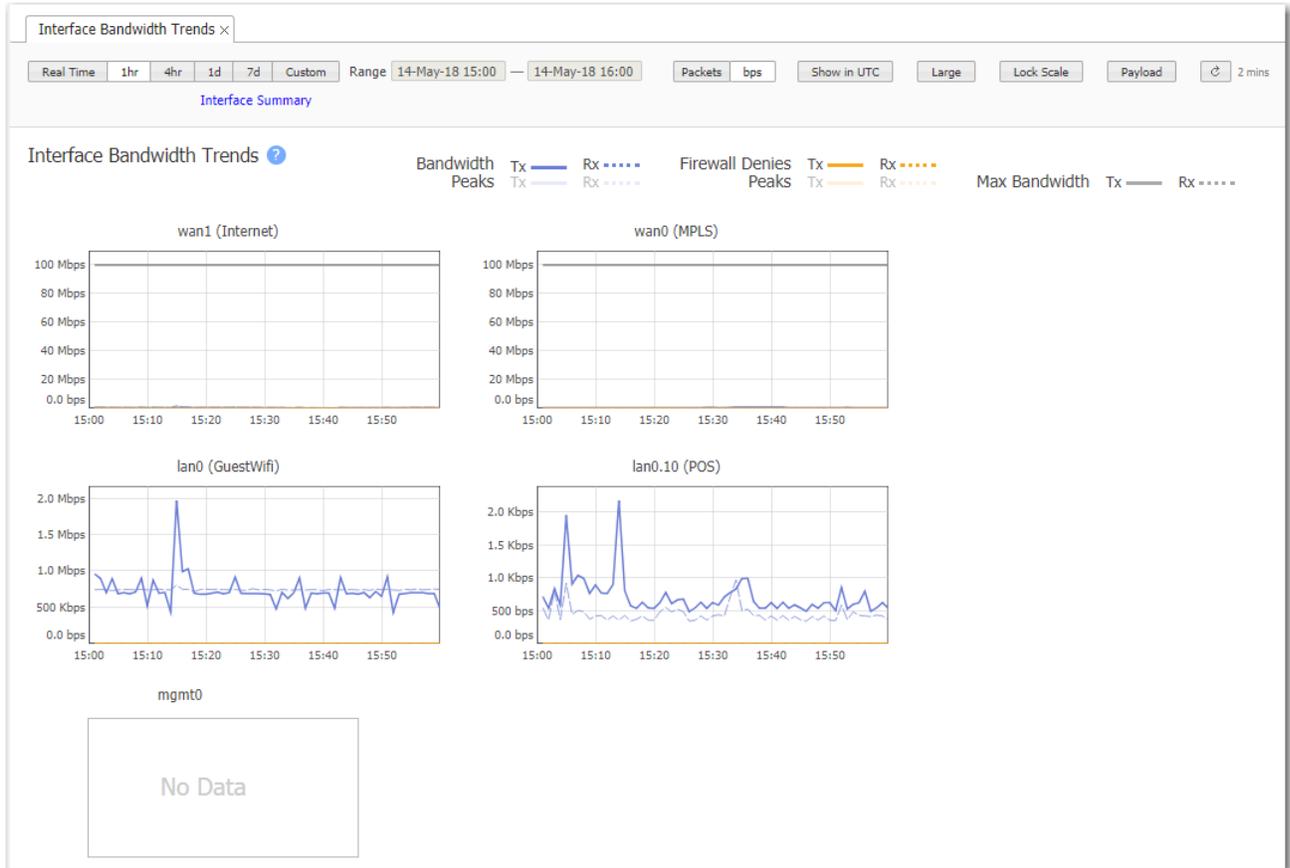




# Interface Bandwidth Trends

Monitoring > [Bandwidth > Overlays & Interfaces] Bandwidth Trends

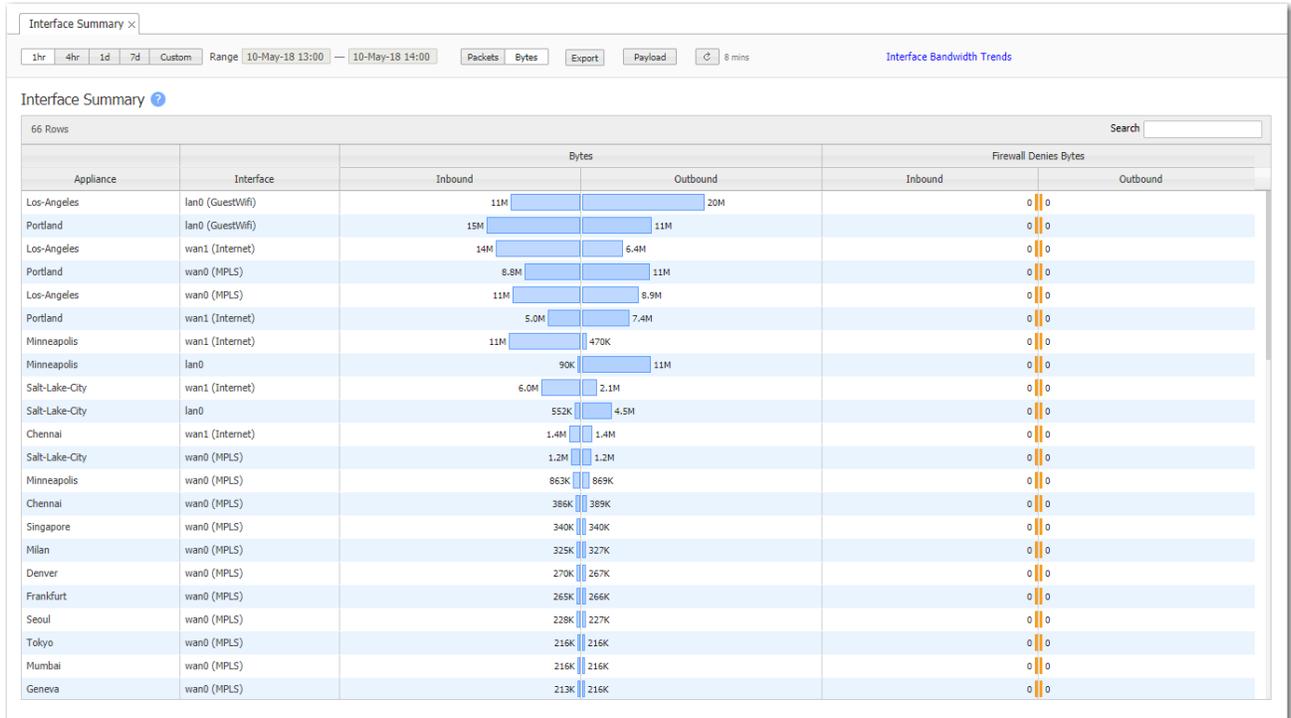
The Interface Bandwidth Trends tab shows interface statistics over time.



# Interface Summary

Monitoring > Overlays & Interfaces > Summary

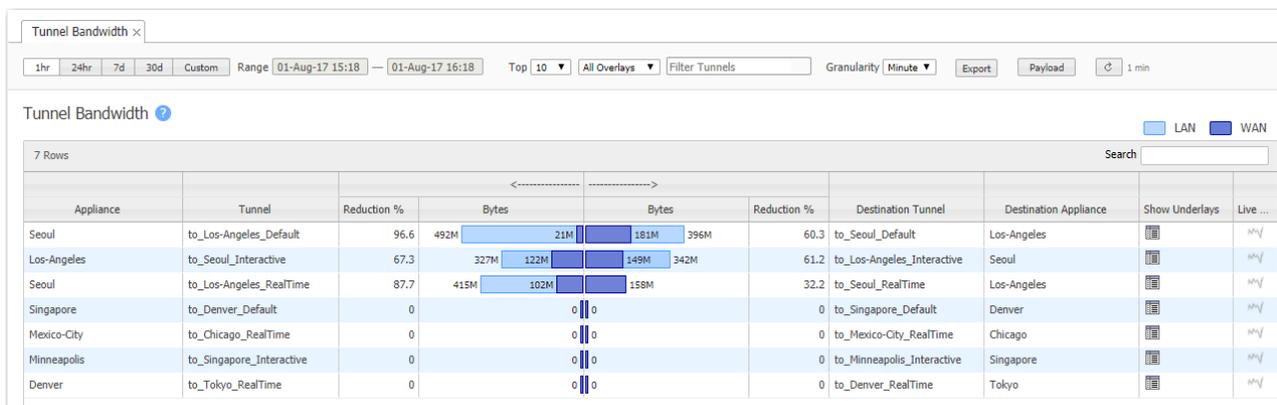
This tab shows interface summary stats, including inbound and outbound Packets or Bytes per interface, as well as Firewall Denies (Drops). The stats are summarized for the selected time period.



# Tunnels Bandwidth

Monitoring > [Bandwidth > Tunnels] Summary

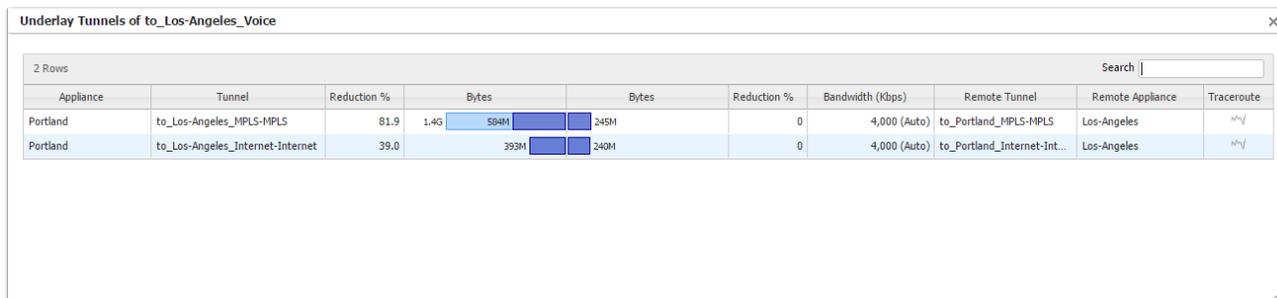
The **Tunnel Bandwidth** chart shows which tunnels are sending the most bytes – that is, the tunnels that are the most active.



## Show Underlays

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

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



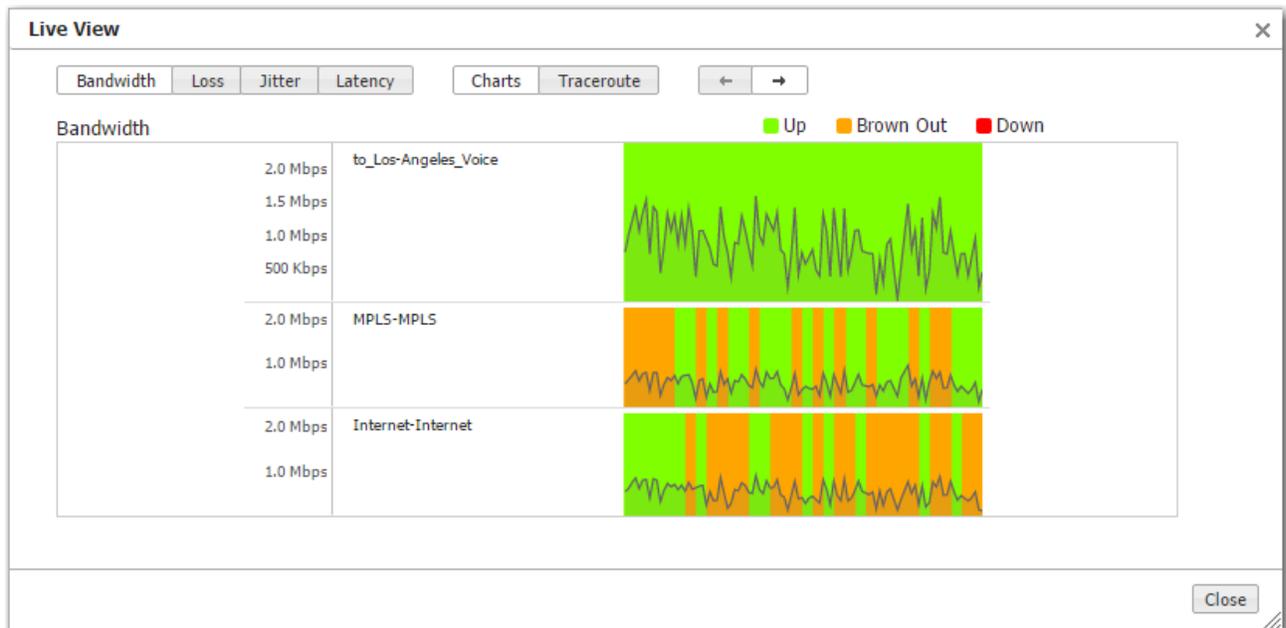
## Traceroute

This shows trace route information between the tunnel source and destination IP addresses. It shows intermediate hops, their IP addresses, and the latency between each hop.



## Live View

Live View shows the live bandwidth, loss, latency, and jitter on all the tunnels. For an overlay, it also shows live tunnel states – **Up**, **Browned Out**, or **Down**.



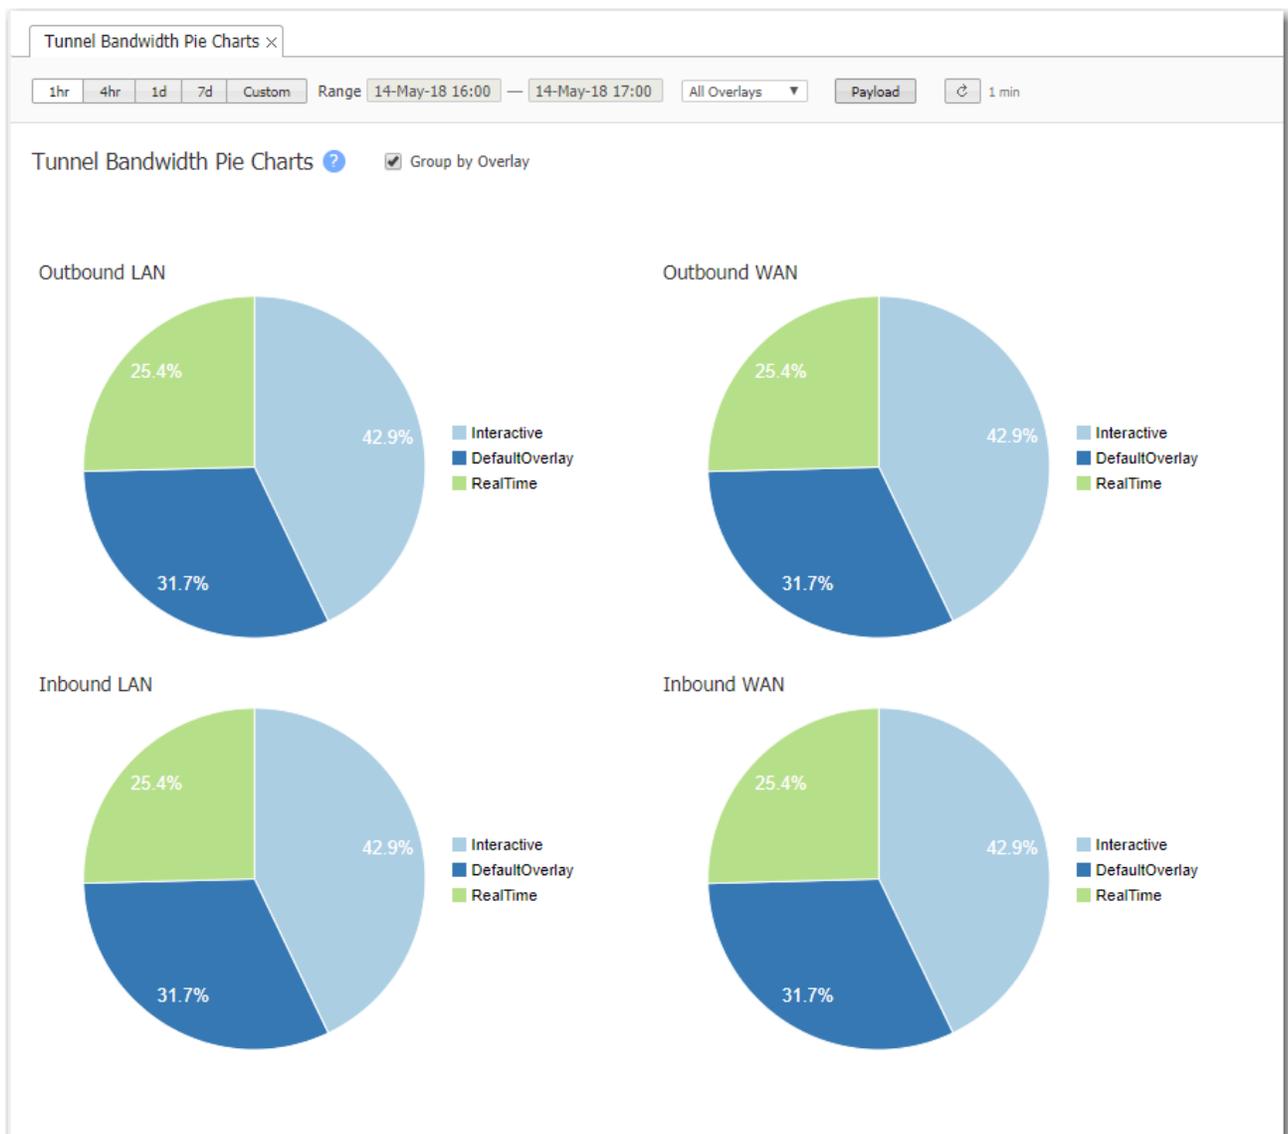
In real-time, LiveView shows how Silver Peak creates synergy to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and delivering consistent application performance while both underlays are in persistent brown-out state.

## Tunnels Pie Charts

Monitoring > [Bandwidth > Tunnels] Pie Charts

The **Tunnel Bandwidth Pie Charts** show what proportion of the bytes a tunnel consumes on the LAN and on the WAN.

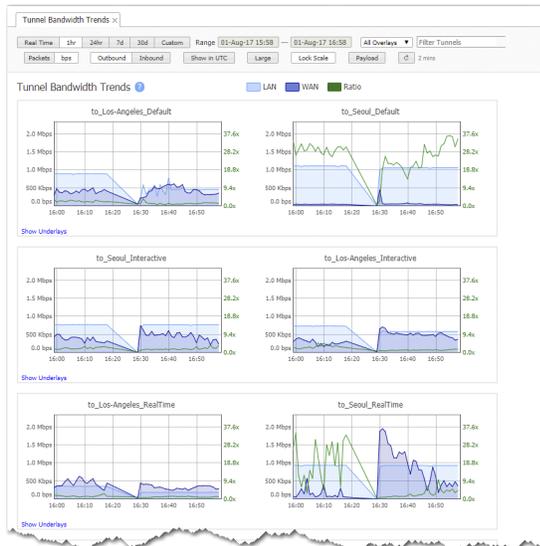
- Mousing over the charts and the legends reveals additional information.
- The WAN charts identify what percentage of the bandwidth the appliance saved by optimizing the traffic.



# Tunnel Bandwidth Trends

Monitoring > [Bandwidth > Tunnels] Trends

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



- For each Business Intent Overlay, the Link Bonding Policy specified determines the bandwidth efficiency.
- To guarantee service quality levels, High Availability requires the most overhead, and High Efficiency requires the least.
- Charts display the total bandwidth used.
- The Payload option shows how much raw data is transmitted. At the same time, it exposes the Peaks option, which enables the viewing of peak transmissions.



**NOTE** Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

# Tunnel Packet Counts

Monitoring > [Bandwidth > Tunnels] Packet Counts

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

Tunnel Packet Counts x

1hr 4hr 1d 7d Custom Range 14-May-18 17:00 — 14-May-18 18:00 Top 25 All Overlays Export Payload ↺

Tunnel Packet Counts ?

25 Rows Search

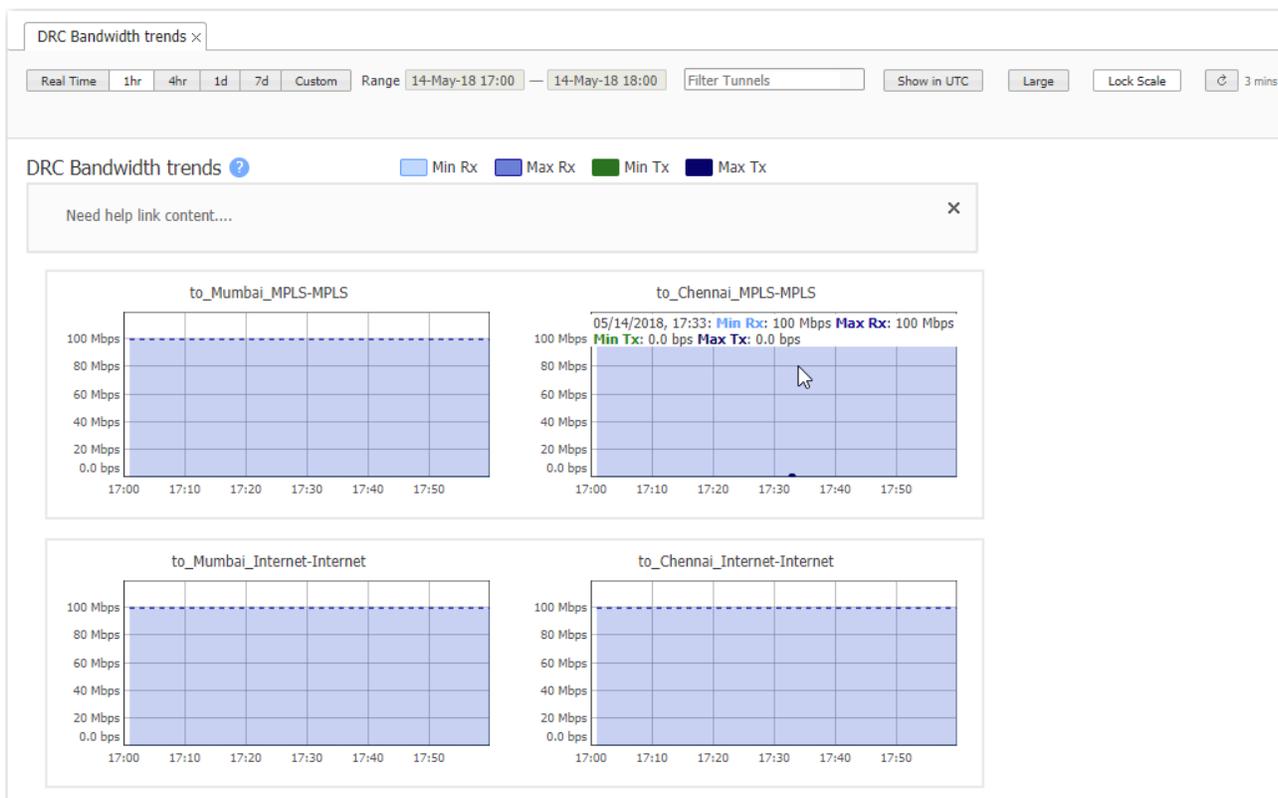
Appliance	Tunnel	Inbound			Outbound		
		LAN Packets	LAN Max pps	WAN Packets	LAN Packets	LAN Max pps	WAN Packets
Portland	to_Los-Angeles_Interactive	205,482	130	180,494	206,522	142	185,660
Los-Angeles	to_Portland_Interactive	204,394	144	183,705	207,239	134	182,053
Portland	to_Los-Angeles_DefaultOv...	140,864	114	115,397	163,154	142	155,597
Los-Angeles	to_Portland_DefaultOverlay	161,561	147	154,082	142,017	115	116,386
Portland	to_Los-Angeles_RealTime	128,076	58	119,651	115,782	54	100,504
Los-Angeles	to_Portland_RealTime	114,563	55	99,442	129,121	60	120,640
Mexico-City	to_Osaka_Interactive	0	0	4	0	0	4
Toronto	to_Frankfurt_DefaultOverlay	0	0	4	0	0	4
Dallas	to_Albuquerque_RealTime	0	0	4	0	0	4
Seoul	to_Singapore_Interactive	0	0	4	0	0	4
Pittsburgh	to_Portland_Interactive	0	0	4	0	0	4
Mexico-City	to_San-Jose_Interactive	0	0	4	0	0	4

## DRC Bandwidth Trends

*Monitoring > [Bandwidth > Tunnels] DRC Trends*

The **DRC Bandwidth Trends** tab shows Dynamic Rate Control statistics over time.

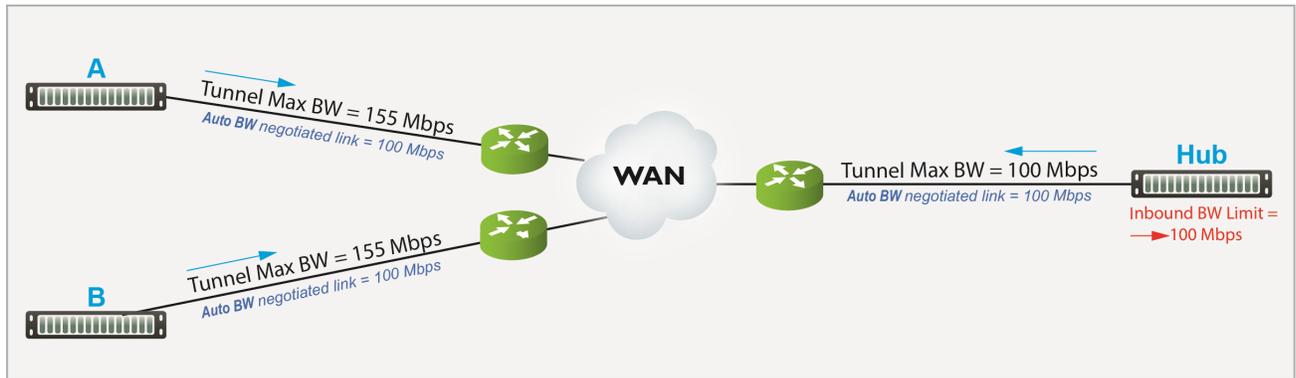
Dynamic Rate Control allows the Hub to regulate the tunnel traffic by lowering each remote appliance's **Tunnel Max Bandwidth**. The smallest possible value is that appliance's **Tunnel Min (imum) Bandwidth**.



## 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 the 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.

## Flows - Active & Recent

*Monitoring > Bandwidth > Flows > Active & Recent Flows*

The **Flows** tab allows you to view, filter, and manage flows for all your appliances. This tab also generates the Active & Recent Flows report, with or without filtering. This report retrieves the maximum number of most recent flows that are evenly distributed among the selected appliances.

Field	Definition
<b>Application</b>	Includes built-in applications, custom applications, and user-created application groups. Select the text field and a list displays. Choose the application you want to apply to your flow or enter the exact application you want to apply.
<b>App Group</b>	Includes the application group created by the user. Select the text field and a list displays. Choose the application group you want to apply to your flow or enter the exact application group you want to apply.
<b>Domain</b>	Includes the domain you can specify to filter your flow. Use the format, <i>*.domain.*</i> or <i>*.domain.[com, info, edu, org, net, ...]</i> . Select the text field and list displays. Choose the domain you want to apply.
<b>Protocol</b>	You can specify the protocol you want to apply to your filter. Select the text field and a list displays. You can select all or specify an individual protocol to apply.
<b>IP/Subnet</b>	This shows the flows that match both SRC IP and DEST IP as the two endpoints if SRC:DEST is enabled. If not enabled, all sources will appear when the filter is applied. You can apply this filter by clicking <b>Enter</b> , without selecting the Apply button if you want to do so.
<b>Port</b>	This displays ports with SRC and DEST as the two endpoints if SRC:DEST is enabled. If not enabled, all ports will appear when the filter is applied.
<b>Zone</b>	You can filter flows to the desired firewall zone. Select the text field and a list displays. Select the text field and a list displays. If <b>From:To</b> checkbox is not enabled, flows are filtered from and to the specified zone. If the checkbox is enabled, the flows are filtered from both the filtered <b>From:To</b> zones.
<b>VLAN</b>	Identifies the Virtual Local Area Network of a packet. Enter the <b>VLAN ID</b> you want to apply to your flow in the text field.
<b>DSCP</b>	Select the desired DSCP from the list. You can choose any or a specified DSCP from the list.
<b>Overlay</b>	The overlay the flow are applied. Overlays are defined in the <b>Business Intent Overlay</b> tab.
<b>Transport</b>	Select any of the three transport types: <b>SD-WAN</b> , <b>Breakout</b> , and <b>Underlay</b> . You can also apply a third-party service in this column if you have one configured.

Field	Definition
<b>Flow Characteristics</b>	You can apply any of the following flow characteristics to your flow: <b>Boosted, Directly Attached, Pass-Through, Stale, Route Dropped, Firewall Dropped, Asymmetric, and Slow Devices</b> . Note: You can only select one flow characteristic at a time.
<b>Include EdgeHA</b>	If not selected, Edge HA flows are excluded (default). If selected, the flows between Edge HA will be included.
<b>Include Built-In</b>	Includes the built-in policy flows. If not selected, they are excluded (default). If selected, they will be included.
<b>Active/Ended</b>	You can select if you want to apply an active or ended flow to as a filter. If selected, you can designate the started or ended time of the flow in the drop down. If <b>Custom</b> is selected from the date widgets will be enabled to specify an exact time frame.
<b>Slow Devices</b>	For debugging. A slow device is one that cannot receive data quickly enough from the Silver Peak appliance. This causes the appliance to expend too many resources for this device, at the expense of accelerating other devices. To counteract this, disable TCP acceleration for the slow devices in the Optimization Policy.
<b>Duration</b>	Shows flows that have lasted through a specific time frame. You can select < (less than) or > (greater than), and enter a specific duration (in minutes).
<b>Bytes</b>	You can specify whether you want to filter flows that have transferred their total bytes or within the last five minutes.
<b>Filter</b>	This list has all the saved filters. When selected, the filter configurations are loaded. See more information below regarding the <b>Filter</b> option.

## Filter

You can configure specific filters in this field. Select the drop-down menu to see a list of default filters you can apply to your flows. Once configured, you can add, edit, or delete filters if you select the edit icon.

Complete the following steps to add a filter:

1. Select the **Edit** icon next to the Filter drop down.
2. Create a filter or select one from the list.
3. Select **+Add**.
4. Select **Save**.

You can also select the history tab with the two arrows next to the **Filter** field if you want to go back to a previously applied filter. A maximum of 20 previously applied filters can be saved.

## Reset or Reclassify Flows

- You can **Reclassify** or **Reset** [Selected / All Returned / All] flows:
  - **Resetting** the flow kills it and restarts it. It is service-affecting.
  - **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.
  - **Selected** flows are individually selected; **All Returned** results from filtering (up to the max number of returnable flows); and **All** refers to all flows, visible or not.
- To export the table as a .csv file, select **Export**.
- **Reduction (%)** refers to reduced WAN traffic, relative to a specific appliance:
  - Reduction (%) for **Outbound** traffic =  $100(\text{Received from LAN} - \text{Transmitted to WAN}) / \text{Received from LAN}$
  - Reduction (%) for **Inbound** traffic =  $100(\text{Transmitted to LAN} - \text{Received from WAN}) / \text{Transmitted to LAN}$
- Flow **Details** are primarily to assist Silver Peak in troubleshooting and debugging.
- To set the column visibility, right-click any header in the Flows table. This will allow you to hide or unhide any selected fields.
- You can also select, drag, and drop any of the columns in the table to the order you want.

The following table represents the values in the Flows report.

Field	Definition
Host Name	Host name of the flow.
Detail	Pop-up that gives more detail regarding the selected flow.
Chart	Displays a real-time flow bandwidth chart with outbound/inbound traffic.
Uptime	The amount of time the flow existed.
Overlay	Name of the overlay to which the flow belongs.
Application	Includes built-in applications, custom applications, and user-created application groups.
Protocol	For selecting an individual or <b>All</b> protocols.
IP1	The IP1 address.
Port 1	The Port 1 address.
IP2	The IP2 address.
Port 2	The Port 2 address.
Inbound Bytes	Traffic received from the WAN.
Outbound Bytes	Traffic received from the LAN.
Inbound Tunnel	The name of the tunnel receiving traffic from the WAN.
Outbound Tunnel	The name of the tunnel receiving traffic from the LAN.

Field	Definition
From Zone	The zone configured by the flow's source endpoint.
To Zone	The zone configured by the flow's destination endpoint.
DSCP LAN RX	The DSCP marking traffic received from the LAN.

# Appliance Flow Counts

*Monitoring > [Bandwidth > Flows] Counts*

The **Appliance Flow Counts** 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 Flow Counts ×

Count Bytes 1hr 4hr 1d 7d Custom Range 12-May-18 17:00 — 13-May-18 17:00 Optimized Traffic Export ↻

Appliance Flow Counts ?

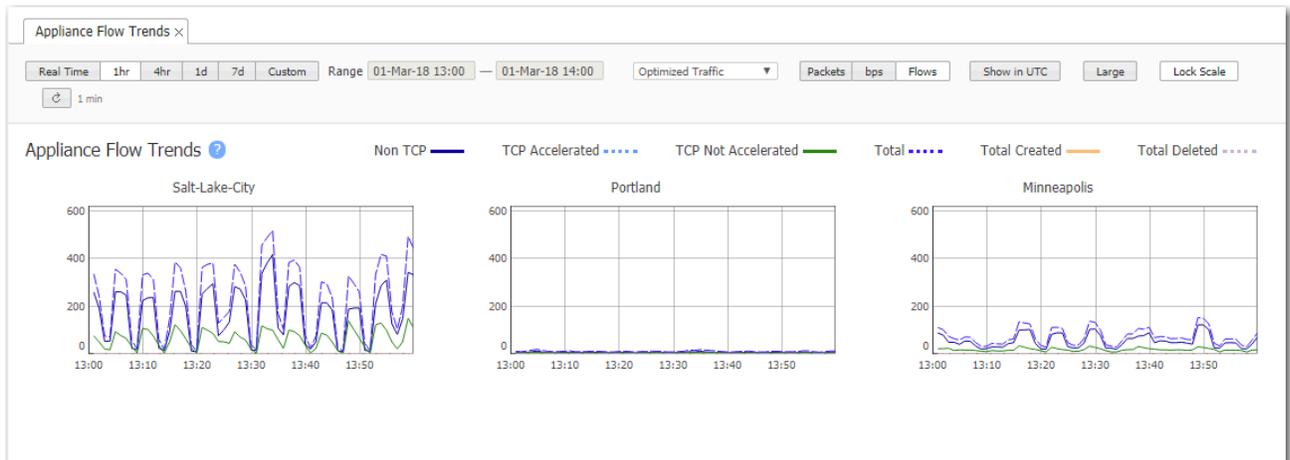
4 Rows Search

Appliance	TCP Accelerated				TCP Unaccelerated				Non TCP			
	Max	Avg	Created	Deleted	Max	Avg	Created	Deleted	Max	Avg	Created	Deleted
Los-Angeles	0	0	0	0	322	150	7,527	7,427	497	31	1	1
Minneapolis	0	0	0	0	9	2	0	0	20	5	0	0
Portland	0	0	0	0	165	140	7,475	7,368	2	0	1	2
Salt-Lake-City	0	0	0	0	282	12	0	0	924	31	0	0

# Appliance Flow Trends

Monitoring > [Bandwidth > Flows] Trends

The **Appliance Flow Trends** charts shows the number of flows, packets, and bits/second through the appliance, over time. It also differentiates among TCP (accelerated and unaccelerated) flows and non-TCP flows.



# Tunnel Flow Counts

Monitoring > [Bandwidth > Flows] Tunnel Counts

The **Tunnel Flow Counts** 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 Flow Counts x

1hr 4hr 1d 7d Custom Range 14-May-18 17:00 — 14-May-18 18:00 Top 25 All Overlays Export ↻

Tunnel Flow Counts ?

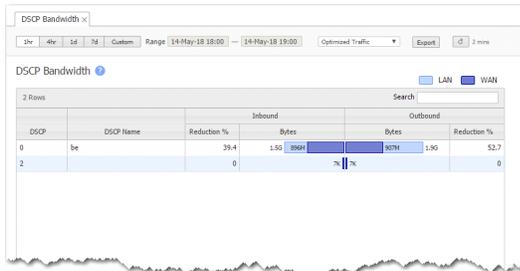
25 Rows Search

Appliance	Tunnel	TCP Accelerated		TCP Unaccelerated		Non TCP	
		Max	Avg	Max	Avg	Max	Avg
Los-Angeles	to_Portland_DefaultOverlay	0	0	65	65	0	0
Portland	to_Los-Angeles_DefaultOverlay	0	0	65	65	0	0
Los-Angeles	to_Portland_Interactive	0	0	55	54	0	0
Portland	to_Los-Angeles_Interactive	0	0	55	54	0	0
Los-Angeles	to_Portland_RealTime	0	0	12	12	0	0
Portland	to_Los-Angeles_RealTime	0	0	12	12	0	0
New-Orleans	to_Toronto_RealTime	0	0	0	0	0	0
Chicago	to_New-Orleans_RealTime	0	0	0	0	0	0
Boston	to_Portland_DefaultOverlay	0	0	0	0	0	0
San-Antonio	to_Chicago_DefaultOverlay	0	0	0	0	0	0
New-Orleans	to_Mexico-City_Interactive	0	0	0	0	0	0
San-Antonio	to_Boston_RealTime	0	0	0	0	0	0
San-Jose	to_Dallas_Interactive	0	0	0	0	0	0
Chicago	to_San-Antonio_RealTime	0	0	0	0	0	0

# DSCP Bandwidth

Monitoring > [Bandwidth > DSCP] Summary

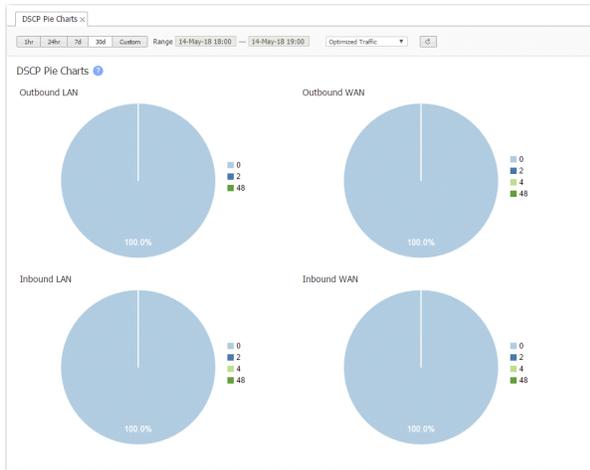
The DSCP Bandwidth chart shows which DSCP classes are sending the most data.



## DSCP Pie Charts

*Monitoring > [Bandwidth > DSCP] Pie Charts*

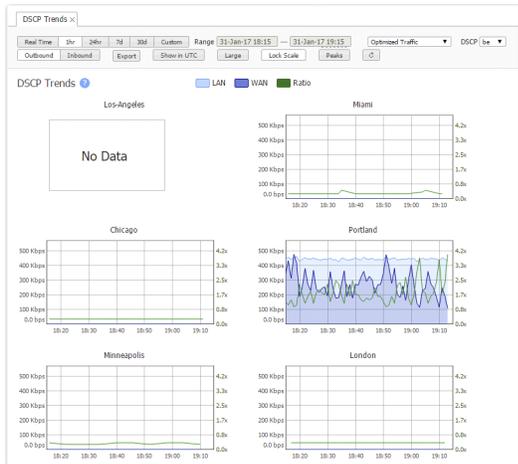
The **DSCP Pie Charts** show the proportion of traffic in each DSCP class. Hovering over the charts and the legends reveals additional information.



# DSCP Trends

Monitoring > [Bandwidth > DSCP] Trends

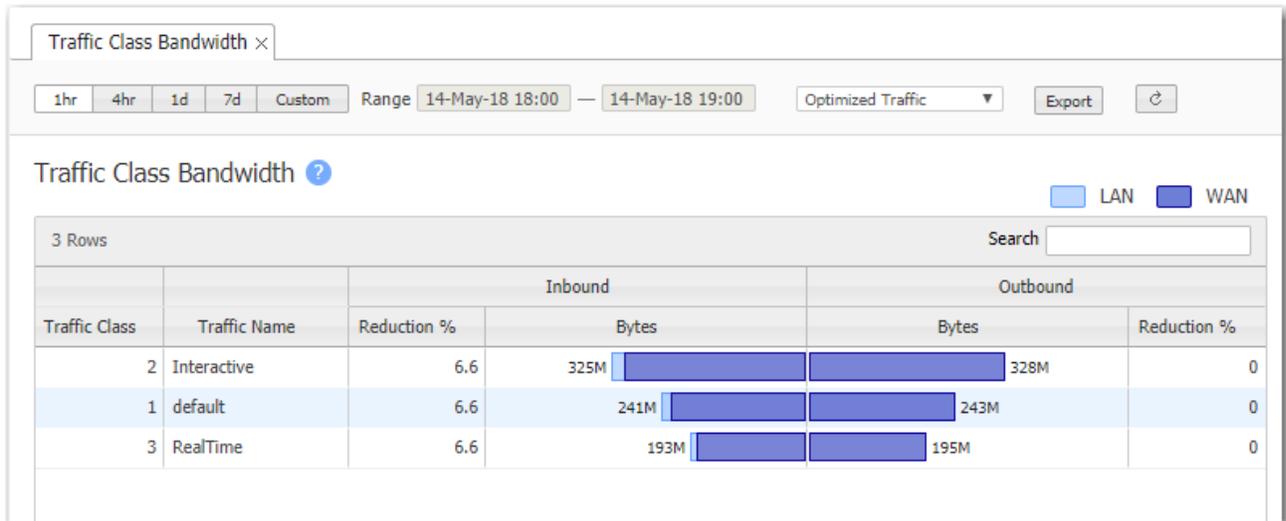
This tab shows DSCP usage over time.



# Traffic Class Bandwidth

Monitoring > [Bandwidth > QoS] Summary

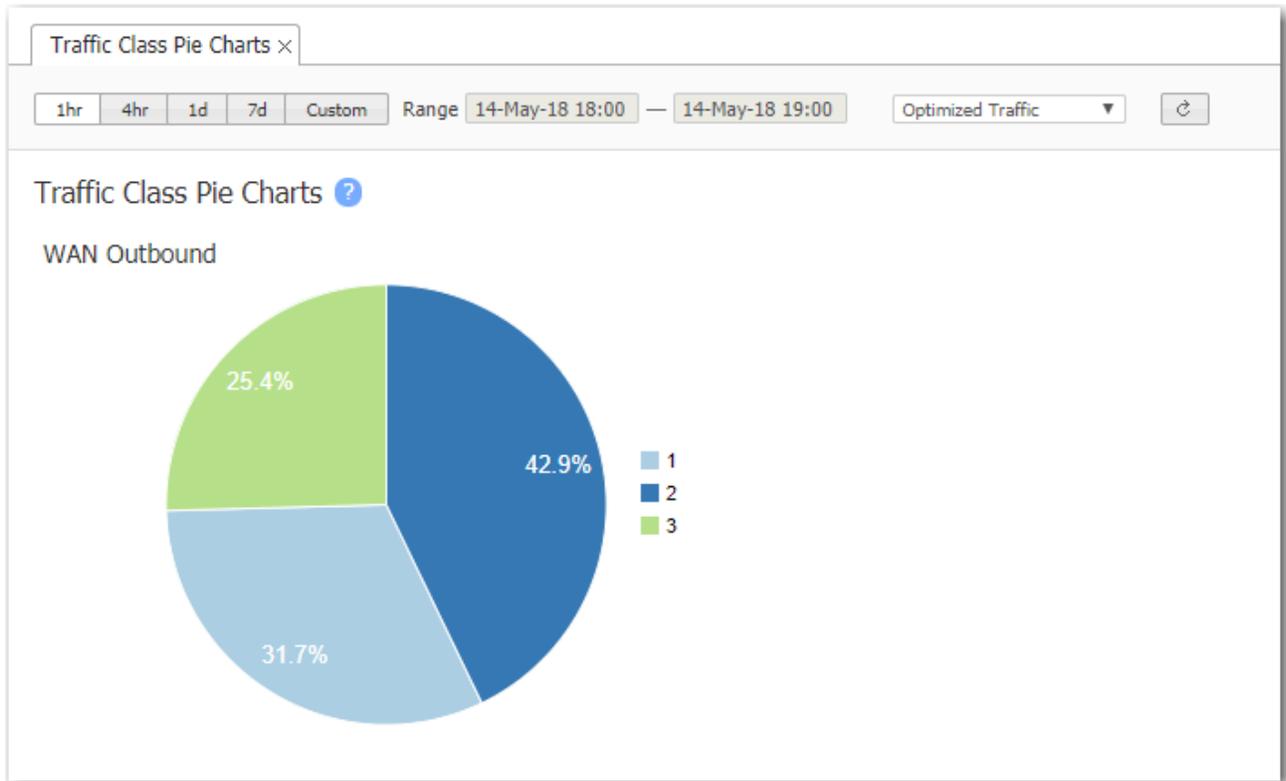
The **Traffic Class Bandwidth** chart shows which QoS traffic classes are sending the most data.



## Traffic Class Pie Charts

*Monitoring > [Bandwidth > QoS] Pie Charts*

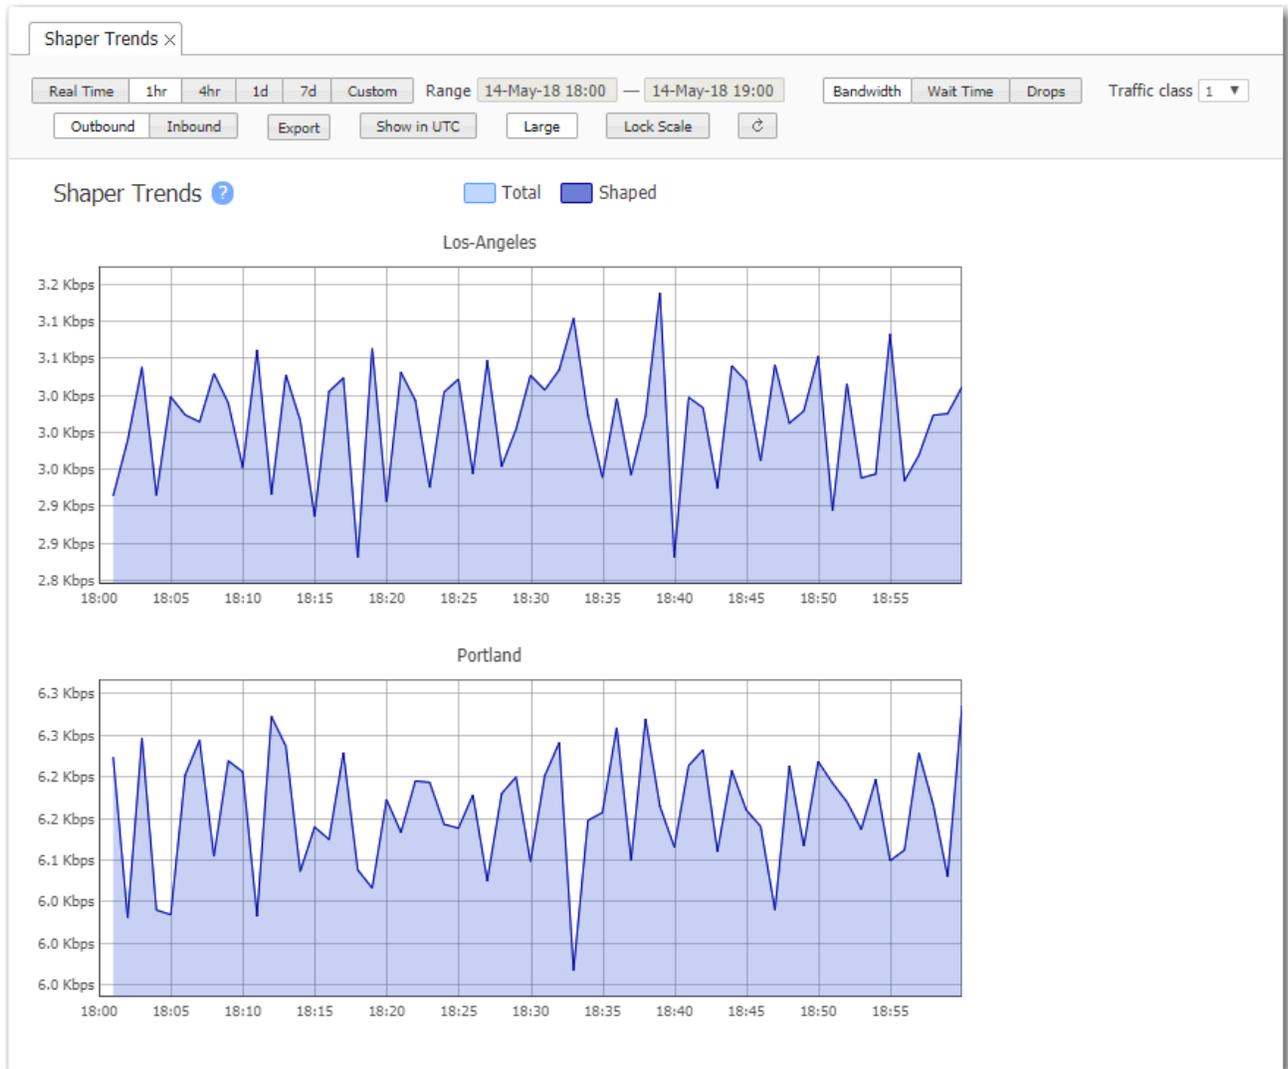
The **Traffic Class Pie Charts** show the proportion of traffic in each Traffic class. Hovering over the charts and the legends reveals additional information.



## QoS (Shaper) Trends

Monitoring > [Bandwidth > QoS] Trends

This tab shows how much bandwidth any traffic class uses over time.



## Works with Office 365

Ensure your overlays have the following options configured to preserve the Works with Office 365 default applications. The table below indicates the default overlays, applications, and preferred policy order configured in the **Business Intent Overlays** tab within Orchestrator. The overlay name indicated in the table below is the default that ships with Orchestrator. This can be modified with user configuration.

**NOTE** However, Skype for Business, SharePoint Online, and Office 365 Exchange **must** break out locally.

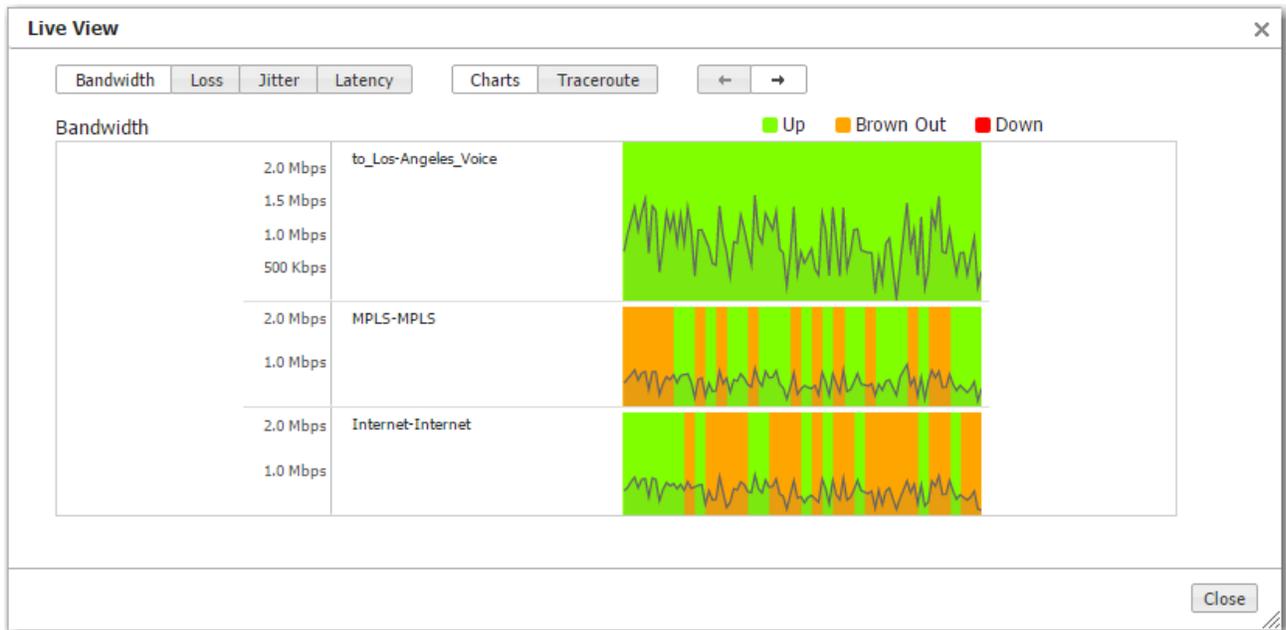
Overlay	Application	Preferred Policy Order (Breakout Traffic to Internet & Cloud Services)	What it Matches
Real-Time	Skype for Business		<ul style="list-style-type: none"> <li>Microsoft Office 365 <b>Optimize</b> and <b>Allow</b> categories for the respective applications.</li> </ul>
CriticalApps	SharePoint Online, Office 365 Exchange		<ul style="list-style-type: none"> <li>Matches Microsoft Office 365 <b>Default</b> categories</li> <li>Office365Common applications</li> </ul> <p><b>NOTE</b> Do not specify other individual Office applications in this group or overlay.</p>
Default	For everything	Any policy order except "Drop"	<ul style="list-style-type: none"> <li>Matches Microsoft Office 365 <b>Default</b> categories</li> <li>Office365Common applications</li> </ul> <p><b>NOTE</b> Do not specify other individual Office applications in this group or overlay.</p>

For more information regarding Works with Office 365 applications, navigate to <https://techcommunity.microsoft.com> and search for the Office 365 blog.

## Live View

*Monitoring > [Tunnel Health] Live View*

Live View shows the live bandwidth, loss, latency, and jitter on all the tunnels. For an overlay, it also shows live tunnel states – **Up**, **Browned Out**, or **Down**.

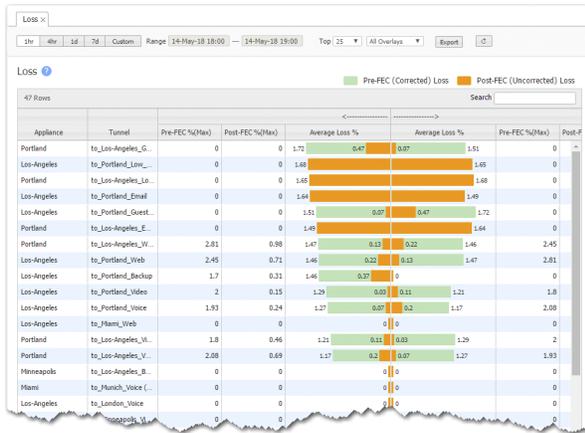


In real-time, LiveView shows how Silver Peak creates synergy to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and delivering consistent application performance while both underlays are in persistent brown-out state.

# Loss

Monitoring > [Tunnel Health > Loss] Summary

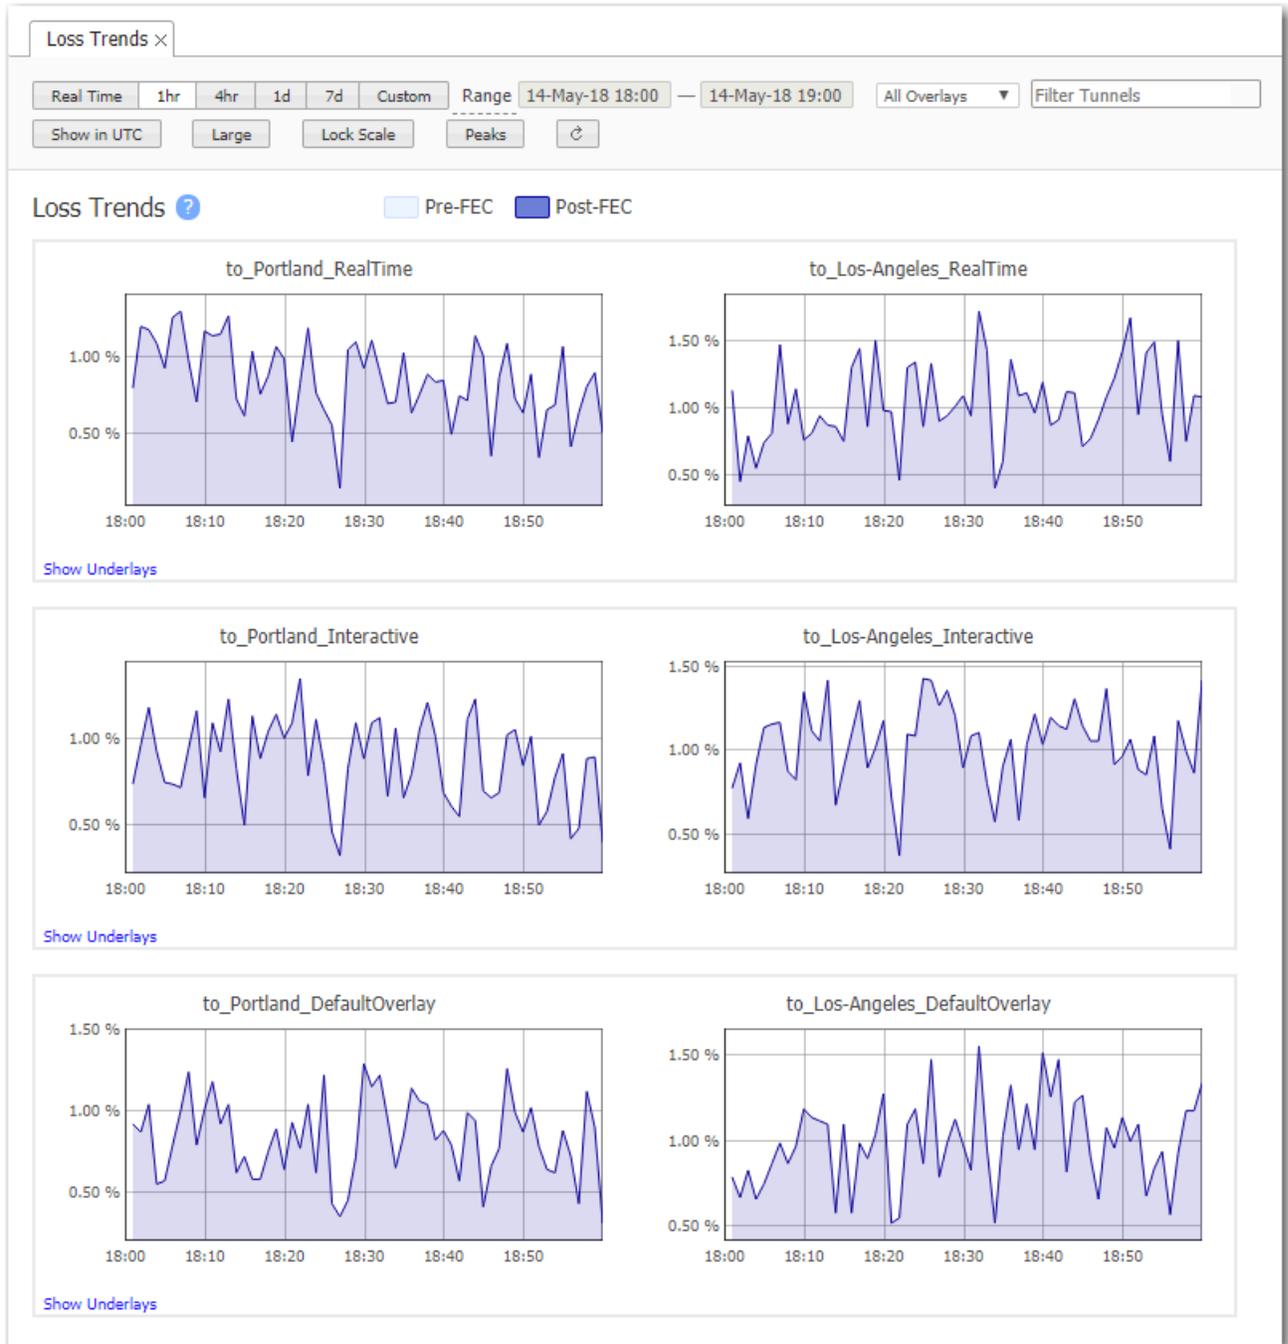
The **Loss** chart shows which tunnels have the most dropped packets.



## Loss Trends

*Monitoring > [Tunnel Health > Loss] Trends*

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



**NOTE** Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

# Jitter Summary

*Monitoring > Tunnel Health > Jitter] > Summary*

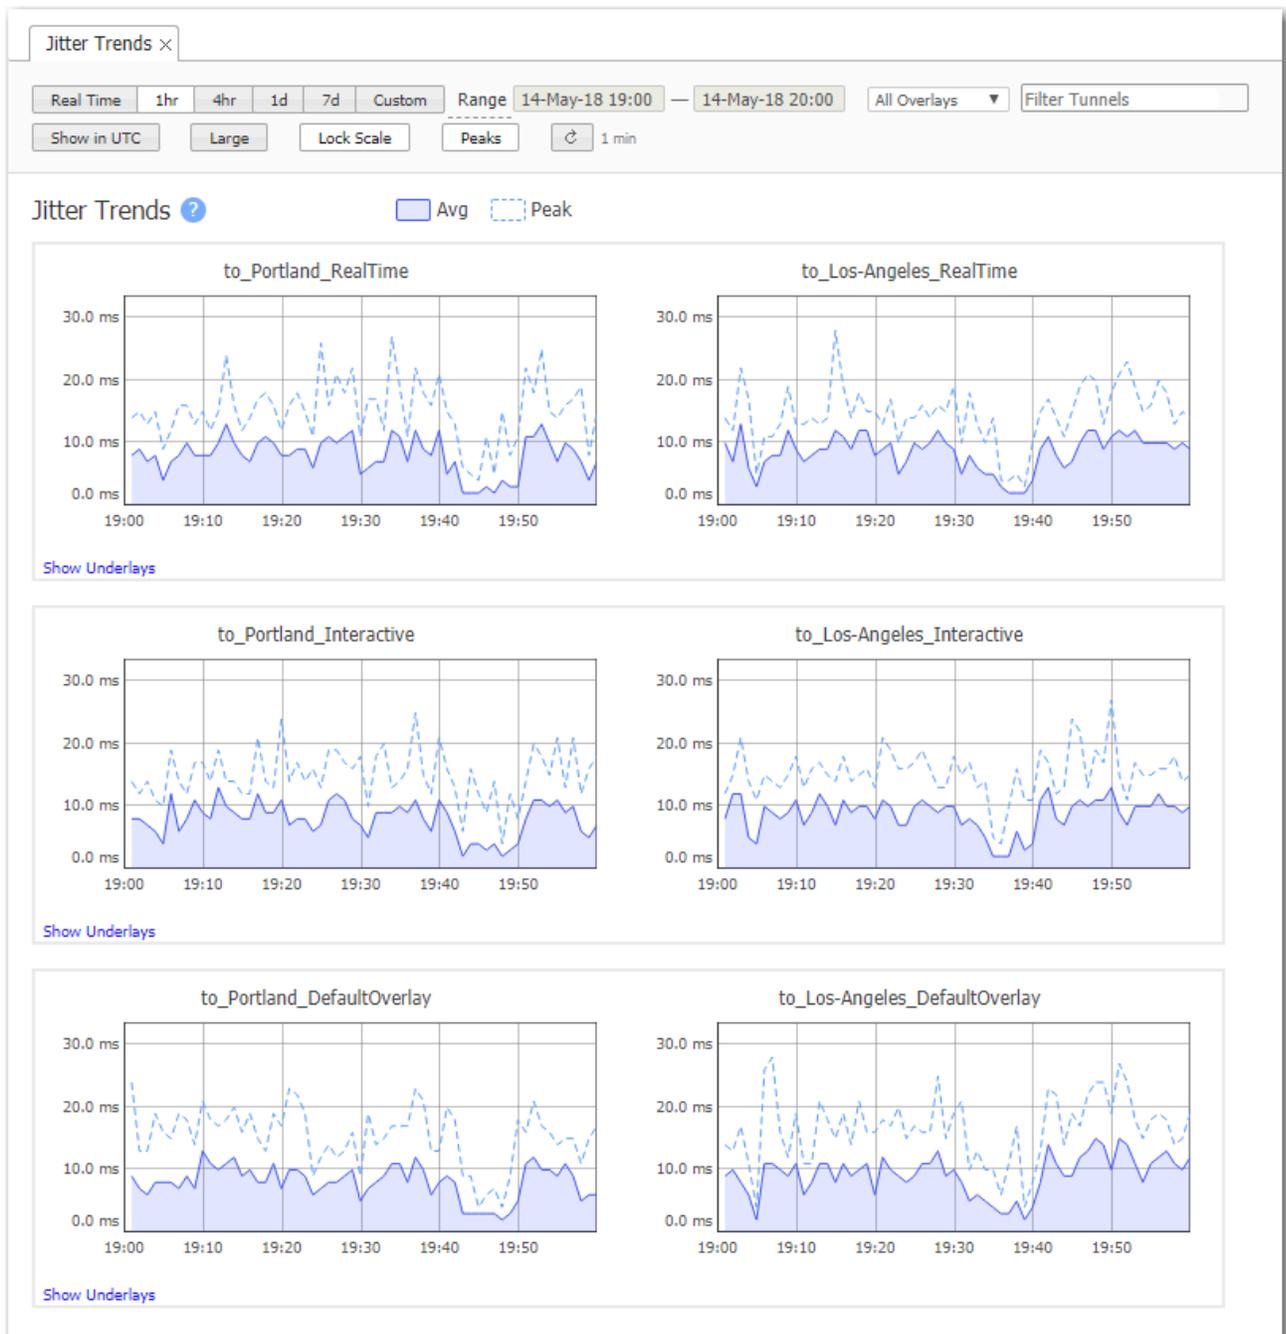
The **Jitter** chart shows which tunnels have the most Jitter. Jitter can be caused by congestion in the LAN, firewall routers, bottleneck access links, load sharing, route flapping, routing table updates, and timing drifts.

Appliance	Tunnel	Max Jitter (ms)	Average Jitter (ms)	Average Jitter (ms)	Max Jitter (ms)	Remote Tunnel	Remote Appliance
Munich	to_Miami_Guest_WFI	318	1	0	0	8 to_Munich_Guest_WFI	Miami
Portland	to_Los-Angeles_Email	36	1	0	0	37 to_Portland_Email	Los-Angeles
Portland	to_Los-Angeles_Video	36	1	0	0	37 to_Portland_Video	Los-Angeles
Portland	to_Los-Angeles_Gold...	36	1	0	0	37 to_Portland_Gold_CI	Los-Angeles
Portland	to_Los-Angeles_Ever...	36	1	0	0	37 to_Portland_Everv...	Los-Angeles
Portland	to_Los-Angeles_Voice	36	1	0	0	37 to_Portland_Voice	Los-Angeles
Portland	to_Los-Angeles_Gues...	36	1	0	0	37 to_Portland_Guest...	Los-Angeles
London	to_Miami_Guest_WFI	9	1	0	0	8 to_London_Guest_WFI	Miami
Chicago	to_Chennai_Guest_WFI	143	1	1	1	9 to_Chicago_Guest_WFI	Chennai
Munich	to_Minneapolis_Gues...	6	1	1	1	5 to_Munich_Guest_WFI	Minneapolis
Munich	to_Chicago_Guest_WFI	6	1	1	1	8 to_Munich_Guest_WFI	Chicago
Portland	to_Mumbai_Guest_WFI	8	1	0	0	37 to_Portland_Guest...	Mumbai
Minneapolis	to_Mumbai_Guest_WFI	6	1	0	0	7 to_Minneapolis_Gues...	Mumbai
Minneapolis	to_Chennai_Guest_WFI	6	1	1	1	9 to_Minneapolis_Gues...	Chennai
Portland	to_Chennai_Guest_WFI	10	1	1	1	19 to_Portland_Guest...	Chennai
London	to_Portland_Guest_WFI	10	1	1	1	8 to_London_Guest_WFI	Portland
Portland	to_Munich_Guest_WFI	10	1	1	1	14 to_Portland_Guest...	Munich
London	to_Minneapolis_Gues...	9	1	1	1	5 to_London_Guest_WFI	Minneapolis
Chennai	to_Miami_Guest_WFI	9	1	1	1	8 to_Chennai_Guest_WFI	Miami

## Jitter Trends

Monitoring > [Tunnel Health > Jitter] Trends

This tab shows tunnel jitter time.





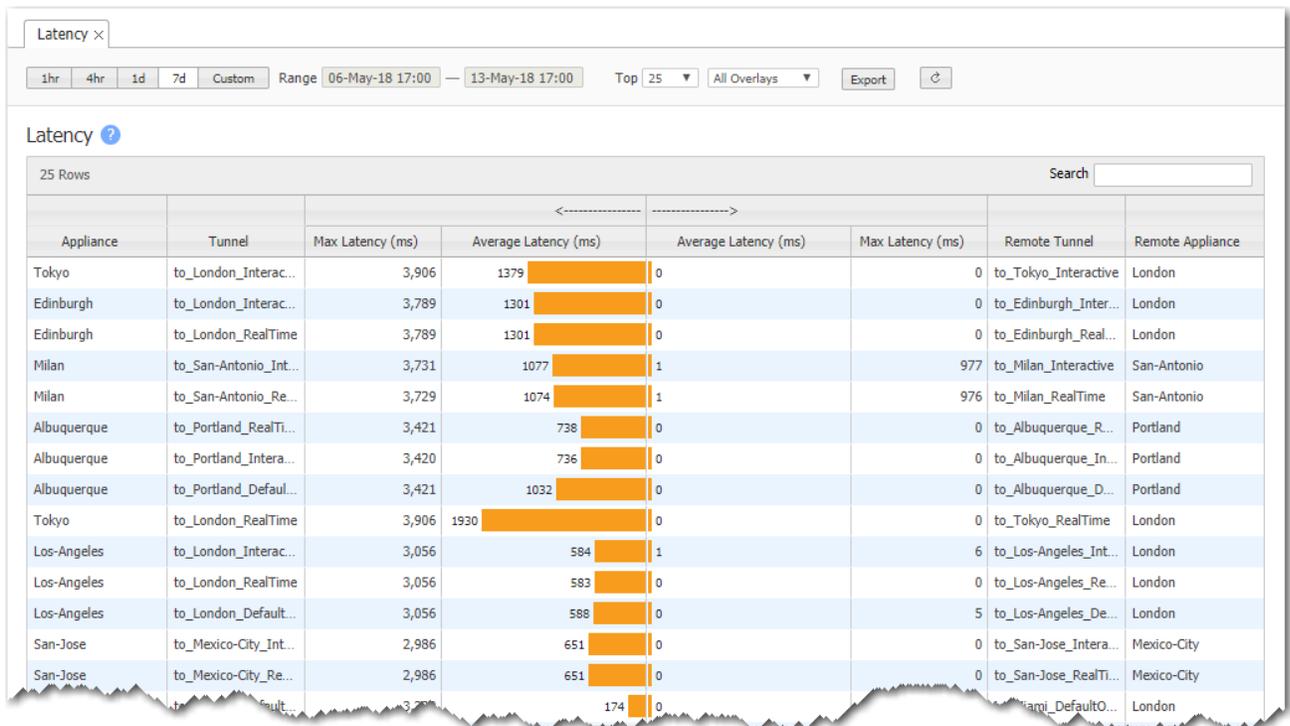
**NOTE** Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

---

# Latency

Monitoring > [Tunnel Health > Latency] Summary

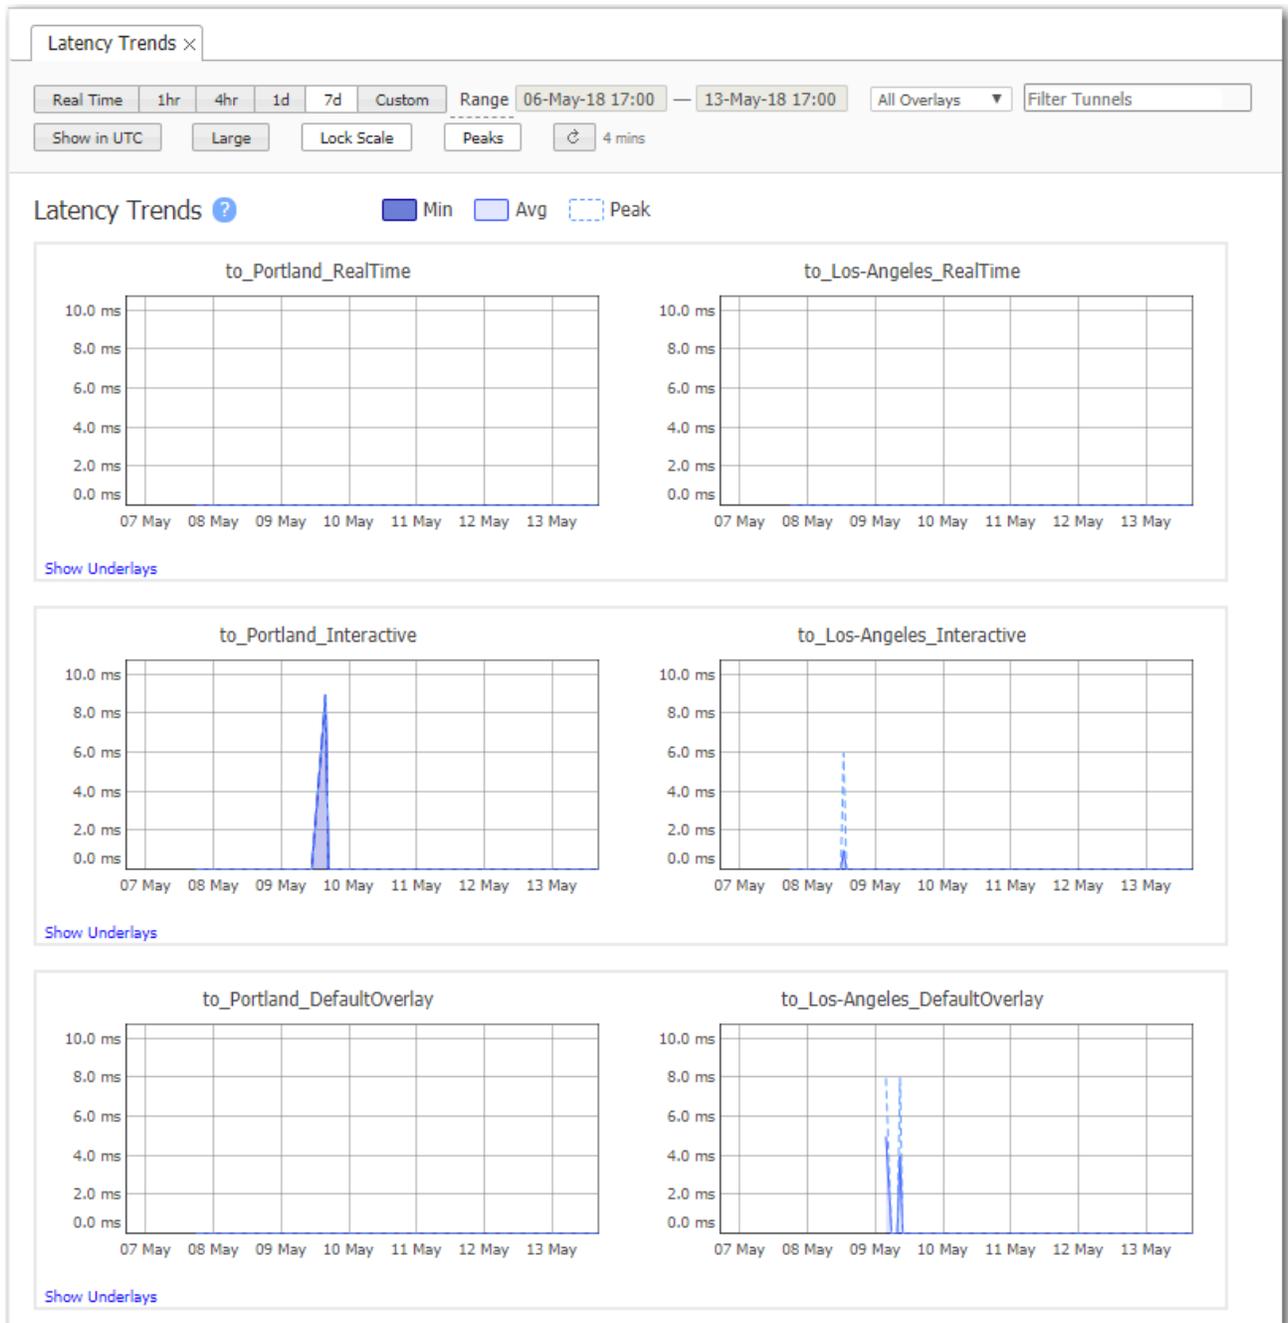
The **Latency** chart shows which tunnels have the most transmission delay, generally as a result of congestion.



# Latency Trends

Monitoring > [Tunnel Health > Latency] Trends

The Latency Trends chart shows tunnel latency over time.





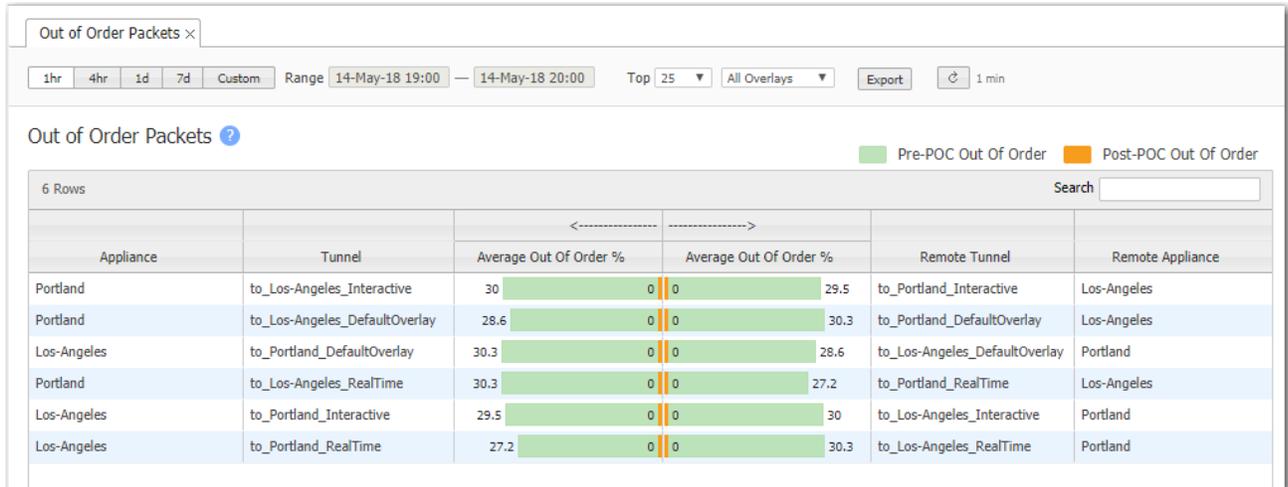
**NOTE** Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

---

# Out of Order Packets

Monitoring > [Tunnel Health > Out of Order Packets] Summary

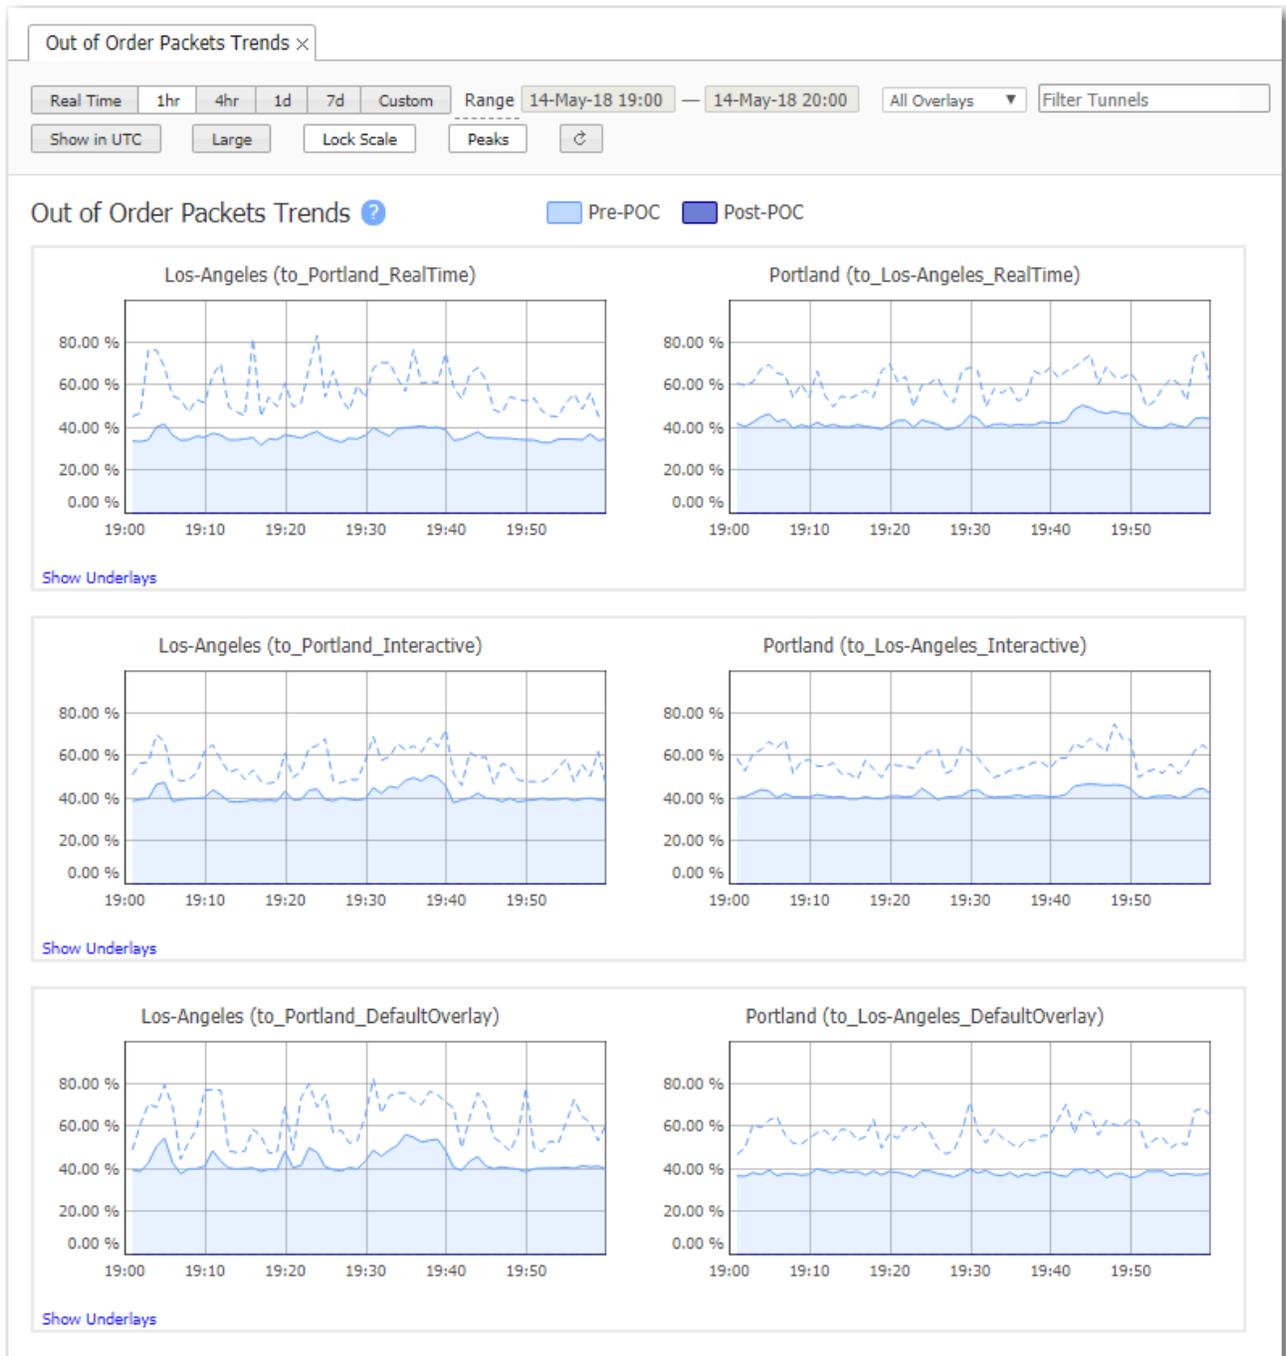
The **Out of Order Packets** chart shows which tunnels receive the most packets out of sequence relative to how they were sent.



# Out of Order Packets Trends

*Monitoring > [Tunnel Health > 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).

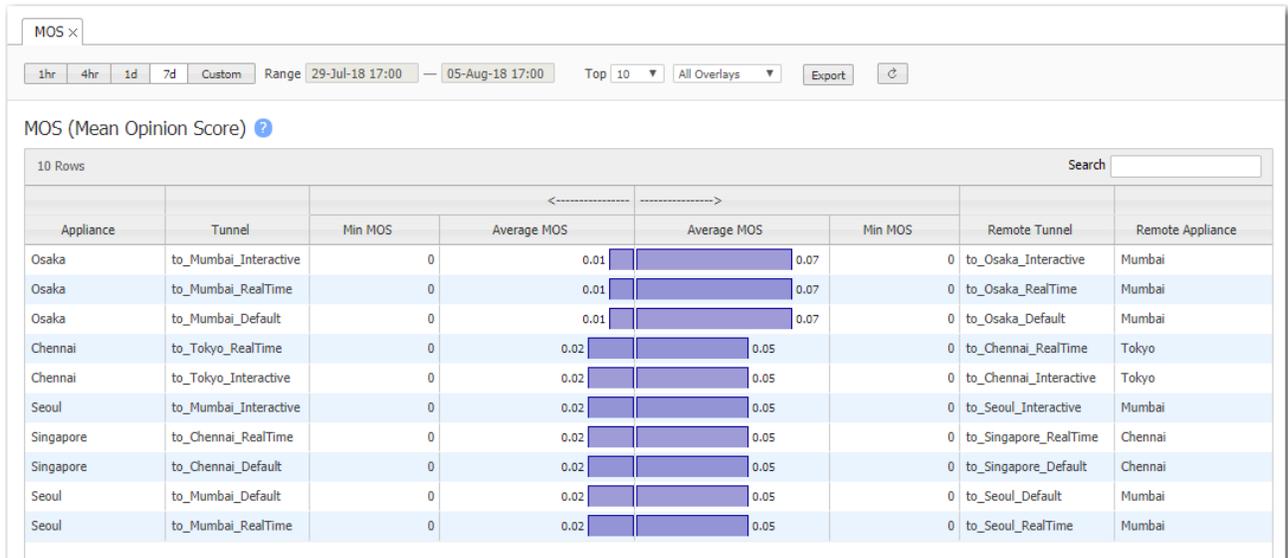


**NOTE** Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

## Mean Opinion Score (MOS) - Summary

*Monitoring > [Tunnel Health > MOS] Summary*

The Mean Opinion Score (MOS) is a commonly used measure for video, audio, and audiovisual quality evaluation. Perceived quality is rated on a theoretical scale of 1 to 5; the higher the number, the better the quality.

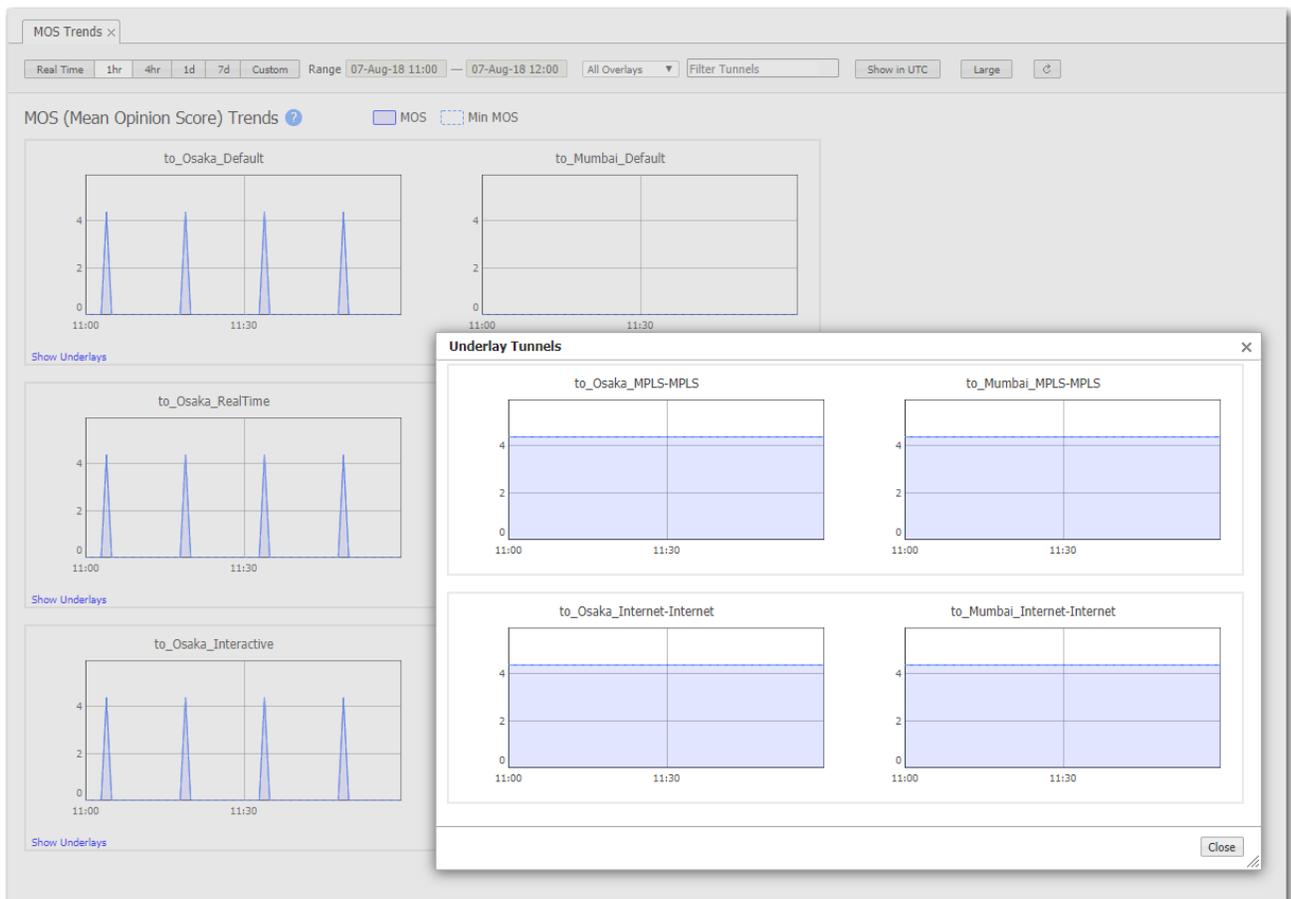


The value can be affected by loss, latency, and jitter. In practice, a value of **4.4** is considered an excellent quality target.

## Mean Opinion Score (MOS) Trends

*Monitoring > [Tunnel Health > MOS] Trends*

The Mean Opinion Score (MOS) is a commonly used measure for video, audio, and audiovisual quality evaluation. Perceived quality is rated on a theoretical scale of 1 to 5; the higher the number, the better the quality.



- The value can be affected by loss, latency, and jitter. In practice, a value of **4.4** is considered an excellent quality target.
- The **Min MOS** value reports the worst score within a minute.

# Tunnels Summary

Monitoring > [Tunnel Health > Other Tunnel Statistics] Tunnel Summary

This tab summarizes tunnel statistics, including reduction, throughput, latency, and packet loss.

The screenshot shows a web interface titled "Tunnels Summary" with a search bar and a table of 50 rows. The table columns are: Tunnel, Status, LAN, WAN, Inbound (Reduction %, LAN Through, WAN Through), WAN (LAN, WAN, Reduction %, LAN Through, WAN Through), Outbound (LAN, WAN, Reduction %, LAN Through, WAN Through), Packets Loss % (Avg, Max), Jitter (ms) (Avg, Max), and Latency (ms) (Avg, Max). The data rows show various tunnels such as "Portland: to\_Los...", "Portland: to\_Ma...", "Minneapolis: to...", "Portland: to\_Chi...", "Los-Angeles: to...", "Portland: to\_Ma...", "Minneapolis: to...", "Portland: to\_Los...", "Chicago: to\_Los...", "Portland: to\_Chi...", "Minneapolis: to...", "Portland: to\_Ma...", "Portland: to\_Ma...", "Portland: to\_Los...", "Chennai: to\_Mu...", "Portland: to\_Los...", "Minneapolis: to...", "Humbolt: to\_Che...", "London: to\_Che...", "Chennai: to\_Mu...", "Portland: to\_Los...", and "Chicago: to\_Ma...".

Tunnel	Status	Inbound				Outbound					Packets Loss %		Jitter (ms)		Latency (ms)		
		LAN	WAN	Reduction %	LAN Through	WAN Through	LAN	WAN	Reduction %	LAN Through	WAN Through	Avg	Max	Avg	Max	Avg	Max
Portland: to_Los...	up - active	2.9 MB	12 MB	0.00	0	962 bps	0	7.2 KB	0.00	390 Mbps	1.6 Mbps	0	0	72.00	32.00	50.02	0.00
Portland: to_Los...	up - active	198 MB	127 MB	35.96	0	0	0	0.00	26 Mbps	17 Mbps	0	0	72.00	32.00	50.02	0.00	
Portland: to_Ma...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	53.00	9.00	99.93	0.00	
Minneapolis: to...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	6.00	3.00	49.98	0.00	
Portland: to_Chi...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	38.00	11.00	99.92	0.00	
Los-Angeles: to...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	5.00	3.00	50.00	0.00	
Portland: to_Ma...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	51.00	9.00	99.93	0.00	
Minneapolis: to...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	59.00	7.00	99.63	0.00	
Portland: to_Los...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	72.00	32.00	50.02	0.00	
Chicago: to_Los...	up - active	0	0	0.00	0	0	0	0.00	0	0	0	0	7.00	17.00	49.97	0.00	
Portland: to_Chi...	up - active	0	5.0 KB	0.00	0	516 bps	0	3.9 KB	0.00	0	671 bps	0	0	38.00	11.00	99.92	0.00
Minneapolis: to...	up - active	0	4.8 KB	0.00	0	657 bps	0	4.9 KB	0.00	0	646 bps	0	0	6.00	3.00	49.98	0.00
Portland: to_Ma...	up - active	0	5.7 KB	0.00	0	587 bps	0	4.4 KB	0.00	0	760 bps	0	0	51.00	9.00	99.93	0.00
Portland: to_Ma...	up - active	0	5.7 KB	0.00	0	587 bps	0	4.4 KB	0.00	0	760 bps	0	0	53.00	9.00	99.93	0.00
Portland: to_Los...	up - active	0	5.5 KB	0.00	0	645 bps	0	4.8 KB	0.00	0	735 bps	0	0	72.00	32.00	50.02	0.00
Chennai: to_Mu...	up - active	0	5.7 KB	0.00	0	704 bps	0	5.3 KB	0.00	0	760 bps	0	0	0.00	1.00	0.00	0.00
Portland: to_Los...	up - active	8.3 KB	27 KB	0.00	0	681 bps	0	5.1 KB	0.00	1.1 Mbps	3.6 Mbps	0	0	72.00	32.00	50.02	0.00
Minneapolis: to...	up - active	0	4.8 KB	0.00	0	458 bps	0	3.4 KB	0.00	0	633 bps	0	0	6.00	3.00	49.98	0.00
Humbolt: to_Che...	up - active	0	5.7 KB	0.00	0	704 bps	0	5.3 KB	0.00	0	760 bps	0	0	0.00	0.00	0.00	0.00
London: to_Che...	up - active	0	5.7 KB	0.00	0	704 bps	0	5.3 KB	0.00	0	760 bps	0	0	0.00	5.00	0.00	0.00
Chennai: to_Mu...	up - active	0	5.7 KB	0.00	0	704 bps	0	5.3 KB	0.00	0	760 bps	0	0	0.00	1.00	0.00	0.00
Portland: to_Los...	up - active	0	5.4 KB	0.00	0	669 bps	0	5.0 KB	0.00	0	722 bps	0	0	72.00	32.00	50.02	0.00
Chicago: to_Ma...	up - active	0	5.7 KB	0.00	0	645 bps	0	4.8 KB	0.00	0	760 bps	0	0	62.00	6.00	99.82	0.00

For each Business Intent Overlay, the Link Bonding Policy specified determines the bandwidth efficiency. To guarantee service quality levels, High Availability requires the most overhead, and High Efficiency requires the least. The table shows the total bandwidth used. The Payload filter removes overhead from the displayed values.

# Appliance Administration Tabs

These menus are related to appliance administration. They include general settings, software management, and tools for troubleshooting and maintenance.

## Appliance User Accounts Tab

Administration > [General Settings > Users & Authentication] Users

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

The screenshot shows the 'Users' tab interface. At the top, there is a 'Users x' tab, a 'Manage Users with Templates' link, an 'Export' button, and a refresh icon. Below this is the 'User Accounts' section with a '34 Rows' indicator and a search box. The main table has the following columns: 'Edit', 'Appliance Name', 'User Name', 'Capability', and 'Enabled'. The table lists 16 rows of data for various appliances, each with an 'admin' and a 'monitor' user account.

Edit	Appliance Name	User Name	Capability	Enabled
✎	Albuquerque	admin	admin	Yes
✎	Albuquerque	monitor	monitor	No
✎	Boston	admin	admin	Yes
✎	Boston	monitor	monitor	No
✎	Chicago	admin	admin	Yes
✎	Chicago	monitor	monitor	No
✎	Dallas	admin	admin	Yes
✎	Dallas	monitor	monitor	No
✎	Denver	admin	admin	Yes
✎	Denver	monitor	monitor	No
✎	Los-Angeles	admin	admin	Yes
✎	Los-Angeles	monitor	monitor	No
✎	Mexico-City	admin	admin	Yes
✎	Mexico-City	monitor	monitor	No
✎	Miami	admin	admin	Yes

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).

- User Names are case-sensitive.
- The table lists all users known to the appliances, whether or not their accounts are enabled.

## Auth/RADIUS/TACACS+ Tab

*Administration > General Settings > Users & Authentication > 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

#### *Authentication and Authorization Fields*

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

### RADIUS and TACACS+

#### *RADIUS and TACACS+ Server Fields*

Field	Definition
<b>Auth Port</b>	For RADIUS, the default value is <b>1812</b> . For TACACS+, the default value is <b>49</b> .
<b>Auth Type</b>	[TACACS+] The options are <b>pap</b> or <b>ascii</b> .
<b>Enabled</b>	Whether or not the server is enabled.
<b>Retries</b>	The number of attempts allowed before lockout.
<b>Server Type</b>	RADIUS or TACACS+

Field	Definition
<b>Timeout</b>	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 <b>Session Management template</b> .

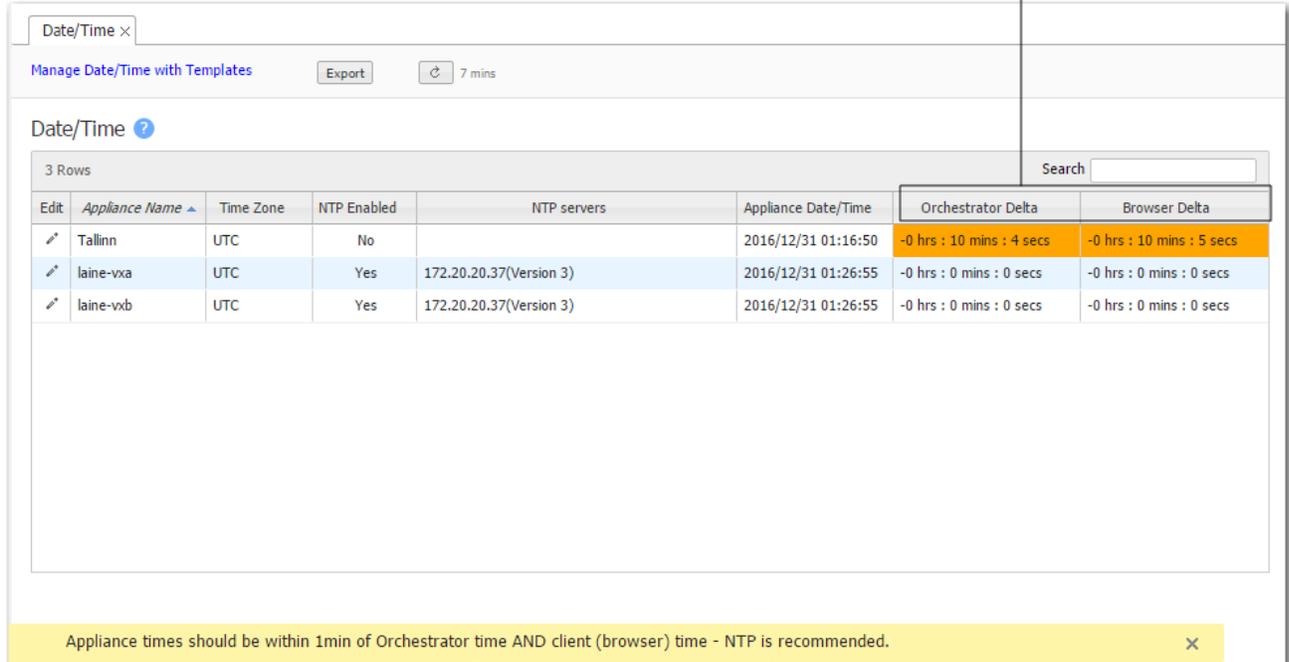
---

## Date/Time Tab

Administration > [General Settings > Setup] Date/Time

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

Relative to the appliance's  
configured time



Date/Time x

Manage Date/Time with Templates Export 7 mins

Date/Time ?

3 Rows Search

Edit	Appliance Name	Time Zone	NTP Enabled	NTP servers	Appliance Date/Time	Orchestrator Delta	Browser Delta
✎	Talinn	UTC	No		2016/12/31 01:16:50	-0 hrs : 10 mins : 4 secs	-0 hrs : 10 mins : 5 secs
✎	laine-vxa	UTC	Yes	172.20.20.37(Version 3)	2016/12/31 01:26:55	-0 hrs : 0 mins : 0 secs	-0 hrs : 0 mins : 0 secs
✎	laine-vxb	UTC	Yes	172.20.20.37(Version 3)	2016/12/31 01:26:55	-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. x

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.

## DNS (Domain Name Servers) Tab

*Administration > General Settings > Setup > DNS*

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

The screenshot shows the DNS configuration interface. At the top, there are tabs for 'Topology' and 'DNS x'. Below the tabs, there is a 'Manage DNS with Templates' link, an 'Export' button, and a refresh icon. The main section is titled 'DNS ?' and contains a table with 5 rows. The table has columns for 'Edit', 'Appliance Na...', 'Primary DNS IP addr', 'Secondary DNS IP addr', 'Tertiary DNS IP addr', and 'Domain Names'. The 'Los-Angeles' appliance has a Primary DNS IP address of 1.1.1.1, while the others have no settings defined.

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 Settings > Setup] SNMP

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

SNMP						Trap Receivers		
Edit	Appliance Name	Enable SNMP Agent	Enable SNMP Traps	Enable SNMP V1/V2	Enabled V3 Users	Trap Receiver 1	Trap Receiver 2	Trap Receiver 3
✎	Tallinn	Yes	Yes	Yes				
✎	laine-vxa	Yes	Yes	Yes				
✎	laine-vxb	Yes	Yes	Yes				

- 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.

### SNMP Settings

Field Name	Description
Enable SNMP	Allows the SNMP application to poll this Silver Peak appliance. (For SNMP v1 and SNMP v2c)

Field Name	Description
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 <b>SNMP v3</b> , instead of using <b>v1</b> or <b>v2c</b> . This provides a way to authenticate without using clear text.
Trap Receiver	IP address of a host configured to receive SNMP traps

## Flow Export Tab

Administration > [General Settings > Setup] NetFlow

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

The screenshot displays the 'Flow Export' configuration page. At the top, there's a search bar and a 'Manage Flow Export with Templates' button. Below is a table with 30 rows, each representing an appliance. The table columns are: Edit, Appliance Name, Flow Exporting, Active Flow Timeout, IPFIX Template Timeout, and Traffic Type. A modal dialog for 'Albuquerque' is open, showing the following configuration:

- Flow Export Configuration:**
  - Enable Flow Exporting:
  - Active Flow Timeout: 1 (1..30) mins
  - IPFIX Template Timeout: 10 (1..1440) mins
  - Traffic Type:
    - Outbound WAN
    - Inbound WAN
    - Inbound LAN
    - Outbound LAN
- Collectors:**
  - Buttons: Add, Collectors, IP Address, Port, Collector Type

- 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 and IPFIX 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/IPFIX 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 **WAN TX**.

## Logging Tab

Administration > [General Settings > Setup] Logging

This tab summarizes the configured logging parameters:

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

The screenshot shows the 'Logging' tab in the Silver Peak Orchestrator interface. At the top, there are tabs for 'Topology' and 'Logging x'. Below the tabs, there is a 'Manage Logging with Templates' link, an 'Export' button, and a refresh icon. The main content area is titled 'Logging' with a help icon. It displays a table with 5 rows and 8 columns. The table is divided into two main sections: 'Log Configuration' and 'Log Facilities Configuration'. The 'Log Configuration' section includes columns for 'Minimum Severity', 'Log File Size Threshold', and 'Number of Logs to Ke...'. The 'Log Facilities Configuration' section includes columns for 'System', 'Audit', and 'Flow'. The table lists five appliances: Tallinn, laine-vxa, laine-vxb, laine2-vxa, and laine2-vxb, all with a 'Notice' severity level, a '50' file size threshold, and '30' logs to keep. The facilities configuration for all appliances is 'local1' for System, 'local0' for Audit, and 'local2' for Flow.

		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

## Severity Levels

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

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

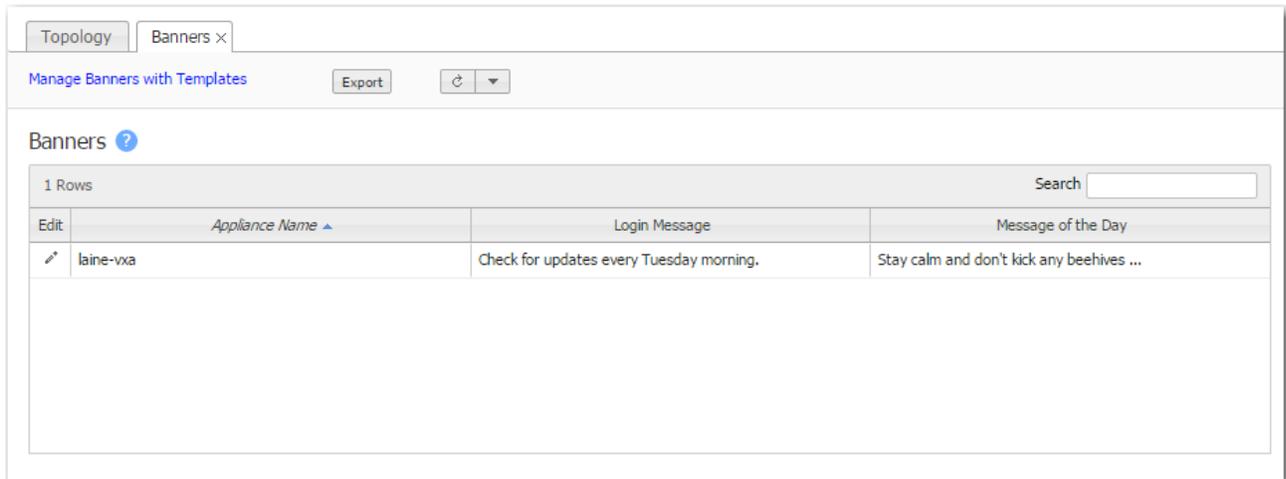
## Remote Logging

- 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**).

## Banners Tab

*Administration > [General Settings > Setup] Banners*

This tab lists the **banner messages** on each appliance.



1 Rows				Search <input type="text"/>	
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 ...		

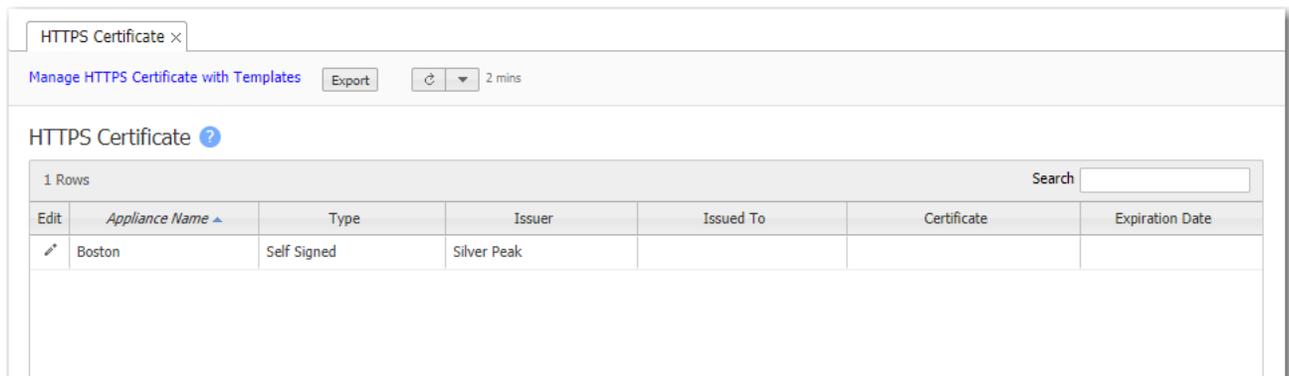
- The **Login Message** appears before the login prompt.
- The **Message of the Day** appears after a successful login.

## HTTPS Certificate Tab

*Administration > [General Settings > Setup] HTTPS Certificate*

The VXOA software includes a self-signed certificate that secures the communication between the user's browser and the appliance.

You also have the option to install your own custom certificate, acquired from a CA certificate authority.



**For a custom certificate, to use with a specific appliance:**

1. Consult with your IT security team to generate a certificate signing request (CSR), and submit it to your organization's chosen SSL Certificate Authority (CA).

Examples of Certificate Authorities include GoDaddy, Verisign, Comodo, Symantec, Microsoft Entrust, GeoTrust, etc.

- For a list of what Silver Peak supports, see [Silver Peak Security Algorithms](#).
  - All certificate and key files must be in **PEM** format.
2. After the Certificate Authority provides a CA-verified certificate:
    - If your IT security team advises the use of an Intermediate CA, then use an **Intermediate Certificate File**. Otherwise, skip this file.
    - Click the Edit icon next to the target appliance, and Upload the **Certificate File** from the CA.
    - Upload the **Private Key File** that was generated as part of the CSR.
  3. To associate the CA verified certificate for use with Orchestrator, click **Add**.

## Orchestrator Reachability Tab

Administration > [General Settings > Setup] Orchestrator Reachability

You can specify how each appliance connects to Orchestrator by designating one of its interface Labels.

### Orchestrator Reachability

Appliances connect to Orchestrator differently based on the characteristics of the interface they're using to communicate (for example, via internal or external networks). Below, you can specify how each type of appliance interface (using its Label) should connect to the Orchestrator.

Default Orchestrator IP or Domain Name  
  Use Orchestrator Management IP

1 Rows, 1 Selected Search

Label	Orchestrator IP or Domain Name	Priority ▲	
MPLS	10.0.185.23	1	✕

## System Information

*Administration > [Software > Upgrade] System Information*

Manage system information with templates, with the exception of the following appliance-specific parameters:

- To edit **System Bandwidth**, use **Configuration > Shaper**.
- To change a **Deployment Mode**, use **Configuration > Deployment**.

Edit	Host Name	Property Key	Property Value
	Albuquerque	Retry count	4
	Albuquerque	Enable SaaS optimization	No
	Albuquerque	IP Id auto optimization	No
	Albuquerque	Bios Version	6.00
	Albuquerque	Bridge Loop Test	Yes
	Albuquerque	Site Name	
	Albuquerque	Interval	10
	Albuquerque	Serial Number	001BBC094101
	Albuquerque	Active S/W Release	0.0.0.0_71358
	Albuquerque	Configured Media Type	ram and disk
	Albuquerque	Hub Site	false
	Albuquerque	Automatically include local subnets	LAN=true, WAN=true
	Albuquerque	SSL optimization for non-IPSec tunnels	No
	Albuquerque	Use shared subnet information	Yes
	Albuquerque	Hold down count	1
	Albuquerque	System Bandwidth	200000
	Albuquerque	Always send pass-through traffic to original sender	No
	Albuquerque	Hardware Version	208001006000 Rev 67090
	Albuquerque	Connection Type	WEBSOCKET

When you click in the **Edit** column for a specific appliance, these two screens are available.

System Summary

The screenshot shows a window titled "System Information - Albuquerque" with a close button (X) in the top right corner. Below the title bar, there are two tabs: "System Summary" (which is selected and highlighted with a blue border) and "System Settings" (with a blue question mark icon). The main content area is divided into three sections: "General", "Configuration", and "Hardware".

General		Configuration	
Appliance Key	29.NE	System Bandwidth	200000 Kbps
Uptime	2d 3h 16m 36s	Mode	router
Active Release	0.0.0.0_71358		
Appliance ID	606465		
Discovery Method	PORTAL		
Connection Type	WEBSOCKET		

Hardware	
Appliance Model	EC-V 208001006000 Rev 67090
BIOS Version	6.00
Serial Number	00-1B-BC-09-41-01

At the bottom right of the window, there are two buttons: "Apply" and "Cancel".

## System Settings

## System Information - Osaka

X

System Summary

System Settings ?

## General

Model EC-V 208001006000 Rev 67090  
 Serial 001BBC0940F9  
 Site Name   
 Hub Site?  Not a hub  Hub  
 Contact Name   
 Contact Email   
 Location

## Optimization

IP Id auto optimization   
 TCP auto optimization   
 Flows and tunnel failure

## Subnet Sharing

Use shared subnet information   
 Automatically include local LAN subnets   
 Automatically include local WAN subnets   
 Metric for local subnets   
 Allow WAN to WAN routing

## Network Memory

Encrypt data on disk   
 Configured Media Type   
 Media Type

## Excess Flow Handling

Excess flow policy

## NextHop Health Check

Enable Health check   
 Retry count  (1..255)  
 Interval  (1..255) seconds  
 Hold down count  (1..255)

## Miscellaneous

SSL optimization for non-IPSec tunnels   
 Bridge Loop Test   
 Enable IGMP snooping   
 Auto Flow Re-Classify  (0..65535) seconds  
 Always send pass-through traffic to original sender   
 IPSec UDP Port   
 Enable default DNS lookup   
 Enable HTTP/HTTPS snooping   
 Quiescent tunnel keep alive time  (1..65535) seconds  
 UDP flow timeout  (1..65535) seconds  
 Non-accelerated TCP Flow Timeout  (1..65535) secs  
 Maximum TCP MSS  (500..9000) bytes  
 NAT-T keep alive time  (0..65535) seconds  
 Tunnel Alarm Aggregation Threshold  Tunnel Alarms  
Raise only 1 alarm above this threshold  
 Maintain end-to-end overlay mapping   
 IP Directed Broadcast

Apply Cancel

## System Information Property Keys

Property Key	Description
Active Release	Specifies the software the appliance is running.
Always send pass-through traffic to original sender	If the tunnel goes down when using WCCP and PBR, traffic that was intended for the tunnel is sent back the way it came.
Appliance ID	A unique identifier for the appliance.
Appliance Key	Orchestrator assigns and uses this key to identify the appliance.
Appliance Model	The specific EC, EC-V, NX, VX, or VRX model.
Auto Flow Reclassify	Specifies how often to do a policy lookup.
Automatically establish tunnels	Reduces configuration overhead by removing the need to manually create tunnels.
Automatically include local [LAN   WAN] subnets	Adds the local subnet(s) to the appliance subnet information. If the appliance is deployed in Bridge mode, then no interface is specified.
BIOS Version	The version of BIOS firmware that the appliance is using.
Bridge Loop Test	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.

Property Key	Description
Configured Media Type	Is either <b>ram and disk</b> (VX) or <b>ram only</b> (VRX). Can change for special circumstances, if recommended by Silver Peak.
Connection Type	The method that Orchestrator uses to communicate with the appliance. Options are WEBSOCKET, PORTAL, and HTTP.
Contact Email	Email address of whom to contact within your organization (optional)
Contact Name	Whom to contact within your organization (optional)
Discovery Method	Specifies how Orchestrator discovered the appliance: <ul style="list-style-type: none"> <li><b>PORTAL</b>            Orchestrator discovered the appliance through the portal account.</li> <li><b>MANUAL</b>            The appliance was added manually.</li> <li><b>APPLIANCE</b>        Orchestrator's IP address was added to the appliance. Portal was not involved.</li> </ul>
Enable default DNS lookup	Allows the appliance to snoop the DNS requests to map domains to IP addresses. This mapping can then be used in ACLs for traffic matching.
Enable Health check	Activates pinging of the next-hop router.
Enable HTTP/HTTPS Snooping	Enables a more granular application classification of HTTP/HTTPS traffic, by inspection of the HTTP/HTTPS header, Host. This is enabled by default.
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.
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.
Encrypt data on disk	Enables encryption of all the cached data on the disks. Disabling this option is not recommended.
Excess flow policy	Specifies what happens to flows when the appliance reaches its maximum capacity for optimizing flows. The default is to <b>bypass</b> flows. Or, you can choose to <b>drop</b> the packets.
Flows and tunnel failure	If there are parallel tunnels and one fails, then <b><i>Dynamic Path Control</i></b> determines where to send the flows. There are three options: <ul style="list-style-type: none"> <li>■ <b>fail-stick</b> - When the failed tunnel comes back up, the flows don't return to the original tunnel. They stay where they are.</li> <li>■ <b>fail-back</b> - When the failed tunnel comes back up, the flows return to the original tunnel.</li> <li>■ <b>disable</b> - When the original tunnel fails, the flows aren't routed to another tunnel.</li> </ul>
Hair-pin traffic	Redirects inbound LAN traffic back to the WAN

Property Key	Description
Hold down count	If the link has been declared down, this specifies how many successful ICMP echoes are required before declaring that the link to the next-hop router is up.
Hub Site	Specifies whether the appliance has been assigned the role, Hub, in Orchestrator. Options are <b>true</b> or <b>false</b> .
Interval	Specifies the number of seconds between each ICMP echo sent.
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).
IPSec UDP Port	Specifies the port that Orchestrator uses to build IPSec UDP tunnels. If the field is blank, Orchestrator uses the default.
Location	Appliance location, optionally specified during appliance setup.
Maximum TCP MSS	(Maximum Segment Size). The default value is 9000 bytes. This ensures that packets are not dropped for being too large. You can adjust the value (500 to 9000) to lower a packet's MSS.
Media Type	Displays the actual media being used.
Metric for local subnets	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.
Mode	Specifies the appliance's deployment mode: Server, Router, or Bridge.
Model	The specific EC, EC-V, NX, VX, or VRX model.
NAT-T keep alive time	If a device is behind a NAT, this specifies the rate at which to send keep alive packets between hosts, in order to keep the mappings in the NAT device intact.
Quiescent tunnel keep alive time	Specifies the rate at which to send keep alive packets after a tunnel has become idle (quiescent mode). The default is 60 seconds.
Region	A user-assigned name created for segmenting topologies and streamlining the number of tunnels created. When regions contain at least one hub, you can choose to connect Regions through hubs only.
Retry count	Specifies the number of ICMP echoes to send, without receiving a reply, before declaring that the link to the WAN next-hop router is down.
Serial Number	Serial number of the appliance
Site / Site Name	Orchestrator won't build tunnels between appliances with the same user-assigned site name.
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 Silver Peak GMS. This activity can apply to the entire distributed network of Silver Peak appliances, or just to a specified group of appliances.
System Bandwidth	The appliance's total outbound bandwidth, determined by appliance model or license.

Property Key	Description
<b>TCP auto optimization</b>	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).
<b>UDP flow timeout</b>	Specifies how long to keep the UDP session open after traffic stops flowing. The default is 120 seconds (2 minutes).
<b>Uptime</b>	The time elapsed since the appliance became operational and available.
<b>Use shared subnet information</b>	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.

# Software Versions

Administration > [Software > Upgrade] Software Versions

This report lists the **software versions** on each appliance.

Software Versions ×

Upgrade appliances software
Export
↻

Software Versions ?

3 Rows
Search

Appliance Name	Partition 1				Partition 2			
	Build Version	Build Date ▲	Active	Next Boot	Build Version	Build Date	Active	Next Boot
Tallinn	8.1.1.0_60681	2016-07-11 11:41:48	No	No	8.1.5.1_65516	2017-05-11 12:29:02	Yes	Yes
Iaine-vxa	8.1.5.1_65516	2017-05-11 12:29:02	No	No	8.1.5.4_67205	2017-09-28 14:28:54	Yes	Yes
Iaine-vxb	8.1.5.1_65516	2017-05-11 12:29:02	Yes	Yes	8.1.5.7_68124	2017-12-11 10:52:17	No	No

## Upgrading Appliance Software

*Administration > [Software > Upgrade] Upgrade Appliances*

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 **Upgrade Appliances** page to upload appliance software to 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.

**Upgrade Appliances**

**Select VXOA Image**

Name	Type	Version	Build Date	
image-8.1.4.2_63369...		8.1.4.2_63...	2017-01-11 15:...	X
image-7.3.3.0_57797...		7.3.3.0_57...	2015-12-09 21:...	X
image-7.3.1.0_56428...		7.3.1.0_56...	2015-09-10 10:...	X
image-7.2.1.0_55514...		7.2.1.0_55...	2015-05-27 14:...	X
image-6.2.5.0_52097...		6.2.5.0_52...	2014-07-22 17:...	X
image-6.2.5.0_51012...		6.2.5.0_51...	2014-06-01 12:...	X

**Target Appliances**

Appliance	Status	Progress
laine-vxb	[Slot0: 8.0.2.0_58453], [Slot1: 8.1.1.0_60681, Current, Next Boot]	
laine-vxa	[Slot0: 8.0.2.0_58453], [Slot1: 8.1.1.0_60681, Current, Next Boot]	
Tallinn	[Slot0: 8.1.1.0_60681, Current, Next Boot], [Slot1: 7.3.6.0_59199]	

Upload VXOA Image

image-8.1.4.2\_63369.img 0.2GB

**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.

The message indicates that this image just finished successfully uploading, as seen in the first line of the VXOA table.

- **Install and reboot** installs the image into the appliance's inactive partition and then reboots the appliance to begin using the new software.

- **Install and set next boot partition** installs the image into the appliance's inactive partition and then points to that partition for the next reboot.
- **Install only** downloads the image into the inactive partition.

## Appliance Configuration Backup

*Administration > [Software > Backup & Restore] Backup Now*

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) in the navigation tree, go to **Administration > Backup Now** and click **Backup**.

**Appliance Backup** ✕

Comment

Search

<i>Mgmt IP</i> ▲	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		



**Appliance Backup** ✕

Comment

Search

<i>Mgmt IP</i> ▲	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...

You cannot delete an appliance backup from Orchestrator.

## Viewing Configuration History

*Administration > [Software > Backup & Restore] Configuration History*

- You can view an appliance's current or previous configuration.
- You can compare any two appliance configuration files.

**Configuration History**

Select any two records to compare.

30 Rows, 1 Selected

Host Name	File Name	Backup Time	Software Versi...	File Conte...	Comment
laine-vxa (10.0.238.71)	initial	28-Dec-16 04:0...	8.1.1.0_60681	<a href="#">View</a>	
laine-vxb (10.0.238.69)	initial	28-Dec-16 04:0...	8.1.1.0_60681	<a href="#">View</a>	
laine-vxa (10.0.238.71)	initial	2-Nov-16 05:00:...	8.1.1.0_60681	<a href="#">View</a>	
laine-vxb (10.0.238.69)	initial	2-Nov-16 05:00:...	8.1.1.0_60681	<a href="#">View</a>	
laine-vxa (10.0.238.71)	initial	1-Nov-16 05:00:...	8.1.1.0_60681	<a href="#">View</a>	
laine-vxb (10.0.238.69)	initial	1-Nov-16 05:00:...	8.1.1.0_60681	<a href="#">View</a>	

Comparison Result :

laine-vxa (10.0.238.71) initial 2-Nov-16 05:00:05		laine-vxb (10.0.238.69) initial 2-Nov-16 05:00:05	
1	##	1	##
2	## Network interface MAC assignment	2	## Network interface MAC assignment
3	##	3	##
4	interface lan0 mac address 00:0C:29:19:53:BF	4	interface lan0 mac address 00:0C:29:ES:96:B4
5	interface mgmt0 mac address 00:0C:29:19:53:A1	5	interface mgmt0 mac address 00:0C:29:ES:96:96
6	interface mgmt1 mac address 00:0C:29:19:53:AB	6	interface mgmt1 mac address 00:0C:29:ES:96:A0
7	interface wan0 mac address 00:0C:29:19:53:B5	7	interface wan0 mac address 00:0C:29:ES:96:AA
8		8	
9	##	9	##
10	## Network interface configuration	10	## Network interface configuration
11	##	11	##
12	interface lan0 create	12	interface lan0 create
13	no interface lan0 dhcp	13	no interface lan0 dhcp
14	interface lan0 displ		
15	interface lan0 ip ac		
16	interface lan0 label		
17	interface lan0 mtu 1		
18	no interface lan0 sh		
19	interface lan0 speed		

**Backup file content**

```

##
## Network interface MAC assignment
##
interface lan0 mac address 00:0C:29:19:53:BF
interface mgmt0 mac address 00:0C:29:19:53:A1
interface mgmt1 mac address 00:0C:29:19:53:AB
interface wan0 mac address 00:0C:29:19:53:B5

##
## Network interface configuration
##
interface lan0 create
no interface lan0 dhcp
interface lan0 display
interface lan0 ip address 0.0.0.0 /0
interface lan0 label ""
interface lan0 mtu 1500
no interface lan0 shutdown
interface lan0 speed-duplex auto/auto
interface wan0 create
no interface wan0 dhcp
interface wan0 display
interface wan0 label ""
interface wan0 mtu 1500
no interface wan0 shutdown
interface wan0 speed-duplex auto/auto

##
## Other IP configuration
##
hostname laine-vxa

##
## Local user account configuration
##
username myself capability admin
no username myself disable
username myself password 7 $1$BXJ0P8Im$FXi1VUu9EMMAega5AYdX1.

##
## System Network Config
    
```

## Restoring a Backup to an Appliance

*Administration > [Software > Backup & Restore] Restore*

You can restore an appliance configuration backup from Orchestrator to any other Silver Peak appliance(s) in your network.

However, be careful to consider any potential conflicts when the backup specifies a static **mgmt0** IP address, as opposed to specifying DHCP.

**Appliance Restore and Reboot** ✕

Source Appliance  -----> Target Appliance    
select an appliance  
laine-vxa  
laine-vxb  
Tallinn

Select the configuration to restore from the table below

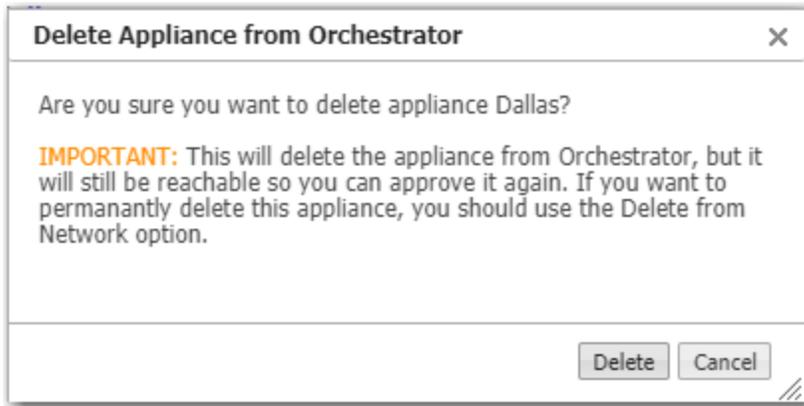
File Name	Backup Time	Software Vers	Tallinn	Comment
initial	28-Dec-16 04:00:05	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	2-Nov-16 05:00:05	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	1-Nov-16 05:00:05	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	15-Sep-16 05:00:05	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	14-Sep-16 05:00:05	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	26-Aug-16 05:00:06	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	25-Aug-16 05:00:05	8.1.1.0_60681	<a href="#">View ↗</a>	
initial	12-Jul-16 05:00:06	8.1.1.0_60681	<a href="#">View ↗</a>	
backup.1432626300007.10.NE	26-May-15 00:45:02	6.2.7.0_53789		Weekly Appliance Backup
backup.1432021500007.10.NE	19-May-15 00:45:02	6.2.7.0_53789		Weekly Appliance Backup

Status Log

## Remove Appliance from Orchestrator

*Administration > [Software > Remove Appliances] Remove from Orchestrator*

Removing an appliance with this action returns the appliance to the **Discovered Appliances** list.



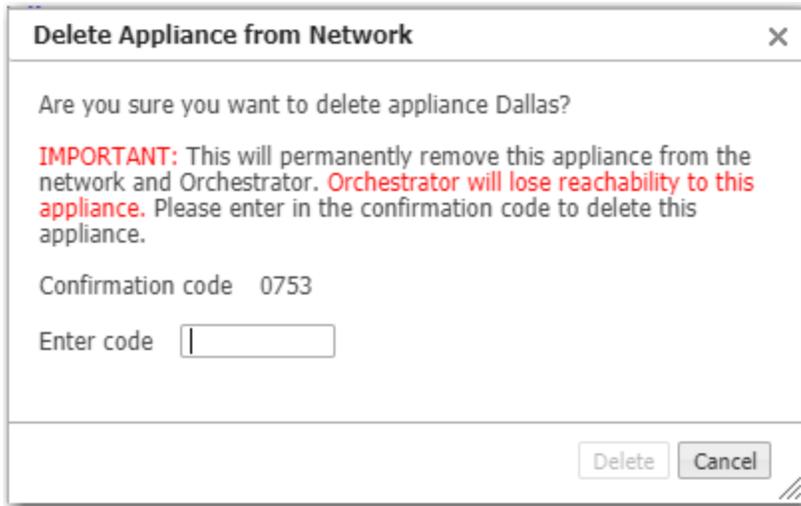
Additionally,

- It deletes the appliance from the navigation tree.
- Orchestrator will break all tunnels/overlays/etc. to this device.

## Remove Appliance from Orchestrator and Account

*Administration > [Software > Remove Appliances] Remove from Orchestrator and Account*

Removing an appliance with this action places the appliance in the **Denied Devices** list, which is located as a link in the **Configuration - Discovered Appliances** menu.



Additionally,

- It deletes the appliance from the navigation tree.
- Orchestrator will break all tunnels/overlays/etc. to this device.
- It tells the Portal to "unlicense" the appliance.

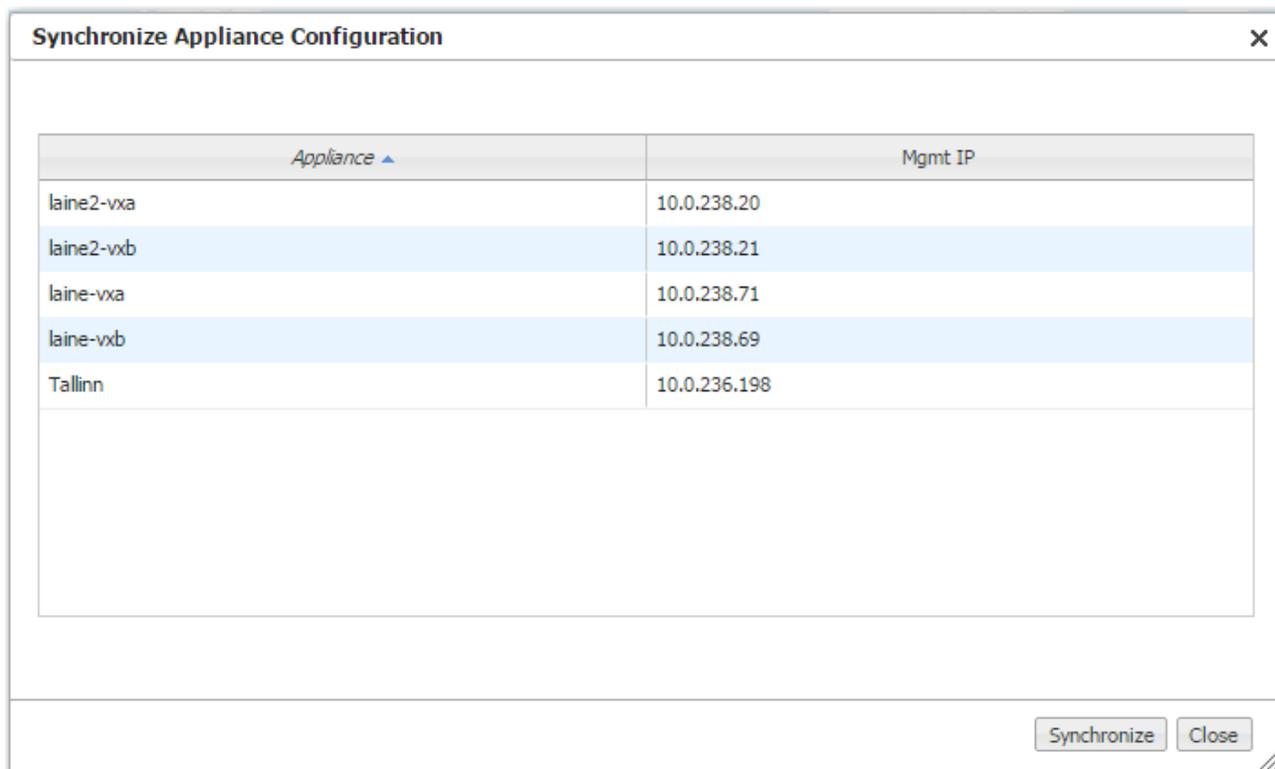
## Synchronizing Appliance Configuration

Administration > [Tools] Synchronize

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 Orchestrator and the appliance in sync.
- Whenever an appliance starts or reboots, Orchestrator automatically inventories the appliances to resync.
- Whenever 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 use this page to manually resync to ensure that Orchestrator has an appliance's current running configuration.



Appliance ▲	Mgmt IP
laine2-vxa	10.0.238.20
laine2-vxb	10.0.238.21
laine-vxa	10.0.238.71
laine-vxb	10.0.238.69
Tallinn	10.0.236.198

## Putting the Appliance in System Bypass Mode

Administration > [Tools] Bypass

System Bypass mode is only available for certain models of Silver Peak physical appliances. Virtual appliances don't support bypass mode.

In **system bypass mode**, the fail-to-wire (or fail-to-glass) card **DOES NOT** receive or process packets.

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 uptime.

- In an in-line deployment (Bridge mode), the LAN interface is physically connected to the WAN interface.
- In Server mode and any Router mode, the appliance is in an open-port state.

**Bypass** ×

*Finished Bypass Enable: operation completed*

Enable  Disable

3 Rows, 1 Selected Search

Appliance ▲	Bypass	Details
laine-vxb	Not Supported	Bypass enable failed: System does not support bypass mode.
laine-vxa	Not Supported	Bypass enable failed: System does not support bypass mode.
<b>Tallinn</b>	<b>ON</b>	Bypass enable successful

When the appliance is in Bypass mode, a message displays in red text at the upper right corner of the user interface.

Name	<b>Tallinn (BYPASS)</b>	Model	<b>NX-3700</b>	
Up Time	<b>13d 13m 35s</b>	VXOA		
Time	<b>2016/03/03 22:04:57 UTC</b>	User	<b>admin</b> <a href="#">[logout]</a>	
Alarms	<b>1 Critical</b>	0 Major	0 Minor	0 Warn

## Broadcasting CLI Commands

*Administration > [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\*](#).

---

## Link Integrity Test

Administration > [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 **tcppperf** (Version 1.4.8).

The screenshot shows the 'Link Integrity Test' configuration window. On the left, the 'laine-vxb' appliance is shown as a server listening on UDP port 5555. On the right, the 'laine-vxa' appliance is shown as a client connecting to 10.1.154.20 on UDP port 5555. The central control panel has the following settings: Bandwidth (laine-vxb to laine-vxa) set to 1000 Kbps; Bandwidth (laine-vxa to laine-vxb) set to 1000 Kbps; Duration (seconds) set to 10; DSCP set to 'any'; Mode set to 'pass-through'; and Test Program set to 'iperf'. Below these settings are 'Start' and 'Clear' buttons. The interface also displays a table of test results for both directions of traffic, showing time intervals, data transferred, throughput, and loss percentage.

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).
- 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 **tcppperf** and should be used cautiously, by advanced users.

## TCPPERF Version 1.4.8

### Basic Mode

Option	Description
-h	<i>help</i>
-s	<i>server</i> : Run tcppperf in server mode (not applicable for file generation). Listens on TCP port 2153 by default. [server_port [server_port [server_port].]]
-sr	<i>server range</i> : <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	Description
-6	<i>ip6</i> . Forces tcppperf to use IPv6 addresses only. Default is IPv4 addresses.
-I	<i>interface IP</i> : Specify source interface IP address. Default is <b>any</b> .
-o	<i>outname</i> : Output filename. Default is <b>stdout</b> .
-u	<i>update &lt;secs&gt;</i> : Frequency of printed updates in seconds. Default is <b>1</b> .
-d	<i>duration &lt;secs&gt;</i> : Set maximum test duration in seconds. Default is <b>infinite</b> .
-w	<i>wait &lt;secs&gt;</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.

Option	Description
-q	<i>quiet &lt;level&gt;</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	Description
-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 &lt;cp&gt;</i> : Sets IP DSCP to <cp> (decimal). Default is 0.
-tr	<i>tcpretries &lt;n&gt;</i> : Sets number of times to retry TCP connections.
-tp	<b>tcppace &lt;n&gt; [mode]</b> : Pace TCP connection setup rate. Limits number of half-open connections to <n>. Valid <mode> types are: <b>preestablish</b> . All connections are established prior to data transmission. Default. <b>simultaneous</b> . Begin data transmission as soon as connection made
-ta	<i>tcpabort</i> : Sends RSTs instead of FINs on close.
-tf	<i>tcpfindelay &lt;secs&gt;</i> : Time to wait after all data sent before sending FIN/RST

### Traffic Generation Parameters

Option	Description
-f	<i>file</i> . Source filename to load. Default is 10MB of random data.
-i	<i>test id &lt;i&gt;</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 &lt;n&gt;</i> : Generate <n> flows. Default is one.
-b	<i>begin &lt;byte&gt;</i> : First byte in transmission. Default is zero.
-e	<i>end &lt;byte&gt;</i> : 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.

Option	Description
<b>-a</b>	<p><i>antipat &lt;mode&gt;</i>: Antipattern mode: default is mutate:</p> <p><b>none</b> Repeats same content verbatim on all flows. Repeats content if end byte exceeds content size.</p> <p><b>mutate</b> Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Destroys cross flow similarity. Destroys original byte code distribution.</p> <p><b>shuffle</b> Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Preserves cross flow similarity. Preserves original byte code distribution.</p> <p><b>fast</b> 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.</p>
<b>-l</b>	<p><i>loopback [mode]</i>: Loopback. Default is unidirectional.</p> <p><b>uni</b> Unidirectional client to server.</p> <p><b>rev</b> Unidirectional server to client.</p> <p><b>bidir</b> Bidirectional, client and server independently send data on the same TCP connection.</p> <p><b>bidir2</b> Bidirectional, client and server independently send data on secondary TCP connections.</p> <p><b>loop</b> Bidirectional, server loops data back to client on the same TCP connection.</p> <p><b>loop2</b> Bidirectional, server loops data back to client on a secondary TCP connection.</p> <p><b>bidir2</b> Bidirectional, transmits one transaction at a time. Client waits for previous transaction to be echoed. Emulates transactional data.</p> <p>NOTES:</p> <ol style="list-style-type: none"> <li>Content source for traffic originating at the server is determined by the server (not client) command line.</li> <li><b>loop2</b> and <b>bidir2</b> modes 2 x &lt;n&gt; TCP connections and requires that the server has even-numbered ports available.</li> </ol>
<b>-r</b>	<i>rate &lt;bps&gt;</i> : Limits aggregate transmission rate to <bps>. Default is no rate limit.
<b>-t</b>	<p><i>trans &lt;min&gt; [max]</i>: Sets size of each socket transaction. Default is 64000.</p> <p>If &lt;min&gt; and &lt;max&gt; are specified, client generates transactions with random sizes between &lt;min&gt; and &lt;max&gt;. This feature is often used with <b>-l</b> and <b>-r</b>. Set the minimum transaction size to 100000 to improve single-flow performance.</p>

Option	Description
-v	<p><i>verify &lt;mode&gt;</i>: Verify integrity of received data. Default is <b>global</b>.</p> <p><b>none</b>            No verification. Fastest/least CPU load.</p> <p><b>global</b>            Single global hash per flow. Fast, but cannot isolate an errored block.</p> <p><b>literal</b>            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 -f file at server.</p> <p><b>embedded</b>        Embedded hashes every 4096 bytes. Slower, can isolate errors to 4096 byte block.</p>
-p	<p><i>repeat &lt;n&gt;</i>: Repeat each content byte n times. Default is <b>1</b> (no repeats). Works for both random data and file content.</p>
-k	<p><i>corrupt &lt;n&gt; &lt;m&gt; &lt;s&gt; [&lt;%change&gt;[&lt;%insert&gt;[&lt;%delete&gt;]]]</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 &lt;b&gt; &lt;e&gt; &lt;l&gt; [s]</i>: Send random excerpts of average &lt;l&gt; length bytes from content between &lt;b&gt;egin and &lt;e&gt;nd bytes. The -b and -e options still specify total bytes to send. Uses random seed s.</p>
-y	<p><i>defred &lt;s% &gt; &lt;m%&gt; &lt;l%&gt; &lt;sb&gt; &lt;smin&gt; &lt;smax&gt; &lt;mb&gt; &lt;mmin&gt; &lt;mmax&gt; &lt;lb&gt; &lt;lmin&gt; &lt;lmax&gt;</i> :</p> <p>Generate content based on defined reduction model. Content is drawn from three data sets: <b>s</b>, <b>m</b>, and <b>l</b>:</p> <p><b>s%</b>            Specifies fraction [50%] of s-type content (short term reducible).</p> <p><b>m%</b>            Specifies fraction [30%] of m-type content (medium term reducible).</p> <p><b>l%</b>            Specifies fraction [20%] of l-type content (long term reducible).</p> <p>Short term content comes from data set of sb Mbytes [100MB] with excerpts uniformly distributed between smin and smax bytes [10K-1M]. Medium term content comes from data set of mb Mbytes [100GB] with excerpts uniformly distributed between lmin and lmax bytes [10K-1M]. Long term content comes from data set of lb Mbytes [100TB] with excerpts uniformly distributed between smin and smax bytes [10K-1M]. The <b>-b</b> and <b>-e</b> options still specify total bytes to send. Performance is best if <b>-b</b> is <b>0</b>. Uses random seed <b>s</b>.</p>
-ssl [param=value ...]	<p>Enable SSL on connection with optional parameters.</p> <p><b>version=2 3 t10 t11 t12</b>. Set the protocol version.</p> <p><b>cipher=OPENSSL-CIPHER-DESC</b>. Set the choice of ciphers.</p> <p><b>ticket=yes no</b>. Enable/disable session ticket extension.</p> <p><b>cert=FILENAME</b>. Use this certificate file.</p> <p><b>key=FILENAME</b>. Use this private keyfile.</p> <p><b>compression=none any deflate zlib rlc</b>. Set the compression method.</p> <p><b>sslcert</b>. Print the SSL certificate in PEM format.</p> <p><b>sslkey</b>. Print the SSL key in PEM format.</p>

# Disk Management

Administration > [Tools] Disk Management

The **Disk Management** tab 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' interface. At the top, there are tabs for 'Topology' and 'Disk Management'. Below the tabs are 'Export' and 'Refresh' buttons. The main area is titled 'Disk Management' and contains a table with 10 rows. The first two rows are for 'Tallinn' appliances. An orange circle highlights the 'Edit' icon in the first row of the 'Tallinn' entries. A callout box points to this icon with the text: 'For example, to access Tallinn's own'. Below the table, a 'Smart data for Tallinn ID 1' window is open, displaying SMART data for the selected disk.

Edit	Appliance Name	Appliance Model	Slot ID	Pairing Slot ID	Status	Size(GB)	Serial Number	Removable	SMART Data
	Tallinn	NX-3700	0	1	OK	465	9XF2JF81	Yes	
	Tallinn	NX-3700	1	0	OK	465	9SP134KN	Yes	
	laine-vxa	VX-1000	0		OK	30		No	N/A
	laine-vxa	VX-1000	1		OK	70		No	N/A
	laine-vxb	VX-1000	0		OK	30		No	N/A
	laine-vxb	VX-1000	1		OK	70		No	N/A
	laine2-vxa	VX-1000	0		OK	30		No	N/A
	laine2-vxa	VX-1000	1		OK	70		No	N/A
	laine2-vxb	VX-1000	0		OK	30		No	N/A
	laine2-vxb	VX-1000	1		OK	70		No	N/A

Attribute	Normalized Value	Worst Value	Raw Value
Read Error Rate	81	61	219,710,224
Spin up time	100	100	0
Start/stop count	100	100	100
Reallocated sector count	100	100	11
Seek error rate	80	60	299,502,215
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	1,078,613
High Fly Writes	1	1	186
Temperature Difference	76	26	4,100,256

### To replace a failed disk:

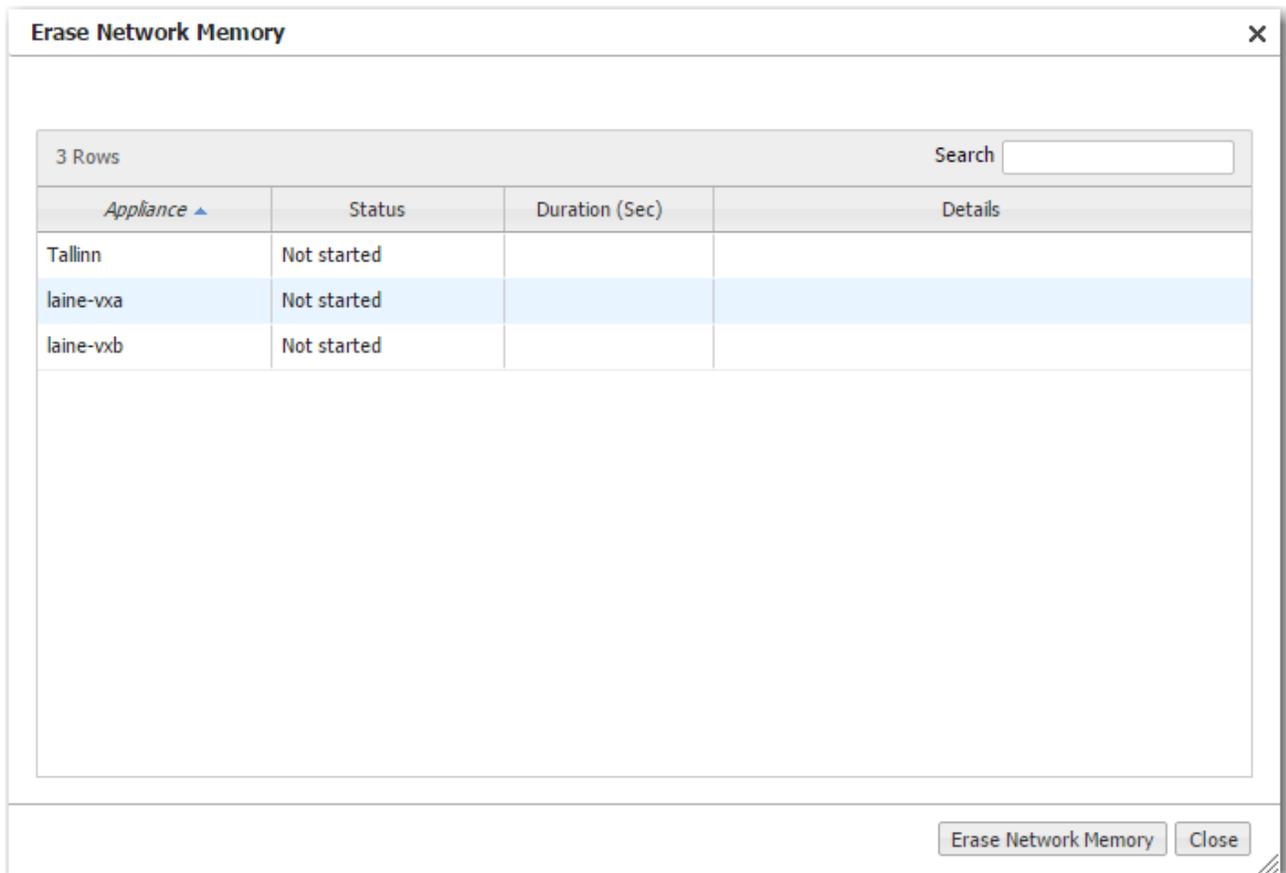
1. Log into your Support portal account, and click **Open a Self Service RMA** for disk replacement.
2. Complete the wizard, using the serial number of the appliance (not the disk).
3. After you receive the new disk, access Appliance Manager by clicking any **Edit** icon that belongs to the appliance in question.
4. Follow the instructions in that page's on-line help.

## Erasing Network Memory

*Administration > [Tools] Erase Network Memory*

Erasing Network Memory removes all stored local instances of data.

No reboot required.



**Erase Network Memory** ×

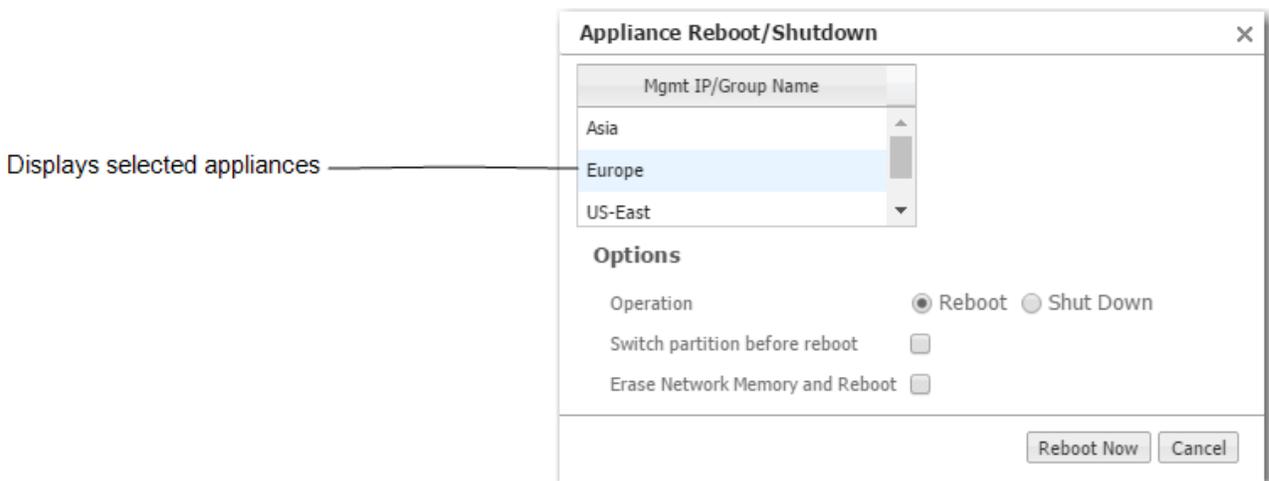
3 Rows Search

<i>Appliance</i> ▲	Status	Duration (Sec)	Details
Tallinn	Not started		
laine-vxa	Not started		
laine-vxb	Not started		

## Rebooting or Shutting Down an Appliance

*Administration > [Tools > Reboot] 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:

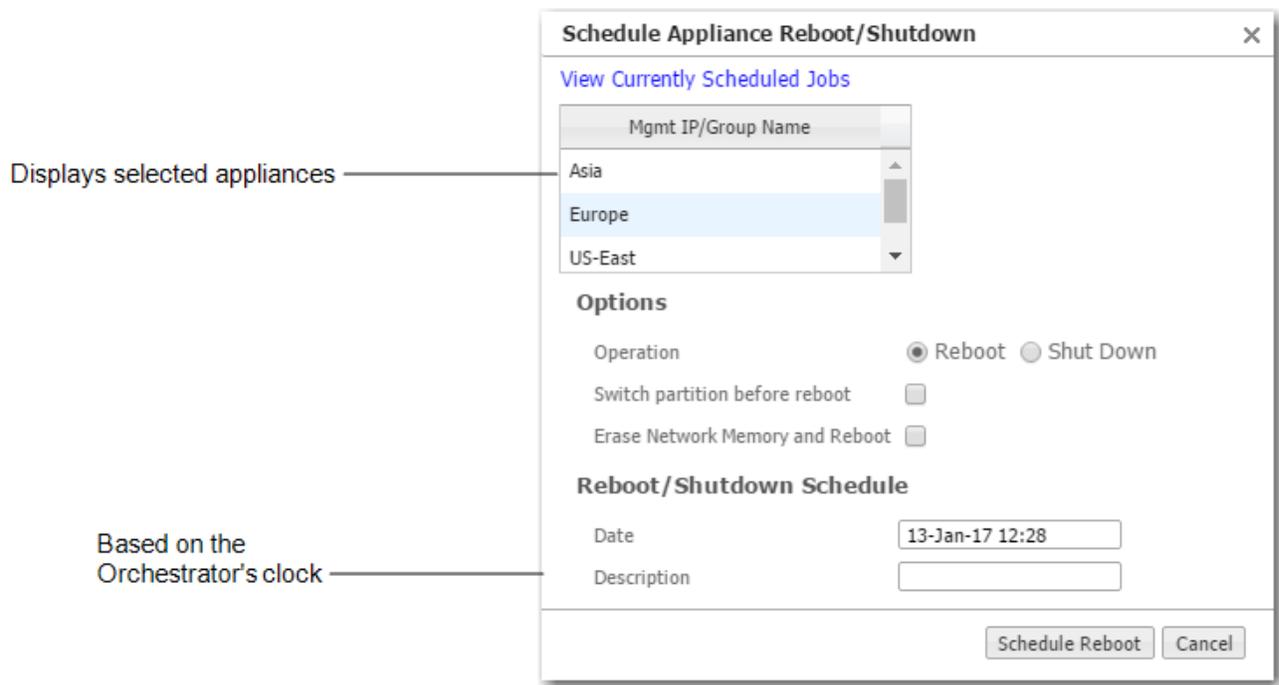
- *hardware bypass*, if deployed in-line (Bridge mode), or
- *an open-port state*, if deployed out-of-path (Router mode or Server mode).

Unless a ***virtual appliance*** is configured for a high availability deployment, all flows are discontinued during reboot.

## Scheduling an Appliance Reboot

*Administration > [Tools > Reboot] Schedule Appliance Reboot*

You can schedule an appliance for any of 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.



To specify the timezone for scheduled jobs and reports, go to **Orchestrator > [Software & Setup > Setup] Timezone for Scheduled Jobs**.

---

# Reachability Status Tab

Administration > [Tools > Monitoring] Reachability Status

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

The top screenshot shows the 'Appliances to Orchestrator' view. The table below represents the data shown in this view:

Appliance Name	Admin Username	Protocol	State
Albuquerque	admin	BOTH	Normal
Boston	admin	BOTH	Normal
Chicago	admin	BOTH	Normal
Dallas	admin	BOTH	Normal
Denver	admin	BOTH	Normal
Los Angeles	admin	BOTH	Normal
Mexico City			
Miami			
Minneapolis			
New Orleans			

The bottom screenshot shows the 'Orchestrator to Appliances' view. The table below represents the data shown in this view:

Appliance Name	Orchestrator IP	Web Socket
Albuquerque	10.0.185.23	Reachable
Boston	10.0.185.23	Reachable
Chicago	10.0.185.23	Reachable
Dallas	10.0.185.23	Reachable
Denver	10.0.185.23	Reachable
Los Angeles	10.0.185.23	Reachable
Mexico City	10.0.185.23	Reachable
Miami	10.0.185.23	Reachable
Minneapolis	10.0.185.23	Reachable
New Orleans	10.0.185.23	Reachable
New York	10.0.185.23	Reachable
Pittsburgh	10.0.185.23	Reachable
Portland	10.0.185.23	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.

## Active Sessions Tab

Administration > [Tools > Monitoring] Active Sessions

These tables display which users are logged in to Orchestrator and which appliances Orchestrator is managing.

Active Sessions ×

Orchestrator Appliance

### Active Sessions

5 Rows Search

User Name	Type	From	Login Time	Idle Time
admin	web	172.20.30.12	17-Jan-17 10:16	6m 9s
admin	web	172.23.43.13	17-Jan-17 10:16	38m 17s
admin	web	172.23.41.56	17-Jan-17 10:44	34m 21s
admin	web	172.23.48.156	17-Jan-17 10:40	4m 2s
admin	web	172.23.41.68	17-Jan-17 11:14	0s

Active Sessions ×

Orchestrator Appliance

### Active Sessions

22 Rows Search

Appliance Name ▲	User Name	Type	From	Login Time	Idle Time
Chennai	admin	web	10.0.239.69	17-Jan-17 10:03	0s
Chicago	Orchestrator	web		17-Jan-17 10:04	0s
Chicago	admin	web	orch-172.23.41.68	17-Jan-17 11:38	26s
London	admin	web	10.0.239.69	17-Jan-17 10:03	0s
Los-Angeles	admin	web	orch-172.23.41.68	17-Jan-17 11:38	26s
Los-Angeles	admin	web	172.23.43.13	17-Jan-17 09:57	7m 0s
Los-Angeles	admin	web	orch-172.23.41.56	17-Jan-17 10:44	55m 38s
Los-Angeles	admin	web	orch-172.23.48.156	17-Jan-17 10:43	56m 7s

# Orchestrator Administration

This section describes items related to managing Orchestrator itself. These activities do not relate to managing appliances.

## Role Based Access Control

*Orchestrator > Orchestrator Server > Users & Authentication > Role Based Access Control (RBAC)*

The Role Based Access Control allows for a more specified experience of your Orchestrator UI. You can assign roles and customize appliance access to a user, as well as specify the menu per user in the Orchestrator UI tree.

### Assign Roles & Appliance Access

Complete the following steps to assign roles and appliance access.

1. In the **Role Based Access Control** tab, select **Assign Roles & Appliance Access**.
2. Select the **User** field and enter the name of the user.
3. Select the **Appliance** field and select the name of the reference to group's appliance that you created in the **Appliance Access** tab.
4. Check the roles you want to apply to your user.
5. Select **Save**.

The following table defines each default role you can select in step 4.

Field	Definition
<b>ConfigAdmin</b>	You can backup and restore appliance configuration and view the configuration history.
<b>OrchestratorAdmin</b>	Allows you to <b>only</b> perform Orchestrator operations, such as settings, tools, user management, and Orchestrator upgrades. Appliance operations are <b>not</b> allowed.
<b>OverlayAdmin</b>	A global role for managing SD-WAN overlays. <b>NOTE</b> Overlay management cannot be specific to a site or region.
<b>SiteMonitor</b>	Read-only permissions equivalent to SiteAdmin.
<b>SiteOperator</b>	Allows for appliance or site specific operations, such as configure appliance specific policies, ACLs, TCAs, SSL certificate; however, this role does <b>not</b> allow you to upgrade or remove an appliance from the network. You also <b>cannot</b> perform global, SD-WAN functions such as overlay management or Zscaler orchestration.
<b>SiteUpgradeAdmin</b>	Allows you to upgrade appliance(s) and remove them from the network.
<b>SuperAdmin</b>	Allows for Read-Write level access to all the menus.

Field	Definition
SiteAdmin	Allows for appliance or site-specific operations, such as configure appliance specific policies, ACLs, TCAs, SSL certificates, and upgrade. You cannot perform global SD-WAN functions like overlay management, Zscaler orchestration, or remove the appliance from the network.
Support	Allows for all support operations.
Monitor	Provides Read-Only level access to all the menus.
SuperAdmin	Provides Read-Write level access to all the menus.

## Roles

There is a set of default roles you can use. You can also create your own role or modify an existing one.

Field	Definition
Role	The name of the default role or the role you created.
Permission	The permission you selected for a given user. <b>Read-Write</b> or <b>Read-Only</b> .
Features	The accessible features for a given user.

To add a role:

1. Select **Create Roles**.
2. Select **Add** or the **Edit** icon from the Roles window.
3. Enter the name for your role.
4. Select a category you want to assign to your user from the following tabs: **Monitoring**, **Configuration**, **Administration**, **Orchestrator**, **Support**, or **Miscellaneous**.
5. Assign an access level to any of the categories: Select **Read Only** or **Read & Write**.
6. Check any of the boxes you want to apply to your role within the designated categories.

**NOTE** You can **Select All** or **Unselect All**.

7. Select **Save**.

## Appliance Access

You can also add specific versions of appliances to a user. Complete the following steps to customize appliance access.

1. Select **Create Appliance Access Groups** in the **Role Based Access Control** tab. The **Appliance Access Group** window opens.
2. Select **Add** or the **Edit** icon to modify or create an existing appliance rule.
3. Select the name field and enter the name of the appliance.
4. Select whether you want to **Select By Groups** or **Select By Region**. You can add all groups or regions or just select a few.
5. Select **Save**.

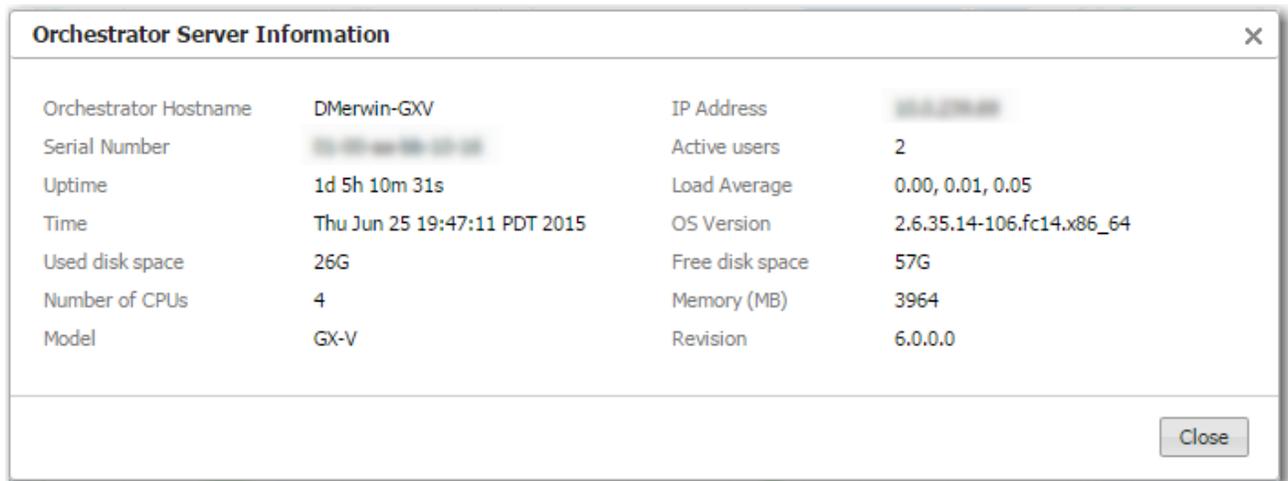
**WARNING** If you are an RBAC user with appliance access only (i.e. without any assigned roles, you will have access to the Appliance Manager, CLI Session, and Broadcast CLI. If you are an RBAC user with any role assigned, access to the Appliance Manager, CLI Session, and Broadcast CLI will be denied.

User	Appliance Access	Roles?	Menu Options
RBAC User	Yes	None assigned	Appliance Manager, CLI Session, Broadcast CLI
RBAC User (non-RBAC User)	No	None assigned	Appliance Manager, CLI Session, Broadcast CLI
RBAC User	No	Any	Appliance Manager, CLI Session, and Broadcast will be denied.

## Viewing Orchestrator Server Information

*Orchestrator [Orchestrator Server > Server Management] Server Information*

This page provides data specific to this Orchestrator server.



The screenshot shows a dialog box titled "Orchestrator Server Information" with a close button (X) in the top right corner. The dialog contains a table of server details:

Orchestrator Hostname	DMerwin-GXV	IP Address	10.0.2.15
Serial Number	00-00-00-00-00-00	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

A "Close" button is located in the bottom right corner of the dialog box.

## Restart, Reboot, or Shutdown

*Orchestrator [Orchestrator Server > Server Management] Restart Orchestrator*

*Orchestrator [Orchestrator Server > Server Management] Reboot Server*

*Orchestrator [Orchestrator Server > Server Management] Shutdown Server*

Orchestrator provides these three convenient actions in the **Orchestrator** menu:

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

## Managing Orchestrator Users

*Orchestrator > [Orchestrator Server > Users & Authentication] User Management*

The **User Management** page allows you to manage who has Read-Write or Read-Only access to Orchestrator.

The screenshot shows the 'User Management' window with the following details:

- Auto Logout:** 30 (1-10080 minutes)
- Max Sessions:** 20 (5-10000)
- Active Sessions:** Add
- Table:** 10 Rows, Search bar, and a table with columns: Edit, User Name, First Name, Last Name, Phone, Email, Two Factor ..., Two Factor ..., Create time, Status, Role, and a delete icon (X).
- Table Data:**

Edit	User Name	First Name	Last Name	Phone	Email	Two Factor ...	Two Factor ...	Create time	Status	Role	
	admin	Admin				No	No	23-Aug-17 1...	Active	Read-Write	
	test				spondugula...	No	No	09-Feb-18 1...	Active	Read-Only	X
	sbheemaraj...					No	No	18-Apr-18 0...	Active	Read-Write	X
	srinivas					No	No	30-Apr-18 1...	Active	Read-Write	X
	syen	Shyh-Pei	Yen		syen@silver-...	No	No	07-May-18 0...	Active	Read-Write	X
	anusha-ro	anusha-ro	read-only			No	No	21-May-18 1...	Active	Read-Only	X
- Buttons:** Apply, Cancel, Close

### Adding a User

- Users can have either **Read-Write** or **Read-Only** privileges. These provide prescribed access to Orchestrator menus.

To further limit the what users can see, you can assign them to customized menu groups in **Orchestrator > User Menu Access**.

- Multi-Factor Authentication (MFA) is a recommended option for each Orchestrator user.
- You cannot modify a Username. You must delete it and create a new user.

## Multi-Factor Authentication

Silver Peak Orchestrators support Multi-Factor Authentication (MFA). This is available on all platforms of the Silver Peak Orchestrator, including on-premise and cloud versions.

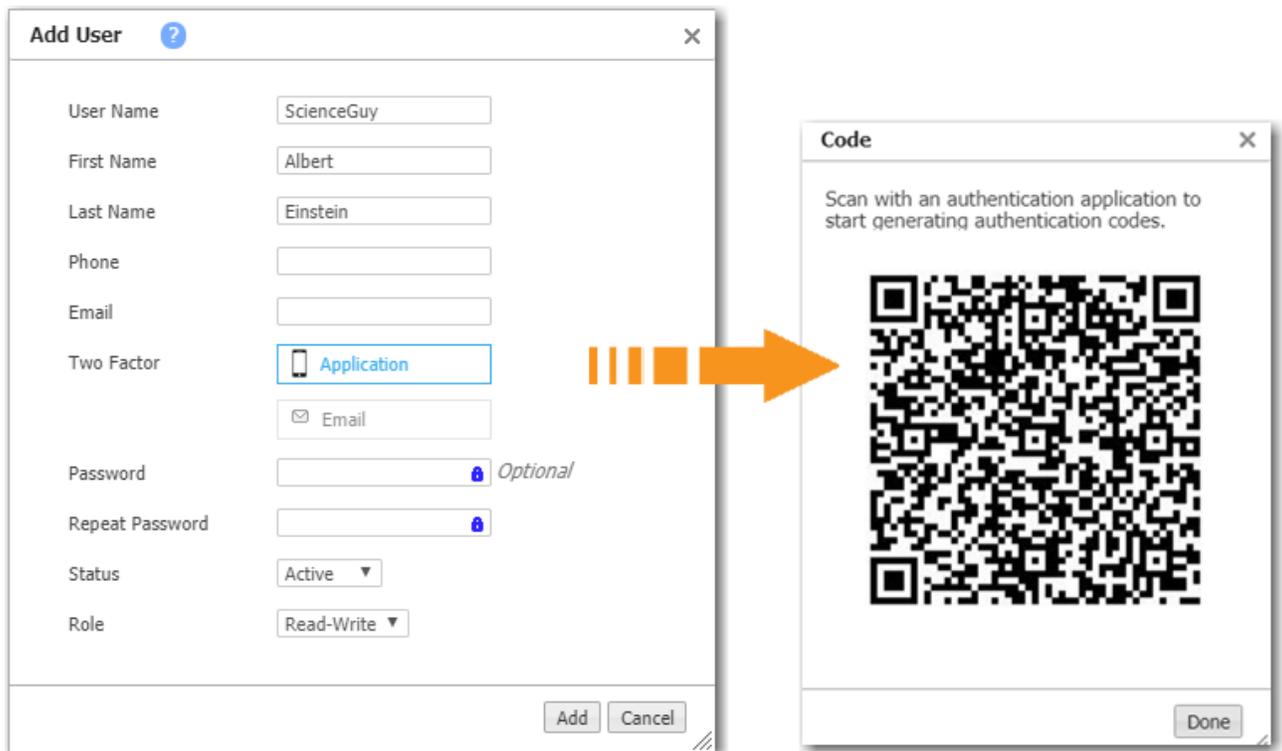
The first step in authentication is always username/password. For added security, users can choose between **Application** or **Email** based authentication, as described below.

**NOTE** Currently, only **admin** users can only configure Multi-Factor Authentication, and only for themselves.

### Configuring Multi-Factor Authentication through an Application

Orchestrator supports applications that provide time-based keys for two-factor authentication and are compliant with RFC 4226 / RFC 6238. Google Authenticator is one such app. The example below uses Google Authenticator on a mobile phone. You can also use a desktop version.

1. To enable Multi-Factor Authentication, go to **Orchestrator > User Management** and click on your username.
2. For **Two Factor**, click **Application**. Orchestrator generates a time-limited QR code.



The image shows two overlapping dialog boxes. The left dialog box is titled "Add User" and contains the following fields:

- User Name: ScienceGuy
- First Name: Albert
- Last Name: Einstein
- Phone: (empty)
- Email: (empty)
- Two Factor: Application (selected), Email
- Password: (empty) Optional
- Repeat Password: (empty)
- Status: Active
- Role: Read-Write

At the bottom of the "Add User" dialog are "Add" and "Cancel" buttons. An orange arrow points from the "Application" option in the "Two Factor" section to the right dialog box. The right dialog box is titled "Code" and contains the text "Scan with an authentication application to start generating authentication codes." Below the text is a large QR code. At the bottom of the "Code" dialog is a "Done" button.

3. With the Google Authenticator app, use the **Scan barcode** function to read the QR code. You

will also be prompted to enter your Orchestrator username and password.

Here you can see Google Authenticator with the new **admin** account added for the Orchestrator, **silverpeak-gxv**.



### Configuring Multi-Factor Authentication through Email

1. To enable Multi-Factor Authentication, go to **Orchestrator > User Management** and click on your username.
2. For **Two Factor**, click **Email** and enter your email address.

If an invalid email address is entered, the account could be locked out and would require password reset procedures.

3. After you click **Add** at the bottom of the dialog, Orchestrator sends you a time-limited authentication code via email. To verify your email address, click that link.

Orchestrator then opens a browser window telling you that your email address has been verified.

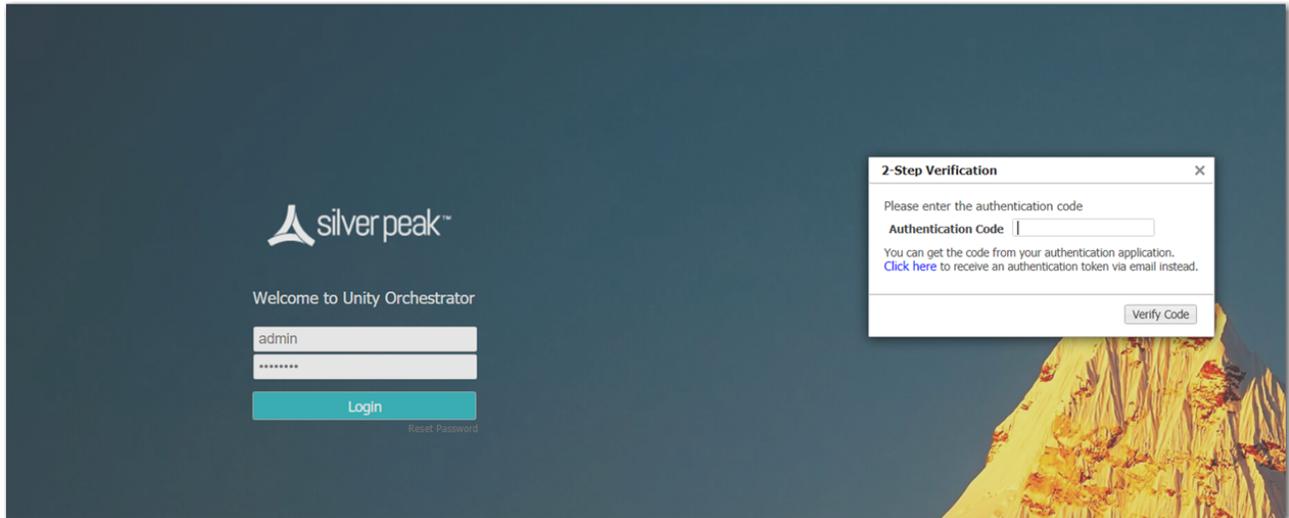
### Using Multi-Factor Authentication

After Multi-Factor Authentication is configured, every login requires two steps—entering the username/password and entering the current token.

Based on which authentication method you chose, do one of the following:

- Use the current token from the Google Authenticator (or other) app.
- Use the code you receive in email.

In both cases, the codes have a specific expiry time.



## Modify User

*Orchestrator > [Users & Authentication] User Management > Edit > Modify User*

The screenshot shows a 'Modify User' dialog box with the following fields and values:

- User Name: test
- First Name: (empty)
- Last Name: (empty)
- Phone: (empty)
- Email: spondugula@silver-peak.
- Two Factor: Application
- Password: \*\*\*\*
- Repeat Password: \*\*\*\*
- Status: Active
- Role: Read-Only

Buttons: Apply, Cancel

- **User Name** is the identifier the user uses to log in.
- **First Name**, **Last Name**, and **Phone Number** are optional information.
- **Email** is required if two factor authentication is enabled.
- **Two Factor Authentication**

This is a second step in the login process, where an authentication code is required.

The code can be obtained in two ways:

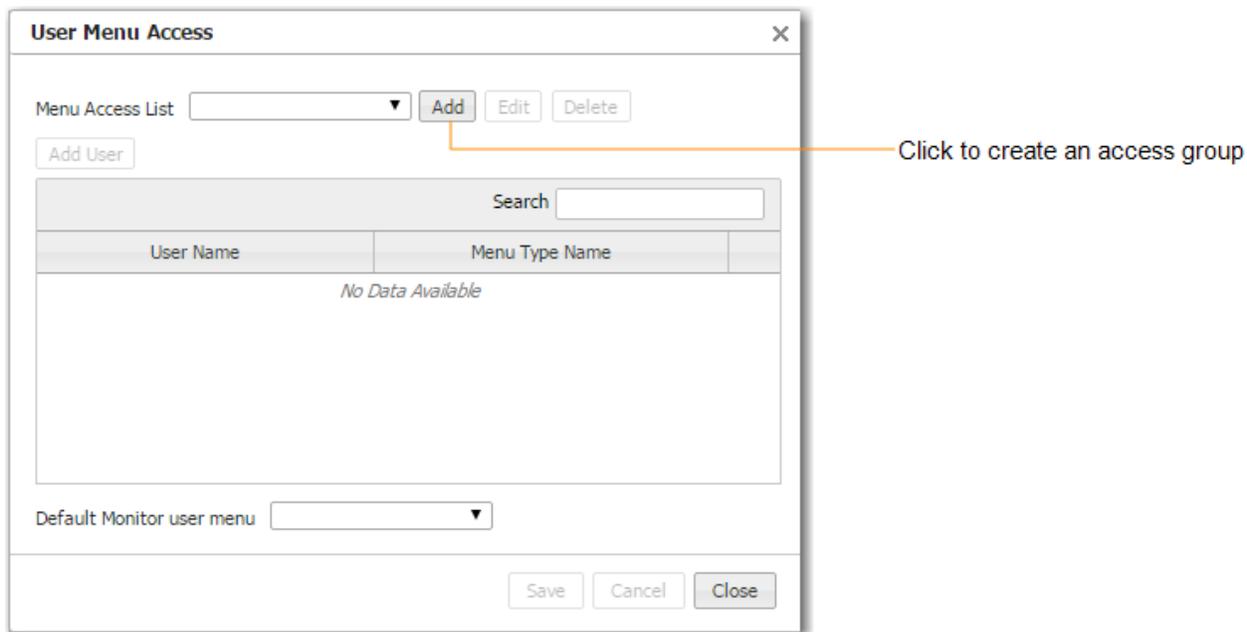
- Using an **Authentication Application** that generates time based authentication codes. If this is activated a Barcode will be generated that can be scanned to set up an authentication app like Google Authenticator for your mobile device.
  - Using your **Email** to receive authentication codes every time you log in. This requires access to your email every time you log in.
- **Password** is used at login.

- **Status** determines whether the user can log in.
- **Role** determines the user's permissions.

## User Menu Access

Orchestrator > [Orchestrator Server > Users & Authentication] User Menu Access

Use the **User Menu Access** page to create groups that have customized menu access privileges. Use these when you want to limit which menus users can see.



For each group you create, select which menus will be visible to assigned users.

**Menu Item Configuration** ✕

Allow users associated with this list to access the selected menus below.

Name

Monitoring Configuration Administration **Orchestrator** Support Miscellaneous

Select all UnSelect all

<p><i>General Settings</i></p> <p><u>Users &amp; Authentication</u></p> <p><input type="checkbox"/> Users</p> <p><input type="checkbox"/> Auth/RADIUS/TACACS+</p> <p><u>Setup</u></p> <p><input type="checkbox"/> Date/Time</p> <p><input type="checkbox"/> DNS</p> <p><input type="checkbox"/> SNMP</p> <p><input type="checkbox"/> Flow Export</p> <p><input type="checkbox"/> Logging</p> <p><input type="checkbox"/> Banners</p> <p><input type="checkbox"/> HTTPS Certificate</p> <p><input type="checkbox"/> Orchestrator Reachability</p>	<p><i>Software</i></p> <p><u>Upgrade</u></p> <p><input type="checkbox"/> System Information</p> <p><input type="checkbox"/> Software Versions</p> <p><input type="checkbox"/> Upgrade Appliances</p> <p><u>Backup &amp; Restore</u></p> <p><input type="checkbox"/> Backup Now</p> <p><input type="checkbox"/> Configuration History</p> <p><input type="checkbox"/> Restore</p> <p><u>Remove Appliances</u></p> <p><input type="checkbox"/> Remove from Orchestrator</p> <p><input type="checkbox"/> Remove from Orchestrator and&lt;br&gt;</p>	<p><i>Tools</i></p> <p><input checked="" type="checkbox"/> Synchronize</p> <p><input checked="" type="checkbox"/> Bypass</p> <p><input checked="" type="checkbox"/> Broadcast CLI</p> <p><input checked="" type="checkbox"/> Link Integrity Test</p> <p><input checked="" type="checkbox"/> Disk Management</p> <p><input checked="" type="checkbox"/> Erase Network Memory</p> <p><u>Reboot</u></p> <p><input type="checkbox"/> Appliance Reboot / Shutdown</p> <p><input type="checkbox"/> Schedule Appliance Reboot</p> <p><u>Monitoring</u></p> <p><input type="checkbox"/> Reachability Status</p> <p><input type="checkbox"/> Active Sessions</p>
--	--	---

OK Cancel

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

**User Menu Access** ✕

Menu Access List

2 Rows Search

User Name	Menu Type Name	
monitor	monitor-only	✕
dave	monitor-only	✕

Default Monitor user menu

## Remote Authentication

*Orchestrator > [Orchestrator Server > Users & Authentication] Authentication*

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

**Remote Authentication**

Local Only  
 RADIUS  
     Read-Write Privilege: 7  
     Read-Only Privilege: 0  
     Authentication Type: PAP  
 TACACS  
     Authentication Type: CHAP

Authentication Order: Remote first

	IP	Port	Secret Key
Primary			*****
Secondary			*****

Save Cancel

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

### To authenticate using RADIUS or TACACS+

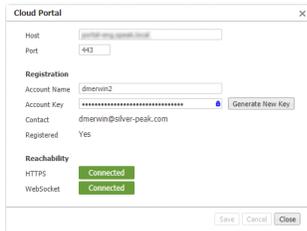
1. Select the access control protocol you want to use.
2. Under **Servers**, enter the information for a Primary server of that type. Entering a Secondary server is optional.

Field	Description
Authentication Order	Whether to use the remote map or the local map first. The default is <b>Remote first</b> .
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.
Read-Write Privilege	The lowest value at which a user has Read-Write privileges. This value must be the same as the value configured in the RADIUS server.
Read-Only Privilege	The lowest value at which a user has Read-Only privileges. This value must be the same as the value configured in the RADIUS server.
Authentication Type	When configuring to use the TACACS+ server, select the type from the drop-down list that matches what's configured on the TACACS+ server.

## Cloud Portal

*Configuration > [Overlays > Licensing] Cloud Portal*  
*Orchestrator > [Orchestrator Server > Licensing] Cloud Portal*

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



The screenshot shows a 'Cloud Portal' configuration window. It contains the following fields and sections:

- Host:** A text input field containing 'portal.silverpeak.com'.
- Port:** A text input field containing '443'.
- Registration:**
  - Account Name:** A text input field containing 'dimerin2'.
  - Account Key:** A text input field containing a masked key (dots) and a 'Generate New Key' button.
  - Contact:** A text input field containing 'dimerin@silver-peak.com'.
  - Registered:** A text input field containing 'Yes'.
- Reachability:**
  - HTTPS:** A green status indicator labeled 'Connected'.
  - WebSocket:** A green status indicator labeled 'Connected'.

At the bottom of the window are 'Save', 'Cancel', and 'Close' buttons.

- 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.

## Orchestration Settings

*Orchestrator > [Orchestrator Server > Tools] Overlay Manager and Tunnel Settings*

The **Orchestration Settings** manage Business Intent Overlays (BIOs) and the properties used to control them. It builds new tunnels and fixes existing ones.

**Orchestration Settings** ? X

Apply Overlays

Reset All Flows

Auto Save Appliance Changes

Apply Templates

Idle Time  Sec

Auto Flow Re-Classify  (0..65535) Sec

**IPSec UDP Settings**

Default Port

Increment Port By

Save Close

Field	Definition
Apply Overlays	When selected, updates all associated appliances when overlay changes are saved.
Reset All Flows	When selected, Orchestrator will automatically reset all flows whenever you edit overlays or change policies or priorities. When deselected, the flows can only be reset manually.
Auto Save Appliance	Selected by default, this automatically saves any changes made to an appliance. If you need a time delay for troubleshooting or testing, you can deselect this option to suspend automatic saving of configuration changes.
Apply Templates	When selected, updates all associated appliances when template changes are saved.
Idle Time	This determines how often the Overlay Manager surveys the network when configuration changes are not being made.

Field	Definition
<b>Auto Flow Re-Classify</b>	Specifies how the Overlay Manager waits before surveying the network when configuration changes are not being made.
<b>IPSec UDP Settings</b>	
<b>Default Port</b>	By default, Business Intent Overlays create IPSec UDP tunnels. <b>Default Port is 10002.</b> If necessary, you can configure this for an individual appliance on its <b>System Information</b> page, under <b>System Settings</b> . This is accessible from the appliance's context-sensitive menu in Orchestrator's navigation pane.
<b>Increment Port By</b>	Referenced when configuring an Edge HA (High Availability) pair. When the value is 1000, the second appliance's default port would become 11002.

## Audit Logs

*Orchestrator > Orchestrator Server > Tools > Audit Logs*

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

You can apply the following filters to your audit logs.

- You can select **Completed**, **In Progress**, or **Queued** filters to determine which actions you want to display in the table.
- You can select the following different log levels: **Debug**, **Info**, **Error** to apply to your filter.
- You can choose either **Auto Refresh** or **Pause** to refresh or pause the table. By default, the table refreshes automatically.
- You can enter in the **Record Count**. This limits the filtering criteria. The default value is 500 and 10,000 is the maximum amount you can filter.
- You can choose the name of the **Appliance** from the lists to apply as a filter.
- You can also search a wild card character (\*) as a user name and all user logs will display. If you enter any value in the user field, there will be no filter applied to the search. The following are true for audit log wild cards:
  - x\*= anything that starts with the entered value
  - \*x= anything that ends with the entered value

Record Count:  (Max 10000)
 Appliance Type to select: 
 From:  To: 
 User:  Type to select:

## Audit Logs Search

User Name	IP Address	Host Name	Action	Task Status	Results	Start Time	End Time	Queued Time	% Completed	Completion Status
OverlayManager		Albuquerque	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Albuquerque	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
Orchestration		Albuquerque	Synchronize	COMPLETED	PARTIAL, 1s, [Config: 1 (1s), State: 0 (0s)] - State before starting synchr...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
Orchestration		San-Jose	Synchronize	COMPLETED	PARTIAL, 1s, [Config: 1 (1s), State: 0 (0s)] - State before starting synchr...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
Orchestration		Salt-Lake-City	Synchronize	COMPLETED	PARTIAL, 1s, [Config: 1 (1s), State: 0 (0s)] - State before starting synchr...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
Orchestration		Paris	Synchronize	COMPLETED	PARTIAL, 0s, [Config: 1 (0s), State: 0 (0s)] - State before starting synchr...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Chennai	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Osaka	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Dallas	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		San-Antonio	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		New-York	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		San-Jose	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Geneva	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		London	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Osaka	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Salt-Lake-City	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Chennai	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Dallas	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Paris	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Tokyo	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		San-Antonio	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Minneapolis	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Miami	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		San-Jose	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Geneva	Add Overlay ACL	COMPLETED	Overlay ACLs added. Data = {"data":{"Overlay_CriticalApps":{"entry":{"1...	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success
OverlayManager		Edinburgh	SaveChanges	COMPLETED	Saved change on appliance successfully	06-Feb-19 12:00	06-Feb-19 12:00	06-Feb-19 12:00	100	Success

Field	Definition
<b>Username</b>	You can filter/search for an audit log by the user name of the appliance.
<b>IP Address</b>	The IP address of the selected appliance.
<b>Host Name</b>	The host name of the appliance the audit log is coming from.
<b>Action</b>	What you want the audit log to do.
<b>Task Status</b>	The status of the audit log task.
<b>Results</b>	The results of the audit log being searched.
<b>Start Time</b>	The time the search of the audit log started.
<b>End Time</b>	The time the search of the audit log ended.
<b>Queued Time</b>	The time the process/task was requested or scheduled in the queue.
<b>Percent Completed</b>	The percent completed of the audit log task.
<b>Completion Status</b>	Whether the task has been completed.

## Pause Orchestration List

*Orchestrator > [Orchestrator Server > Tools] Pause Orchestration List*

When troubleshooting, you can pause Orchestration for the appliances in question.

**Pause Orchestration List** [X]

Orchestration will be suspended for these appliances, but their status will continue to be monitored.

Type to select

Host Name	IP	Version
<i>No Data Available</i>		

## Tunnel Settings Tab

*Orchestrator > Orchestrator Server > Tools > Tunnels Settings*

Use this page to manage the properties for those tunnels created by Orchestrator. This tab provides tunnel settings for General, IKE, and IPsec for MPLS, Internet, and LTE WAN Interface labels.

*Tunnel Settings for Overlays and Tunnel Groups*

General	
<b>Mode</b>	Indicates whether the tunnel protocol is <b>ipsec</b> , <b>ipsec_udp</b> , <b>udp</b> , or <b>gre</b> . If you select IPsec, you can specify the IKE version in the <b>IKE</b> tab.
<b>Auto Max BW Enabled</b>	Allows the appliances to auto-negotiate the maximum tunnel bandwidth.
<b>Auto Discover MTU Enabled</b>	Allows the appliances to auto-negotiate the maximum tunnel bandwidth.
<b>MTU</b>	(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. <b>Auto</b> allows the tunnel MTU to be discovered automatically, and it overrides the MTU setting.
Packet	
<b>Reorder Wait</b>	The maximum time the appliance holds an out-of-order packet when attempting to reorder. The packets can come from either the same or different path, or from the FEC correction engine. 100ms is the default value and should be adequate for most situations. If the reorder wait time exceeds 100ms (or the set value), the packet will be delivered out of order.
<b>FEC</b>	(Forward Error Correction) can be set to <b>enable</b> , <b>disable</b> , and <b>auto</b> .
<b>FEC Ratio</b>	When FEC is set to <b>auto</b> , this specifies the maximum ratio. The options are 1:2, 1:5, 1:10, or 1:20.
Tunnel Health	
<b>Retry Count</b>	Number of failed keep-alive messages that are allowed before the appliance brings the tunnel down.
<b>DSCP</b>	Determines which DSCP marking the keep-alive messages should use.
Fastfail Thresholds	

---

**Fastfail Thresholds**

Fastfail thresholds determine how quickly to disqualify a tunnel from carrying data when multiple tunnels are carrying data between two appliances.

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:

$$Base = 200ms$$

$$N = 2$$

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 keep-alive signal doesn't receive a reply. The options are **disable**, **enable**, and **continuous**. If the disqualified tunnel subsequently receives a keep-alive reply, its recovery is instantaneous.
  - If set to **disable**, keep-alives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost.
  - If set to **enable**, keep-alives are sent every second, and a missed reply increases the rate at which keep-alives are sent from 1 per second to 10 per second. Failover occurs after 1 second.
  - When set to **continuous**, keep-alives 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.

---

**IPsec Encryption Algorithm**

For encrypting tunnel data. Choose from **auto**, **AES-256**, or **AES-128**.

---

**Latency**

The amount of latency measure in MS.

---

<b>Loss</b>	The amount of data lost measured in percent.
<b>Jitter</b>	The amount of jitter measured in MS.
<b>FastFail Wait-Time Base Offset</b>	The base time used when you calculate the fastfail timeout.
<b>FastFail RTT Multiplication Factor</b>	The multiplier in the formula used to calculate the fastfail timeout.

<b>IKE</b>	
<b>Authentication Algorithm</b>	This is for setting tunnel authentication. Choose from <b>SHA-1</b> , <b>SHA2-256</b> , <b>SHA2-384</b> , or <b>SHA2-512</b> .
<b>Encryption Algorithm</b>	Specifies the encryption algorithm used for the Phase 1 negotiation. Choose from <b>AES-256</b> , <b>AES-128</b> , or <b>auto</b> .
<b>Diffie-Hellman Group</b>	The Diffie-Hellman group used for IKE SA negotiation.
<b>Lifetime</b>	The lifetime of IKE SA.
<b>Dead Peer Detection</b>	Delay time: the interval in seconds to check the IKE peer. Retry Count: amount of times you specify the system to retry if a connection has failed.
<b>Phase 1 Mode</b>	Defines the exchange mode for Phase 1. The options are <b>Main</b> or <b>Aggressive</b> . If IKEv2 is selected, the default mode is aggressive.
<b>IKE Version</b>	The IKE major version. Select either <b>IKEv1</b> or <b>IKEv2</b> .

<b>IPSec</b>	
<b>Authentication Algorithm</b>	The authentication algorithm used by IPSec SA. Choose from <b>SHA-1</b> , <b>SHA2-256</b> , <b>SHA2-384</b> , or <b>SHA2-512</b> .
<b>Encryption Algorithm</b>	Specifies the encryption algorithm used for the Phase 1 negotiation. Choose from <b>AES-256</b> , <b>AES-128</b> , or <b>auto</b> .
<b>Enable IPsec Anti-replay Window</b>	Select if you want to enable the IPSec anti-replay window. If selected, protection is provided against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet. The default window size is 64 packets.
<b>Lifetime</b>	The lifetime of IKE SA.

---

**Perfect Forward Secrecy Group**

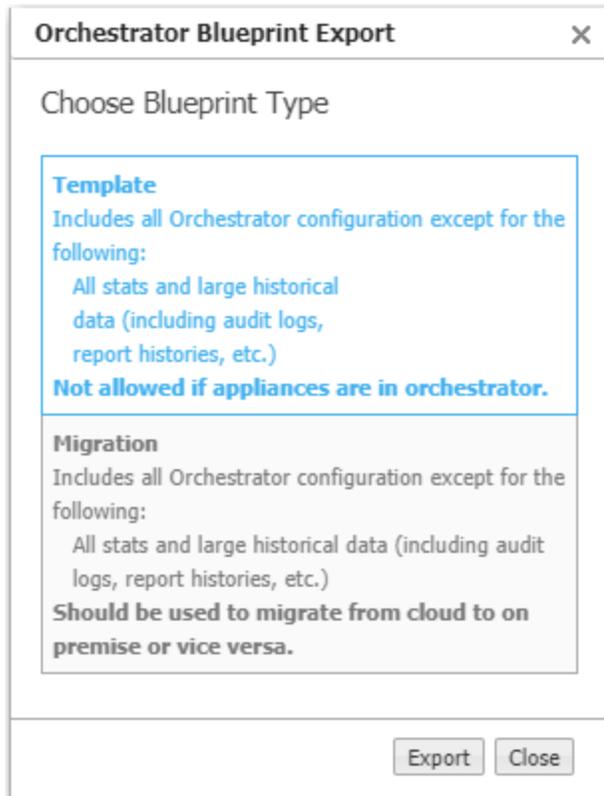
Specifies the Diffie Hellman Group  
exponentiations used for IPSec SA negotiation.

---

## Orchestrator Blueprint Export

*Orchestrator > [Orchestrator Server > Tools] Orchestrator Blueprint Export*

Use this page to create and export a configuration that Orchestrator-SP can use as a template for other Orchestrators.



## Brand Customization

*Orchestrator > [Orchestrator Server > Tools] Brand Customization*

Use this menu to customize the branding elements of Orchestrator's user interface.

**Brand Customization** [X]

**Login Page**

Background Image  
Silver Peak Upload Custom

Logo Image 210 x 70px  
Silver Peak Upload Custom

Welcome Text  
Show Hide

Patent Text  
Show Hide

**Browser**

Favicon  
Silver Peak Upload Custom

Tab Title  
Silver Peak Custom

**Application/Header/Footer**

Logo Image 200 x 60px  
Silver Peak Upload Custom

"Unity Orchestrator" Text  
Show Hide

Tree Icon  
Silver Peak Upload Custom

Copyright/Footer  
Silver Peak Custom

**Support Menu**

Support Link URL  
Silver Peak Custom

Apply Close

## Maintenance Mode

You can put one or more appliances in maintenance mode by selecting the specific appliance in the tree. Upon approval, the appliances are added to the maintenance list. You can also put an appliance in maintenance mode by searching "**Maintenance Mode**" in the search bar or by right-clicking on any appliance and selecting **Maintenance Mode**. Complete the following steps to add an appliance to maintenance mode.

1. Navigate to **Maintenance Mode** in Orchestrator.
2. Select **Add**. The **Configure Maintenance Mode** window opens.
3. Check **Pause Orchestration** if you want to pause orchestration.
4. Check **Suppress Alarms** if you want to suppress alarms associated with this appliance while in maintenance mode.
5. Select **OK**.
6. Select **Save**.

**NOTE** The appliance goes into maintenance mode if you pause orchestration and/or suppress all alarms.

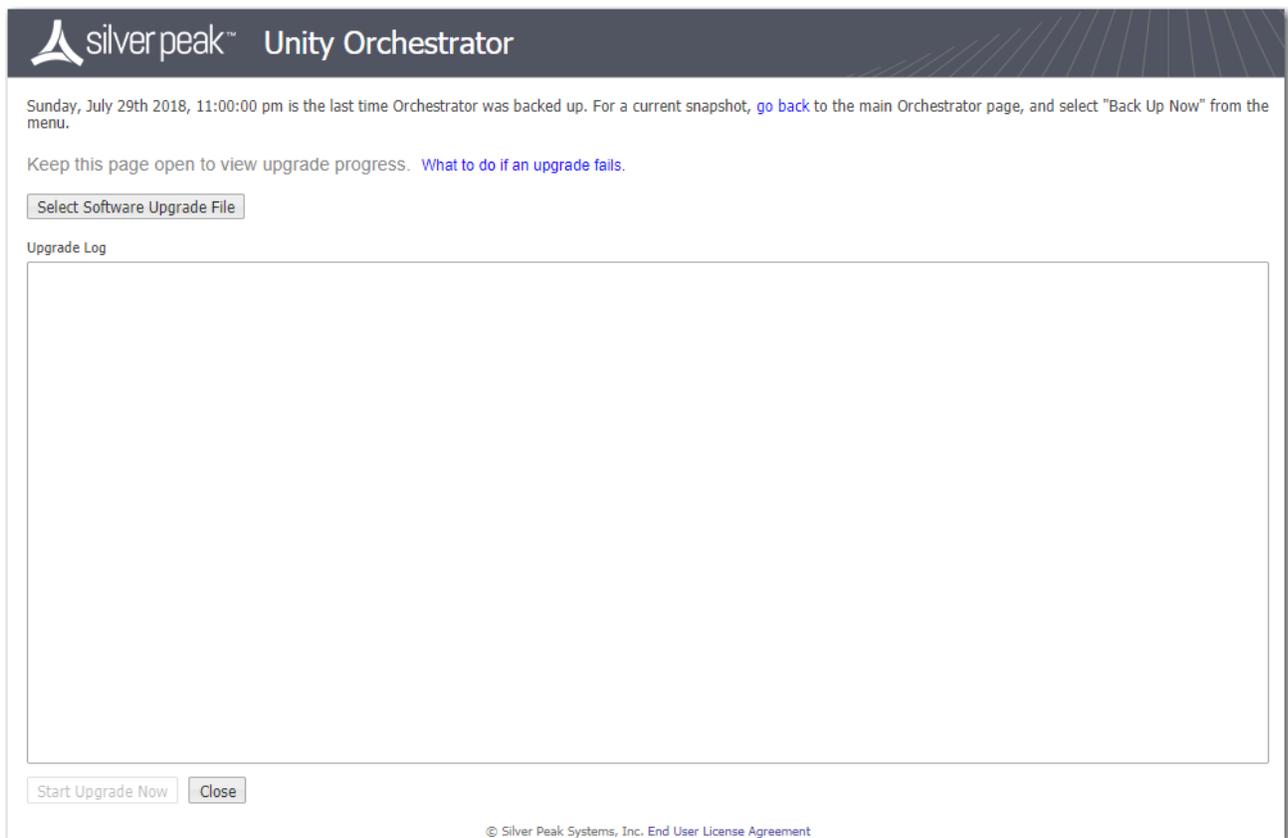
Field	Definition
Host Name	The host name of the appliance you are adding to maintenance mode.
Alarms	Whether you chose to suppress or not suppress your alarms while the appliance is in maintenance mode.
Orchestration	If paused, all orchestration is paused on the selected appliance, except IPSec UDP Tunnel Key material.
IP	The IP address of the appliance in maintenance mode.
Version	The current version of the appliance.

## Upgrading Orchestrator Software

*Orchestrator > [Software & Setup > Upgrade] Upgrade Orchestrator*

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

We recommend that you allocate enough RAM and CPU Cores prior to upgrading to appropriately increase the processing power required to support the latest features.



The screenshot shows the 'Unity Orchestrator' software interface. At the top, there is a dark header with the Silver Peak logo and the text 'Unity Orchestrator'. Below the header, a message states: 'Sunday, July 29th 2018, 11:00:00 pm is the last time Orchestrator was backed up. For a current snapshot, [go back](#) to the main Orchestrator page, and select "Back Up Now" from the menu.' Below this, a note says: 'Keep this page open to view upgrade progress. [What to do if an upgrade fails.](#)' There is a button labeled 'Select Software Upgrade File'. Underneath, the text 'Upgrade Log' is followed by a large, empty rectangular box intended for displaying the upgrade progress. At the bottom of the page, there are two buttons: 'Start Upgrade Now' and 'Close'. A small copyright notice at the very bottom reads: '© Silver Peak Systems, Inc. End User License Agreement'.

## Checking for Orchestrator and Appliance Software Updates

*Orchestrator > [Software & Setup > Upgrade] Check for Updates*

These pages show what appliance and Orchestrator server software is available for download.

**Check for Updates** ✕

Orchestrator Releases

Release	Type	Release Date	Description	Release Notes
8.3.0.00000	BETA	03-Nov-17 00:00		
8.4.0.35900	BETA	13-May-18 00:00	Test for orchestrator feature GMS-11402	<a href="#">Download</a>
99.99.99.35870	BETA	13-May-18 00:00	Test for orchestrator feature GMS-11402	<a href="#">Download</a>
99.99.99.36894	BETA	13-May-18 00:00	Test for orchestrator feature GMS-11402	<a href="#">Download</a>

VXOA Releases

Release	Type	Release Date	Description	Release Notes
0.0.0.0_67610	BETA	09-Nov-17 00:00	KR test vxoa image	
0.0.0.0_67847	BETA	27-Nov-17 00:00	KR test	
8.1.7.1_68811	GA	07-Feb-18 00:00		
8.1.7.3_69551	BETA	09-Apr-18 00:00	rma testing	
8.1.7.7_70949	GA	15-May-18 00:00	test for upgrade from portal image	
8.1.7.7_72000	BETA	15-May-18 00:00	For testing purpose only	

Go to Support Portal to Download Close

## Backing Up on Demand

*Orchestrator > [Software & Setup > Backup] Backup Now*

Use this page to backup the Orchestrator database on demand.

Orchestrator Backup

Protocol: FTP

Hostname: 172.23.43.13

Username: root

Password: \*\*\*\*\*

Directory: /

Port: 21

Max backups to retain: 10

Status Log

Test Backup Now Close

FTP Connection to Host: 172.23.43.13, Port No: 21, Directory: / with username: root was successful.

## Scheduling Orchestrator Database Backup

*Orchestrator > [Software & Setup > Backup] Schedule Backup*

Use this page to schedule a backup the Orchestrator database.

### Schedule Orchestrator Backup

[View Currently Scheduled Jobs](#)

**Destination**

Protocol:

Hostname:

Username:

Password:

Directory:

Port:

Max backups to retain:

**Schedule**

Schedule:  [Edit](#)

Description:

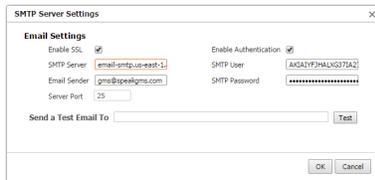


**TIP** To specify the timezone for scheduled jobs and reports, go to **Orchestrator > [Software & Setup > Setup] Timezone for Scheduled Jobs**.

## SMTP Server Settings

*Orchestrator > [Software & Setup > Setup] SMTP Server Settings*

For permanent, private email delivery, change the SMTP (Simple Mail Transfer Protocol) server and settings to your company's SMTP settings.

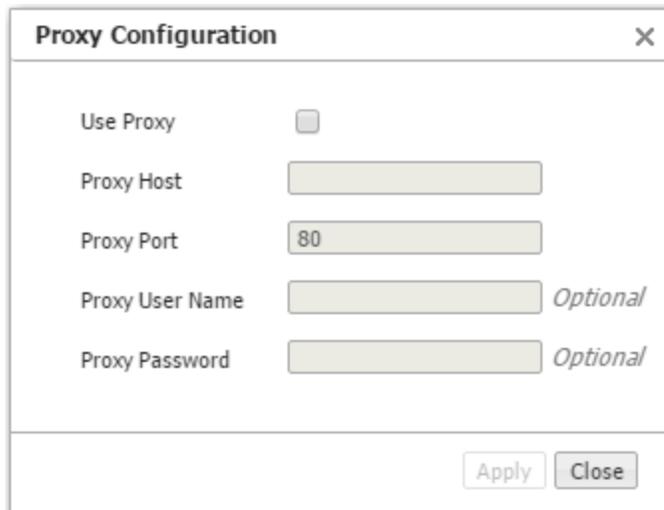


- If a test email doesn't arrive within minutes, check your firewall.
- After configuring the SMTP settings, you can specify email recipients for:
  - **alarms** (Monitoring > Alarms > Alarm Recipients), and
  - **reports** (Monitoring > [Reporting] Schedule & Run Reports)

## Proxy Configuration

*Orchestrator > [Software & Setup > Setup] 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 'Proxy Configuration' dialog box with a title bar containing the text 'Proxy Configuration' and a close button 'X'. 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 containing the value '80'.
- Proxy User Name**: A text input field followed by the text *Optional*.
- Proxy Password**: A text input field followed by the text *Optional*.

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

## Orchestrator's HTTPS Certificate

*Orchestrator > [Software & Setup > Setup] HTTPS Certificate*

Orchestrator includes a self-signed certificate that secures the communication between the user's browser and Orchestrator. You also have the option to install your own custom certificate, acquired from a CA authority.

**HTTPS Certificate**

Intermediate Certificate File  No file chosen

Certificate File  No file chosen

Private Key File  No file chosen

**Current Certificate**

Issuer *Silverpeak GMS*

Issued to *Silverpeak GMS*

Expiration *Fri Nov 23 21:05:40 GMT 2114*

For a custom certificate, to use with Orchestrator:

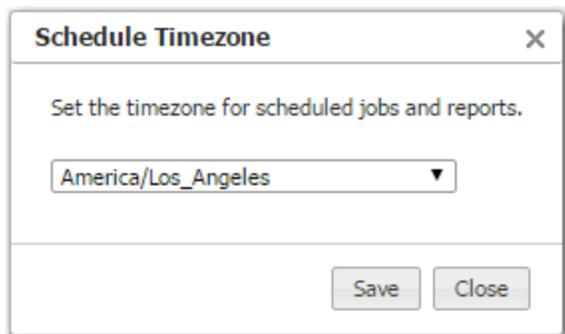
1. Consult with your IT security team to generate a certificate signing request (CSR), and submit it to your organization's chosen SSL Certificate Authority (CA).
  - Examples of Certificate Authorities include GoDaddy, Verisign, Comodo, Symantec, Microsoft Entrust, GeoTrust, etc.
  - For a list of what Silver Peak supports, see [Silver Peak Security Algorithms](#).
  - All certificate and key files must be in **PEM** format.
2. After the Certificate Authority provides a CA-verified certificate:
  - If your IT security team advises the use of an Intermediate CA, then use an **Intermediate Certificate File**. Otherwise, skip this file.
  - Load the **Certificate File** from the CA.
  - Upload the **Private Key File** that was generated as part of the CSR.

3. To associate the CA verified certificate for use with Orchestrator, click **Upload**.

## Timezone for Scheduled Jobs

*Orchestrator > [Software & Setup > Setup] Timezone for Scheduled Jobs*

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



The image shows a dialog box titled "Schedule Timezone" with a close button (X) in the top right corner. The dialog contains the instruction "Set the timezone for scheduled jobs and reports." Below this is a dropdown menu with "America/Los\_Angeles" selected. At the bottom of the dialog are two buttons: "Save" and "Close".

## Orchestrator Statistics Configuration

Orchestrator > [Software & Setup > Setup] Statistics Configuration

Use this tab to specify which types of statistics you want the Orchestrator to collect from the appliances.

Statistics Configuration ×

Statistics Configuration ? ↻

16 Rows Search

Stats Type	Enable
Tunnel	<input checked="" type="checkbox"/>
Application	<input checked="" type="checkbox"/>
Flow	<input checked="" type="checkbox"/>
Drops	<input checked="" type="checkbox"/>
Port	<input checked="" type="checkbox"/>
Jitter	<input checked="" type="checkbox"/>
Interface	<input checked="" type="checkbox"/>
Differentiated Services Code Point (DSCP)	<input checked="" type="checkbox"/>
Domain Name Server (DNS)	<input checked="" type="checkbox"/>
Dynamic Rate Control (DRC)	<input checked="" type="checkbox"/>
Shaper	<input checked="" type="checkbox"/>
TopTalkers	<input checked="" type="checkbox"/>
TrafficClass	<input checked="" type="checkbox"/>
Behavioral	<input checked="" type="checkbox"/>
Overlay-Interface-Transport	<input checked="" type="checkbox"/>
Mean Opinion Score (MOS)	<input checked="" type="checkbox"/>

Apply Restore Defaults Cancel

## Appliance Statistics Configuration

Orchestrator > [Software & Setup > Setup] Appliance Stats Configuration

This screen displays the default values for appliance properties.

**IMPORTANT:** Changing the default values of these settings is not recommended without consulting Silver Peak.

### Stats Configuration

IMPORTANT: Changing the default values of these settings is not recommended without consulting Silver Peak.

Property Name	Property Value
minuteRetention	1440
verticalRetention	2
app max_items	100
app evict_enable	false
port max_items	100
port evict_enable	false
dns max_items	100
dns evict_enable	false
ip max_items	100
ip evict_enable	false
behavioral max_items	100
behavioral evict_enable	false
flows_csv_enable	false

Apply Restore Defaults Close

## Orchestrator Advanced Properties

*Orchestrator > [Software & Setup > Setup] Advanced Properties*

**IMPORTANT:** Changing the default values of these settings is not recommended without consulting Silver Peak.

**Orchestrator Advanced Properties** ✕

**IMPORTANT:** Changing the default values of these settings is not recommended without consulting Silver Peak.

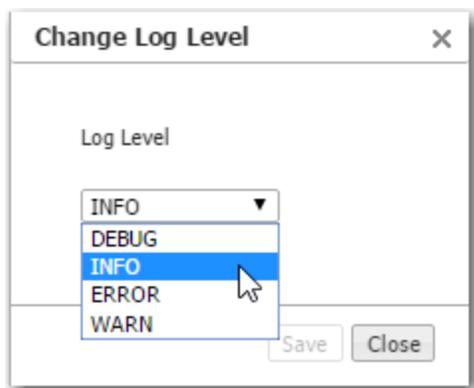
37 Rows Search

<i>Property Name</i> ▲	Property Value
ParallelActionTasks	50
ParallelOrchestrationTasks	50
ParallelReachabilityTasks	20
ParallelStatsTasks	20
bridgeCacheExpireTime	120
dbPoolConnectionTimeout	30000
dbPoolIdleTimeout	120000
dbPoolLeakDetectionThreshold	300000
dbPoolMaxConnectionLifeTime	3000000
dbPoolMaxConnections	1000
dbPoolMinimumIdleConnections	10
dbPoolValidationTimeout	3000
denyApplianceOnDelete	true
emailImagesMaxSize	10
excludeTables	true
excludedTableNames	dailyapp,dailydrc,dailydrops,dailydscp,d...
failedLoginAttemptThreshold	5
fastRecordGenRate	100
jettyAcceptQueueSize	1000
jettyIdleTimeout	60000

## Changing Orchestrator's Log Level

*Orchestrator > [Software & Setup > Setup] Change Log Level*

Use this form to change what level of server-side Orchestrator logs are retained. The default is **INFO**.



### Minimum Severity Levels

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

Level	Definition
<b>ERROR</b>	An error. This is a non-urgent failure.
<b>WARNING</b>	A warning condition. Indicates an error will occur if action is not taken.
<b>INFORMATIONAL</b>	Informational. Used by Silver Peak for debugging.
<b>DEBUG</b>	Used by Silver Peak for debugging

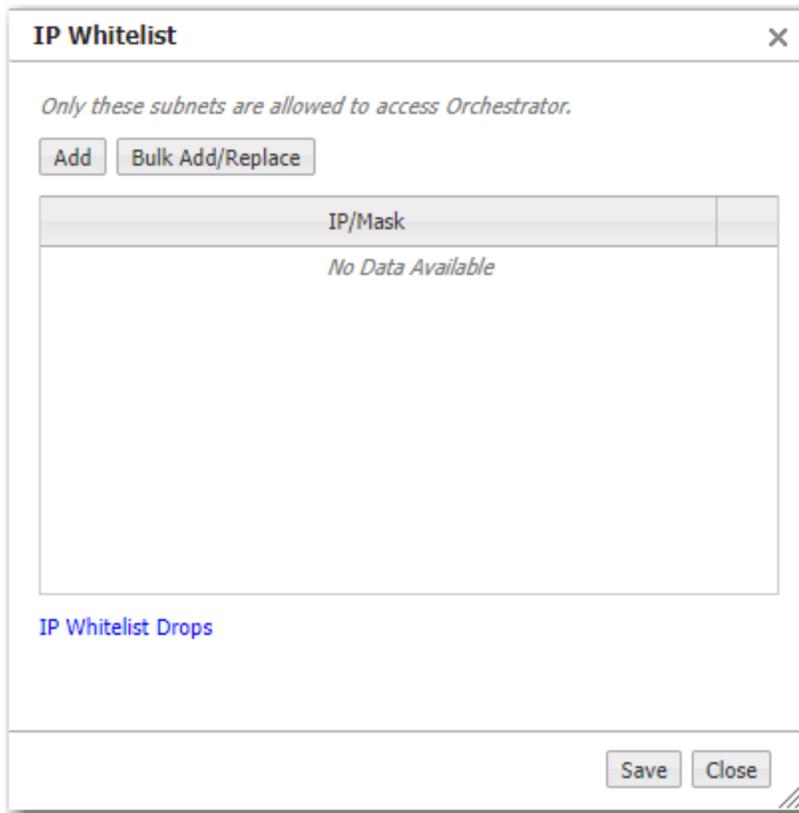
- The bolded part of the name is what displays in Silver Peak's logs.
- If you select **INFO** (the default), then the log records any event with a severity of **INFO**, **WARNING**, and **ERROR**.
- These are purely related to event logging levels, **not** alarm severities, even though some naming conventions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

## IP Whitelist

Orchestrator > [Software & Setup > Setup] IP Whitelist

**IP Whitelist** is a feature that restricts access to Orchestrator to a specified list of source subnets.

If a source IP address changes (for example, with NAT IP), then users can get locked out Orchestrator.



To view a list of traffic that's been dropped because of these restrictions, click **IP Whitelist Drops**.

## Orchestrator's Getting Started Wizard

*Orchestrator > [Software & Setup > Setup] Configuration Wizard*

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

This 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 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 Orchestrator and the 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 page, Orchestrator's wizard reappears at your next login.

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

## Customer and Technical Support

When working with Customer Support, these tabs facilitate your opening a support case. They also provide Customer Support with data and reports needed for troubleshooting network issues.

## Tech Support - Appliances

*Support > Technical Assistance > Tech Support - Appliances*

Use this tab to open/create cases, upload files to Silver Peak Support, and download selected files to Orchestrator.

You can filter between the five different file types: All, Logs, Sys Dump, Snapshot, and TCP Dump. The table in this tab displays the following:

Field	Description
Appliance Name	The name of the appliance that the logs are coming from
File type	The type of file.
File Name	The name of the file.
Last Modified	The date the file was modified last.
File Size	The size of the file.

### Download to Orchestrator

Complete the following steps if you want to download selected files to your local Orchestrator server.

1. Go to the **Tech Support - Appliance** tab in Orchestrator.
2. Select a file in the in the table you want to download.
3. Select **Download to Orchestrator**.

The Monitor Transfer Progress window opens that provides a status of the file downloads. You can also cancel a download at any time by selecting **Cancel**.

Field	Definition
Source	The source where the files are coming from.
Files	The files selected to download.
Start Time	The start time the files were downloaded.
End Time	The end time the files were downloaded.
Transferred	The percentage representing how much the files have been downloaded.
Status	The status of the download (in progress or download finished).
Cancel	Select cancel at any time to interrupt and stop the download.

## Tech Support - Orchestrator

*Support > Technical Assistance > Tech Support - Orchestrator*

Use this tab to view and manage logs for Orchestrator, create cases, and upload files to Silver Peak Support. On this page, you can also select files you want to download to your local desktop and can filter between the five different file types: All, Logs, Sys Dump, and Appliances.

Tech Support - Orchestrator ?

Filter All Logs Sys Dump Appliances Create Case Upload Selected Files Download selected Files Generate Sys Dump 🕒 1 min

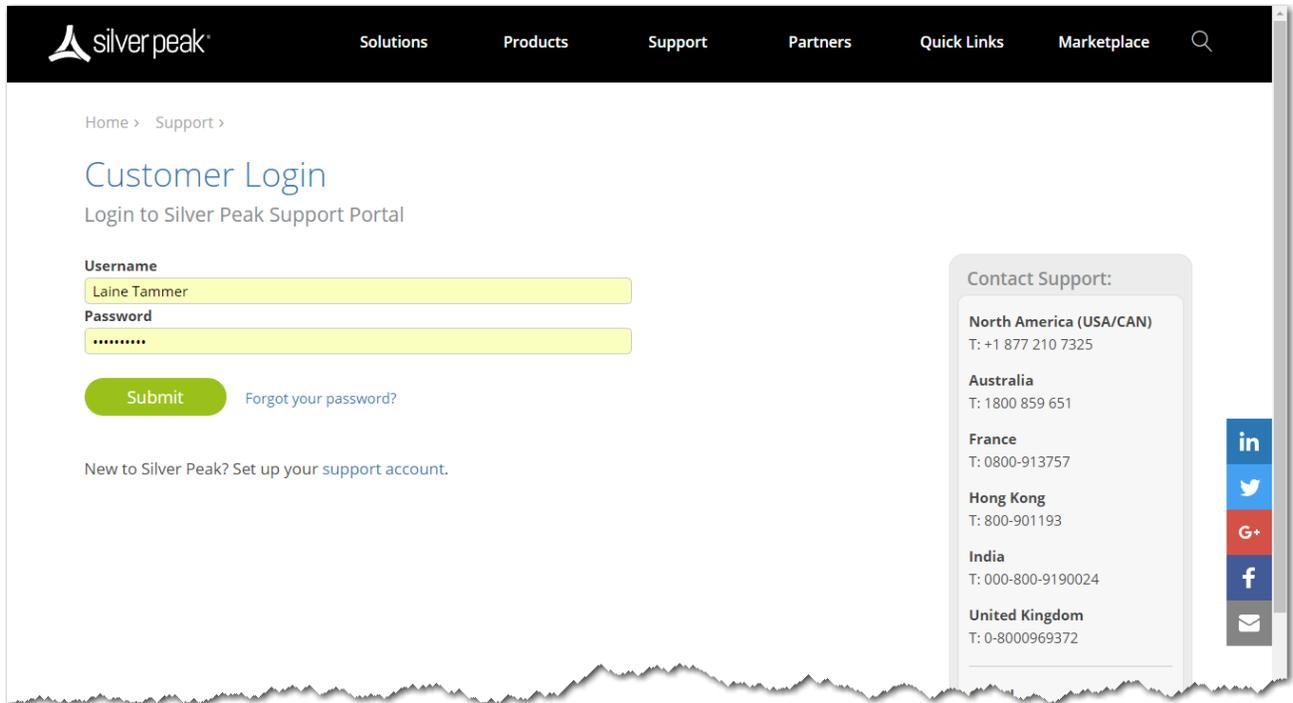
Source	File Type	File Name	Last Modified	File Size
Orchestrator	Logs	server.log	20-Feb-19 12:36	289.7M
Orchestrator	Logs	jetty-request.log	20-Feb-19 12:36	22.1M
Orchestrator	Logs	mysqld.log	20-Feb-19 12:36	2.2G
Orchestrator	Logs	metrics.log	20-Feb-19 12:36	3.4M
Orchestrator	Logs	quartz.log	20-Feb-19 11:49	1.9M
Orchestrator	Logs	mysql_slow_queries.log	20-Feb-19 11:19	9.1K
Orchestrator	Logs	ifsetup.log	20-Feb-19 11:16	212
Orchestrator	Logs	reportjob.log	20-Feb-19 00:52	428.3K
Orchestrator	Logs	https_localhost_0_localstorage	20-Feb-19 00:52	65.5K
Orchestrator	Logs	server-2019-02-19-1.log.gz	20-Feb-19 00:00	28.6M
Orchestrator	Logs	jetty-request-2019-02-19-1.log.gz	20-Feb-19 00:00	980.4K
Orchestrator	Logs	metrics-1.log	19-Feb-19 20:23	10.5M
Orchestrator	Logs	server-2019-02-18-1.log.gz	19-Feb-19 00:00	28.7M
Orchestrator	Logs	jetty-request-2019-02-18-1.log.gz	19-Feb-19 00:00	714.6K
Orchestrator	Logs	server-2019-02-17-1.log.gz	18-Feb-19 00:00	29.5M

1. Select a file you want to download to your local under **File Type** in the table.
2. Select **Download Selected Files**.
3. The **Download Selected Files** window appears. Select **Download**.

## Logging into the Support Portal

*Support > [Technical Assistance] Support Portal Log-in*

When you have a Silver Peak account and need technical assistance or customer support, select **Support > Tech Support**. The following page opens in a separate browser tab.

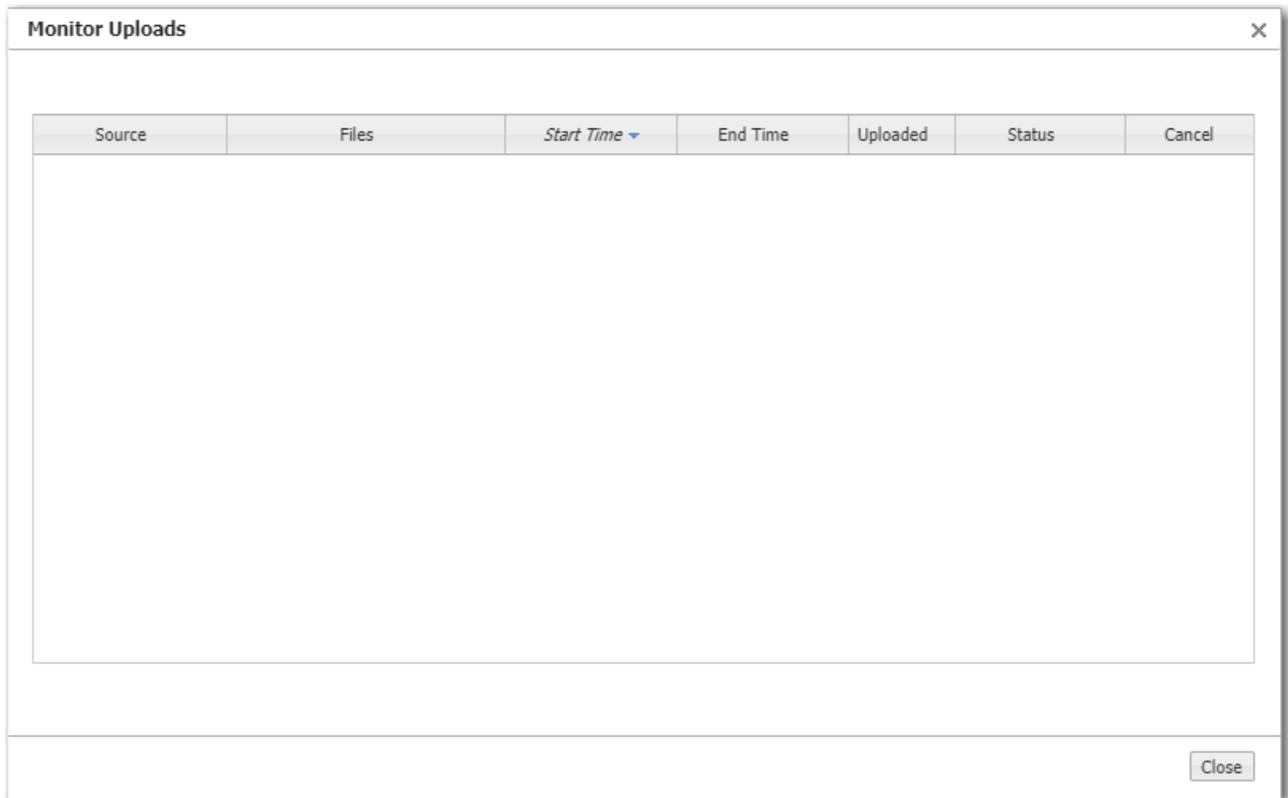


You can also access this page directly by going Silver Peak's web page and selecting **Support > Customer Login** from the menu bar.

## Monitoring Uploads

*Support > [Technical Assistance] Monitor Uploads*

This table displays the current status of any files being uploaded to Support.

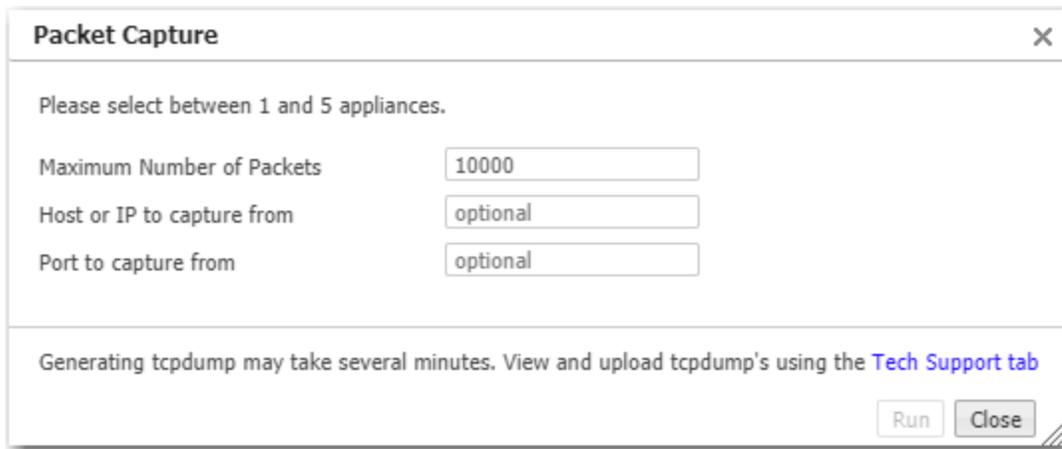


Source	Files	Start Time	End Time	Uploaded	Status	Cancel
--------	-------	------------	----------	----------	--------	--------

## Packet Capture

*Support > [Technical Assistance] Packet Capture*

When requested by Support, use this screen to capture packets from one to five appliances, selected in the navigation pane.

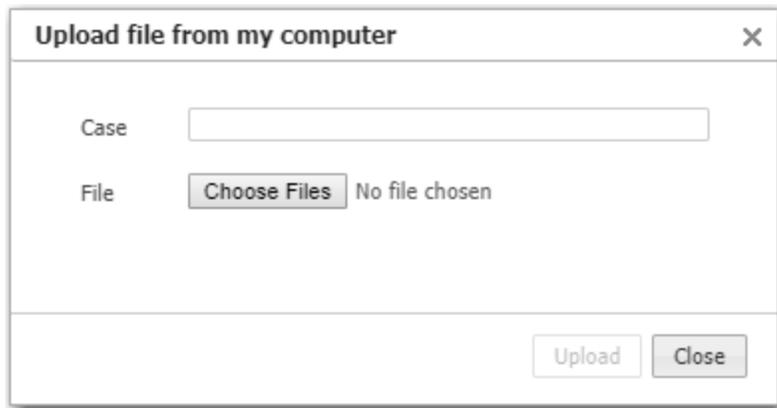


The screenshot shows a dialog box titled "Packet Capture" with a close button (X) in the top right corner. The main content area contains the instruction "Please select between 1 and 5 appliances." followed by three input fields: "Maximum Number of Packets" with the value "10000", "Host or IP to capture from" with the value "optional", and "Port to capture from" with the value "optional". At the bottom of the dialog, there is a note: "Generating tcpdump may take several minutes. View and upload tcpdump's using the [Tech Support tab](#)". In the bottom right corner, there are two buttons: "Run" and "Close".

## Upload Local Files

*Support > [Technical Assistance] Upload Local Files*

Use this dialog to upload files related to your Support case from your computer.



The image shows a dialog box titled "Upload file from my computer" with a close button (X) in the top right corner. The dialog contains the following elements:

- A "Case" label followed by an empty text input field.
- A "File" label followed by a "Choose Files" button and the text "No file chosen".
- At the bottom right, there are two buttons: "Upload" and "Close".

## Create a Support Case

*Support > [Technical Assistance] Create Case*

Use this file to create an Support case.

You'll receive a case number and instructions for what to do next.

**Create Case** [X]

Name

Email

Phone

Appliance

Case Priority

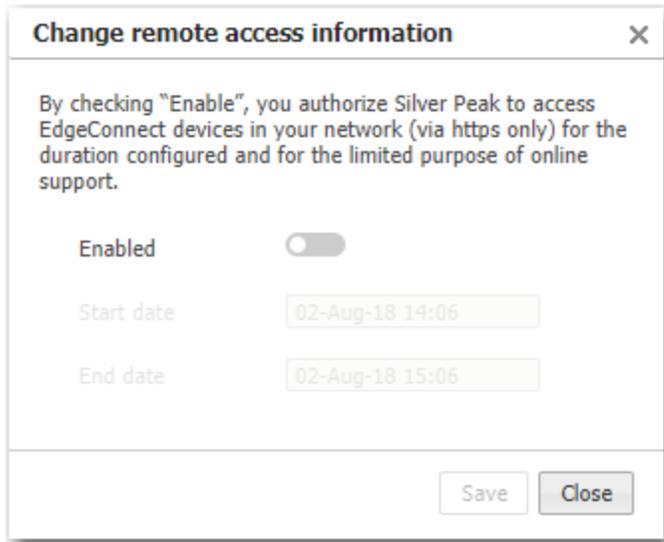
Subject

Description

## Remote Access

*Support > [Technical Assistance] Remote Log Receiver*

When working with Silver Peak Support to troubleshoot, you may be asked to allow access to you EdgeConnect devices during the online support session.



The dialog box is titled "Change remote access information" and has a close button (X) in the top right corner. It contains the following text: "By checking 'Enable', you authorize Silver Peak to access EdgeConnect devices in your network (via https only) for the duration configured and for the limited purpose of online support." Below this text are three fields: "Enabled" with a toggle switch currently turned off, "Start date" with a text input field containing "02-Aug-18 14:06", and "End date" with a text input field containing "02-Aug-18 15:06". At the bottom right of the dialog are two buttons: "Save" and "Close".

# Partition Management

Support > [Technical Assistance] Partition Management

You can use this table to regain Orchestrator disk space by selectively eliminating stats you no longer need.

Partition Management x

Partition Management ? 2 mins

807 Rows						Search <input style="width: 100px;" type="text"/>
Table Name	Partition Name	Rows	Size	Start Time	End Time	
actionlog	defaultPartition	1400427	2.8 GB			
actionlog	p1524096000	0	180 KB	20-Oct-17 17:00	18-Apr-18 17:00	
actionlog	p1508544000	0	180 KB	23-Apr-17 17:00	20-Oct-17 17:00	×
actionlog	p1492992000	0	180 KB	25-Oct-16 17:00	23-Apr-17 17:00	×
actionlog	p1477440000	0	180 KB	28-Apr-16 17:00	25-Oct-16 17:00	×
actionlog	p1461888000	0	180 KB	31-Oct-15 17:00	28-Apr-16 17:00	×
actionlog	p1446336000	0	180 KB	04-May-15 17:00	31-Oct-15 17:00	×
actionlog	p1430784000	0	213 KB		04-May-15 17:00	×
dailyapp	defaultPartition	0	74 KB			
dailyapp	p1524096000	0	74 KB	20-Oct-17 17:00	18-Apr-18 17:00	
dailyapp	p1508544000	0	74 KB	23-Apr-17 17:00	20-Oct-17 17:00	×
dailyapp	p1492992000	0	74 KB	25-Oct-16 17:00	23-Apr-17 17:00	×
dailyapp	p1477440000	0	74 KB	28-Apr-16 17:00	25-Oct-16 17:00	×
dailyapp	p1461888000	0	74 KB	31-Oct-15 17:00	28-Apr-16 17:00	×
dailyapp	p1446336000	0	74 KB	04-May-15 17:00	31-Oct-15 17:00	×

## Remote Log Receivers

*Support > Technical Assistance > Remote Log Receiver*

This table lists all configured remote log receivers that are sent and managed by Orchestrator. You can choose between sending your data with four different types of receivers: HTTP, HTTPS, KAFKA, and SYSLOG. Each receiver employs a different mechanism for supporting asynchronous notifications. After you determine which remote receiver you want to use to send your data, you can configure specific settings for that receiver.

Edit	Log Type	Name	Receiver Type	Hostname/URL	Enabled	
	Audit Log	Splunk	SYSLOG	10.17.9.10:514	Yes	
	Alarm	Splunk	SYSLOG	10.17.9.10:514	Yes	

Complete the following instructions for adding a receiver.

1. Select **Add Receiver**.
2. Select the type of receiver you want to use from the list.
3. Depending on which receiver you choose, a settings pop-up will appear. Enter the appropriate information for each receiver. See the following tables below for each receiver's settings.
4. Select **Save**.

### HTTP Receiver Settings

Field	Definition
<b>Enable Receiver</b>	Check this box to enable the selected receiver.
<b>Name</b>	The name of the receiver the logs are going to.
<b>Log type</b>	Select the type of log from the list you want to apply.
<b>URL</b>	The URL served by HTTP/HTTPS log server that Orchestrator will send log data with POST REST calls.
<b>User name</b>	The user name used in Basic Authentication when making REST calls (Optional)
<b>Password</b>	Password used in Basic Authentication when making REST calls (Optional)
<b>Repeat Password</b>	Your password repeated.

## HTTPS Receiver Settings

Field	Definition
Enable Receiver	Check this box to enable the selected receiver.
Name	The name of the receiver the logs are going to.
Log type	Select the type of log from the list you want to apply.
URL	The URL of the HTTPS Receiver.
User name	The user name used in Basic Authentication when making REST calls (Optional).
Password	The password used in Basic Authentication when making REST calls (Optional).
Repeat Password	Your password repeated.

## KAFKA Receiver Settings

Field	Definition
Enable Receiver	Check this box to enable the selected receiver.
Name	The name of the receiver the logs are going to.
Log type	Select the type of log from the list you want to apply.
Topic	The topic name on KAFKA Receiver
Bootstrap Servers	The domain name served by KAFKA Receiver. e.g. "xxx.com:9092", "1.1.1.1:9092"
Acks	Defines the amount of KAFKA servers that acknowledge a message before considering the message delivered. <ul style="list-style-type: none"> <li>■ acks=0: expect no acknowledge</li> <li>■ acks=1: only leader server must acknowledge</li> <li>■ ack=all: all servers must acknowledge.</li> </ul>
Retries	The amount of times KAFKA will try before returning an error.
Batch Size	The multiple messages KAFKA will produce until the batch size is exceeded.
Buffer Size	The maximum memory size that can be used for buffering messages. When buffer size is exceeded, a message will be blocked.
Linger Time	The amount of time that KAFKA will wait before sending next message batch.

## SYSLOG

Field	Definition
Enable Receiver	Check this box to enable the selected receiver.
General Settings	
Log Type	The type of log being sent to the SYSLOG receiver.

Field	Definition
Protocol	The protocol being used between devices.
Hostname	The hostname of the SYSLOG receiver to identify the device.
Port	The port number of the SYSLOG receiver that accepts incoming events.
Custom Data	The custom data embedded inside the SYSLOG message.
Facility Settings	The facility code defined in the RFC5424 protocol.
Audit Log	The type of audit log.
Audit Log Severity Settings	
Error	The severity level of the error: select from the drop-down menu.
Info	The severity level of the information: select from the drop-down menu.
Debug	The severity level of the debug: select from the drop-down menu.

## Routing Peers Table

*Support > Technical Assistance > Routing Peers Table*

The **Routing Peer Table** page can be used to track the communication between multiple peers within a network and for troubleshooting purposes. This page also reflects the details of the subnet information being shared between each set of peers.

The following table describes the values for the Routing Peers table.

Field Name	Description
Appliance Name	The name of the appliance.
Peer ID	The ID of the peer.
Peer Name	The name of the peer.
Role	Whether the hub or spoke topology is being used for the specified peer.
Last Transmission Count	The last transaction count the peer was sent.
Time since Last Transmission	How many seconds have elapsed since the last subnet update was sent to the peer.
Last Received Count	The last transaction count from the peer that was received.
Time since Last Received	The amount of time since the last received update.
MainVer and Region	The main version and the region of the designated peer.
Message	Peer information to assist in troubleshooting for Silver Peak Support.

## RMA Wizard

Support > [Technical Assistance] RMA

Use this screen as instructed by Support to prepare a Return Merchandise Authorization.

The screenshot shows a window titled "RMA Wizard" with a close button in the top right corner. The main heading is "RMA Discovery and Restore". To the right of the heading are two numbered steps: "1" (highlighted in blue) and "2". Below the heading, there are two columns of fields. The left column is titled "Appliance to Replace" and has a "Type to select" dropdown menu. Below it are labels for "IP", "Model", "Hostname", "Serial Number", and "Software Version". The right column is titled "New Appliance" and also has a "Type to select" dropdown menu. Below it are labels for "IP", "Model", "Hostname", "Serial Number", and "Software Version". At the bottom of the window, there are three buttons: "< Previous", "Next >", and "Apply".

# Built-in Policies

Support > [User Documentation] Build-in Policies

This table displays read-only built-in policies, which are executed before any other policies.

Built-in Policies		7 mins		660 Rows		Search	
Appliance Name	Map	Priority	Match Criteria	Action	Comment		
Mumbai	map1	65500	Source IP any local ip, Destination IP any, Source Port any, Destination Port any, ...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	Next hop monitoring pings, IPSLA...		
Mumbai	map1	65508	Source IP any local ip, Destination IP any, Source Port any, Destination Port any, ...	Force Bypass false, Drop false, Passthrough-unshaped false, Bind To Tunnel true...	ICMPV6 Destination Unreachable ...		
Mumbai	map1	65509	Source IP any local ip, Destination IP any, Source Port any, Destination Port any, ...	Force Bypass false, Drop false, Passthrough-unshaped false, Bind To Tunnel true...	ICMPV6 Time Exceeded error traffic		
Mumbai	map1	65510	Source IP any local ip, Destination IP any, Source Port any, Destination Port any, ...	Force Bypass false, Drop false, Passthrough-unshaped false, Bind To Tunnel true...	ICMP Destination Unreachable err...		
Mumbai	map1	65511	Source IP any local ip, Destination IP any, Source Port any, Destination Port any, ...	Force Bypass false, Drop false, Passthrough-unshaped false, Bind To Tunnel true...	ICMP TTL Expired error traffic		
Mumbai	map1	65512	Source IP any ipv4 address, Destination IP 52.38.28.122 netmask 255.255.255...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	Silver Peak cloud portal HTTPS		
Mumbai	map1	65513	Source IP any ipv4 address, Destination IP 52.38.28.122 netmask 255.255.255...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	Silver Peak cloud portal HTTP		
Mumbai	map1	65514	Source IP any local ip, Destination IP any, Source Port 4500, Destination Port an...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	IPsec NAT traffic		
Mumbai	map1	65515	Source IP any local ip, Destination IP any, Source Port any, Destination Port 450...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	IPsec NAT traffic		
Mumbai	map1	65516	Source IP any local ip, Destination IP any, Source Port 500, Destination Port any...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	IPsec traffic		
Mumbai	map1	65517	Source IP any local ip, Destination IP any, Source Port any, Destination Port 500...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	IPsec traffic		
Mumbai	map1	65518	Source IP any local ip, Destination IP any, Source Port 2048, Destination Port 20...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	WCCP protocol		
Mumbai	map1	65519	Source IP any local ip, Destination IP any, Source Port 4164, Destination Port an...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	UDP flow redirection		
Mumbai	map1	65520	Source IP any local ip, Destination IP any, Source Port any, Destination Port 416...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	UDP flow redirection		
Mumbai	map1	65521	Source IP any local ip, Destination IP any, Source Port 4164, Destination Port an...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	TCP flow redirection		
Mumbai	map1	65522	Source IP any local ip, Destination IP any, Source Port any, Destination Port 416...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	TCP flow redirection		
Mumbai	map1	65523	Source IP any local ip, Destination IP any, Source Port 179, Destination Port any...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	BGP routing protocol		
Mumbai	map1	65524	Source IP any local ip, Destination IP any, Source Port any, Destination Port 179...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	BGP routing protocol		
Mumbai	map1	65525	Source IP any local ip, Destination IP any, Source Port any, Destination Port any, ...	Force Bypass false, Drop false, Passthrough-unshaped true, Bind To Tunnel false...	BGP routing protocol		

## Realtime Charts

*Support > [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.

The screenshot displays the 'Realtime Charts' interface. At the top, there is a tab labeled 'Realtime Charts x'. Below the tab, the title 'Realtime Charts' is followed by a help icon. The interface is divided into two main sections: 'Filter' and 'Metric to plot'.  
 In the 'Filter' section, there are three dropdown menus: 'Appliance' set to 'laine-vxa (10.0.238.71)', 'Type of stats' set to 'Aggregate Tunnel Stats', and 'Traffic Type' set to 'Optimized'. Below these is a 'Dashboards' section with a 'Dashboard' dropdown menu and three buttons: 'Save', 'Save As', and 'Delete'.  
 In the 'Metric to plot' section, the 'Metric' dropdown is set to 'Inbound WAN Rate bps', and the 'Show in local time' checkbox is checked. A 'Plot Chart' button is located to the right of this section.  
 Below the filter and metric sections is a line chart. The chart title is 'Inbound WAN Rate bps, Traffic: Optimized, laine-vxa (10.0.238.71)'. The y-axis represents the rate in bits per second (bps), ranging from 0.00 to 3.0M with major grid lines every 1.0M. The x-axis shows time in HH:MM format, with labels for 15:18, 15:19, 15:20, 15:21, and 15:22. The chart displays a blue line that fluctuates slightly around a mean value of approximately 3.5M bps.

1. Select the filters you want, then click **Plot**.

The chart appears at the bottom of the page.

2. To save as a dashboard, click **Save As**, then enter a name for your dashboard. Don't include spaces in your name. Click **Save**.

If successful, a green Success bar appears and the dashboard name shows up in the

**Dashboard** field.

To retrieve it later, go to this page and choose the dashboard from the drop down list.

## Historical Charts

*Support > [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.

Historical Charts ×

### Historical Charts ?

**Filter**

Appliance:

Type of Stats:

Tunnel:

**Metric to plot**

Metric:

Plot Rate:

Show in local time:

**Dashboards**

Dashboard:

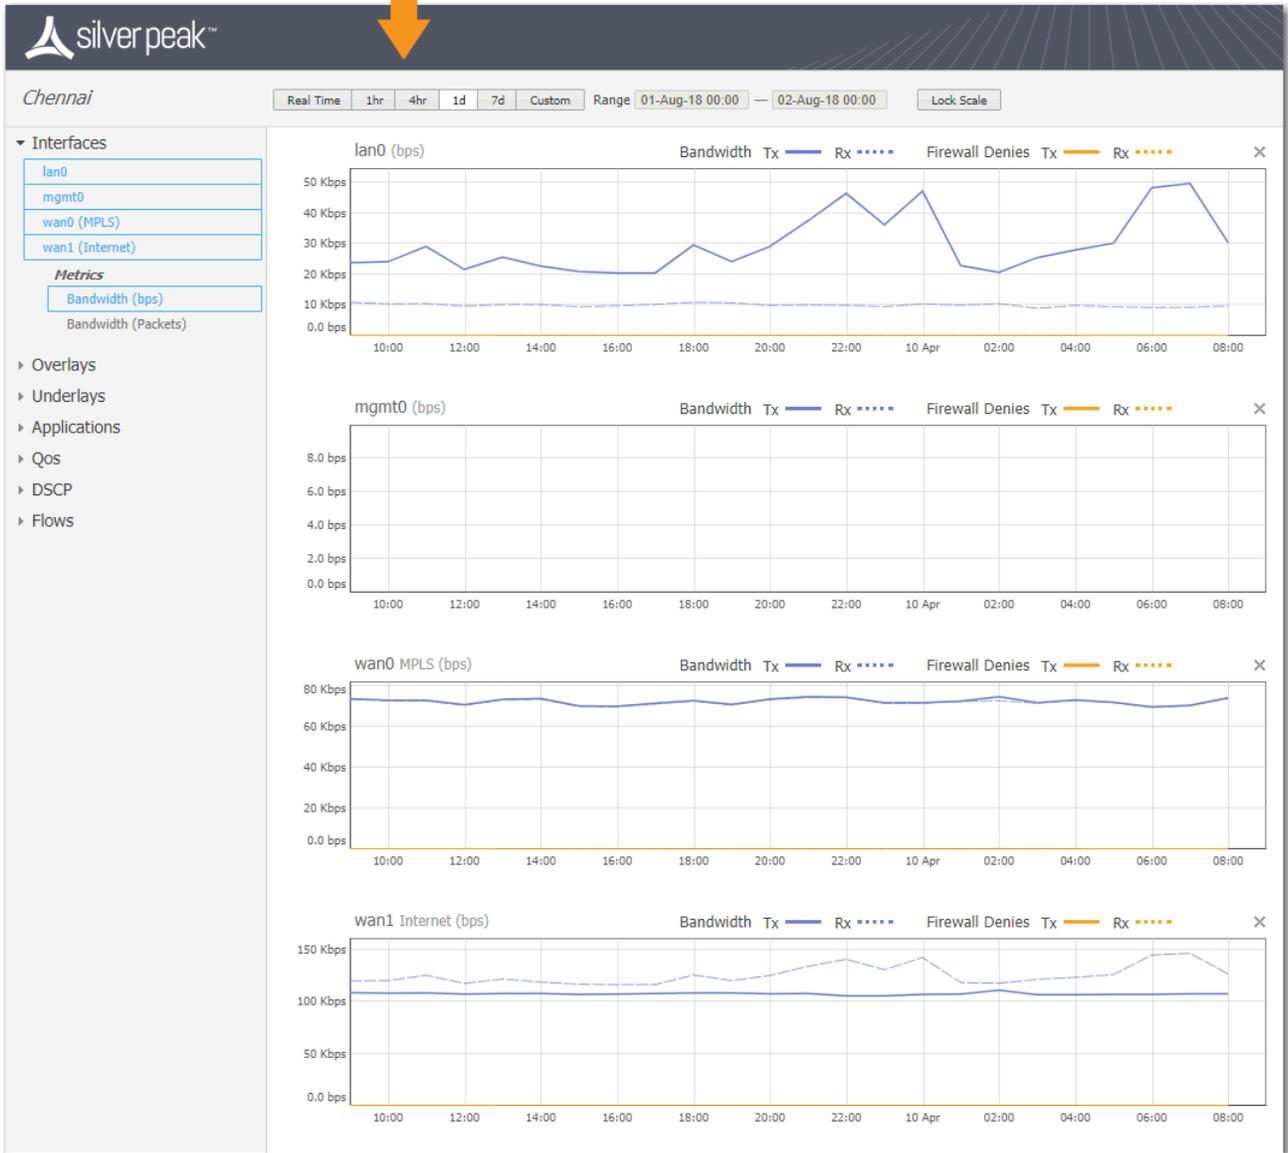
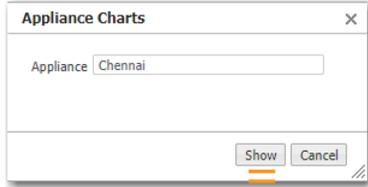
Outbound WAN Bytes, Tunnel: auto\_tun\_10.1.153.20\_to\_10.1.154.20, 1 minute cumulative (laine-vxa (10.0.238.71)) ×

Dec 2016 Jan 2017

## Appliance Charts

*Support > [Reporting] Appliance Charts*

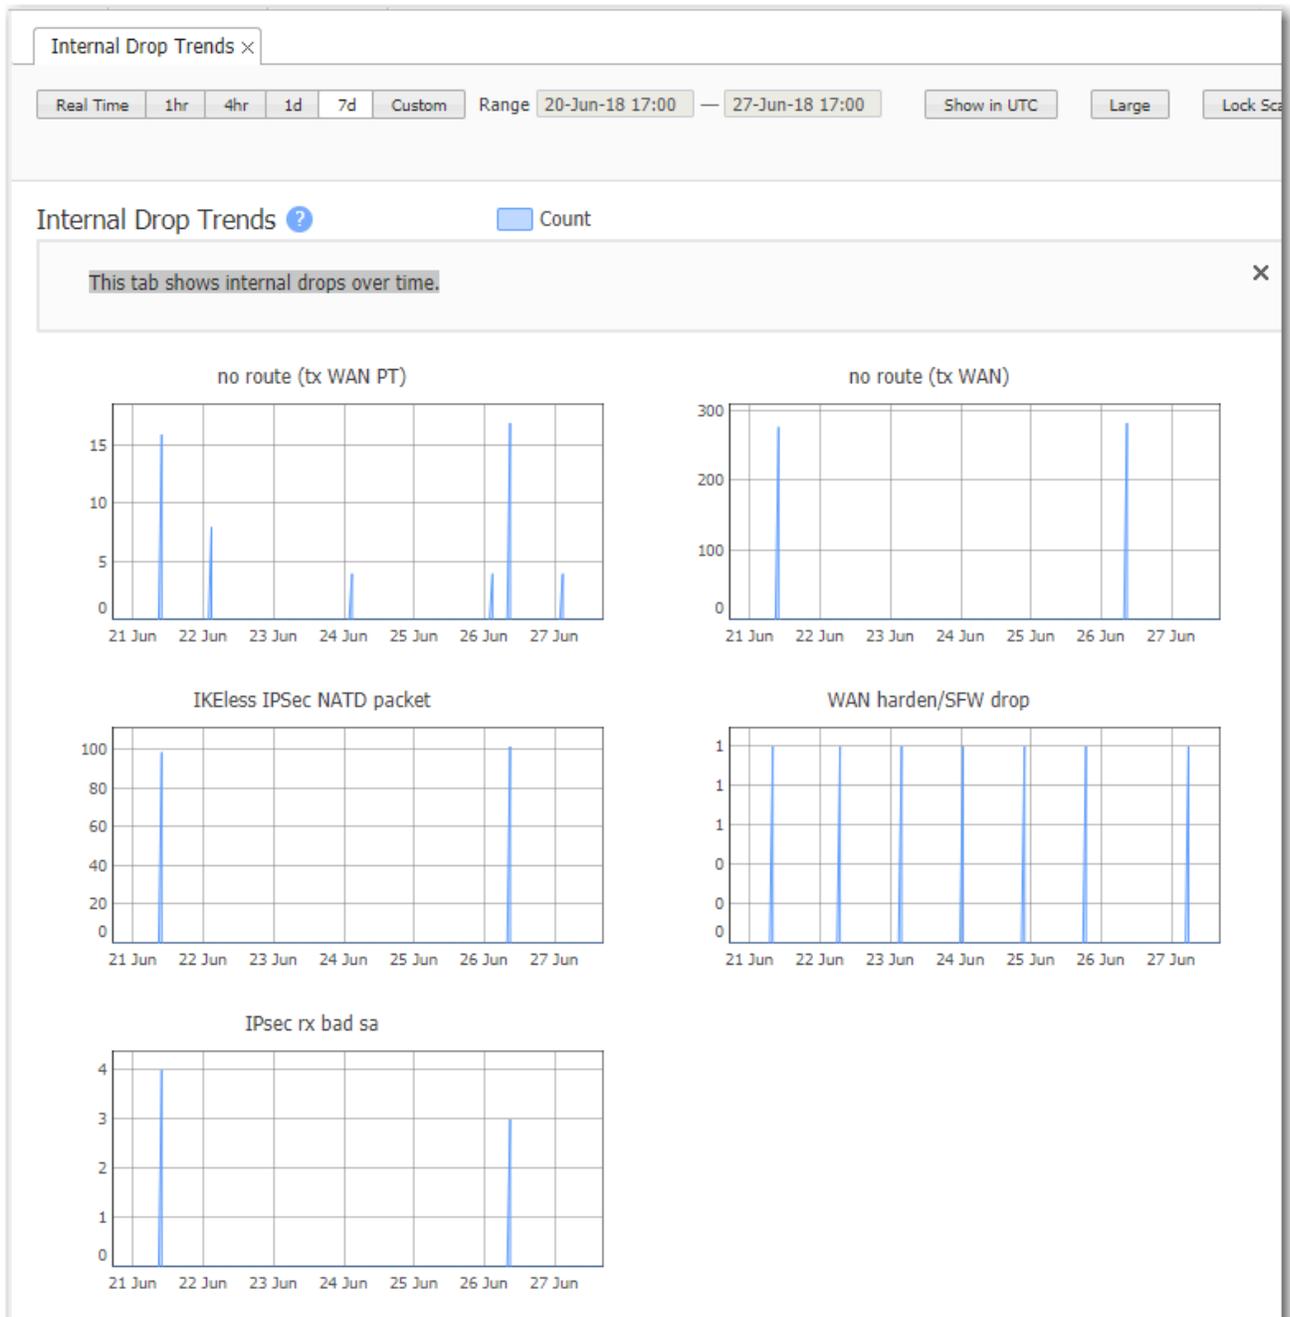
Use this screen to access an individual appliance's realtime and historical charts.



# Internal Drop Trends

Support > [Reporting] Dropped Packet Trends

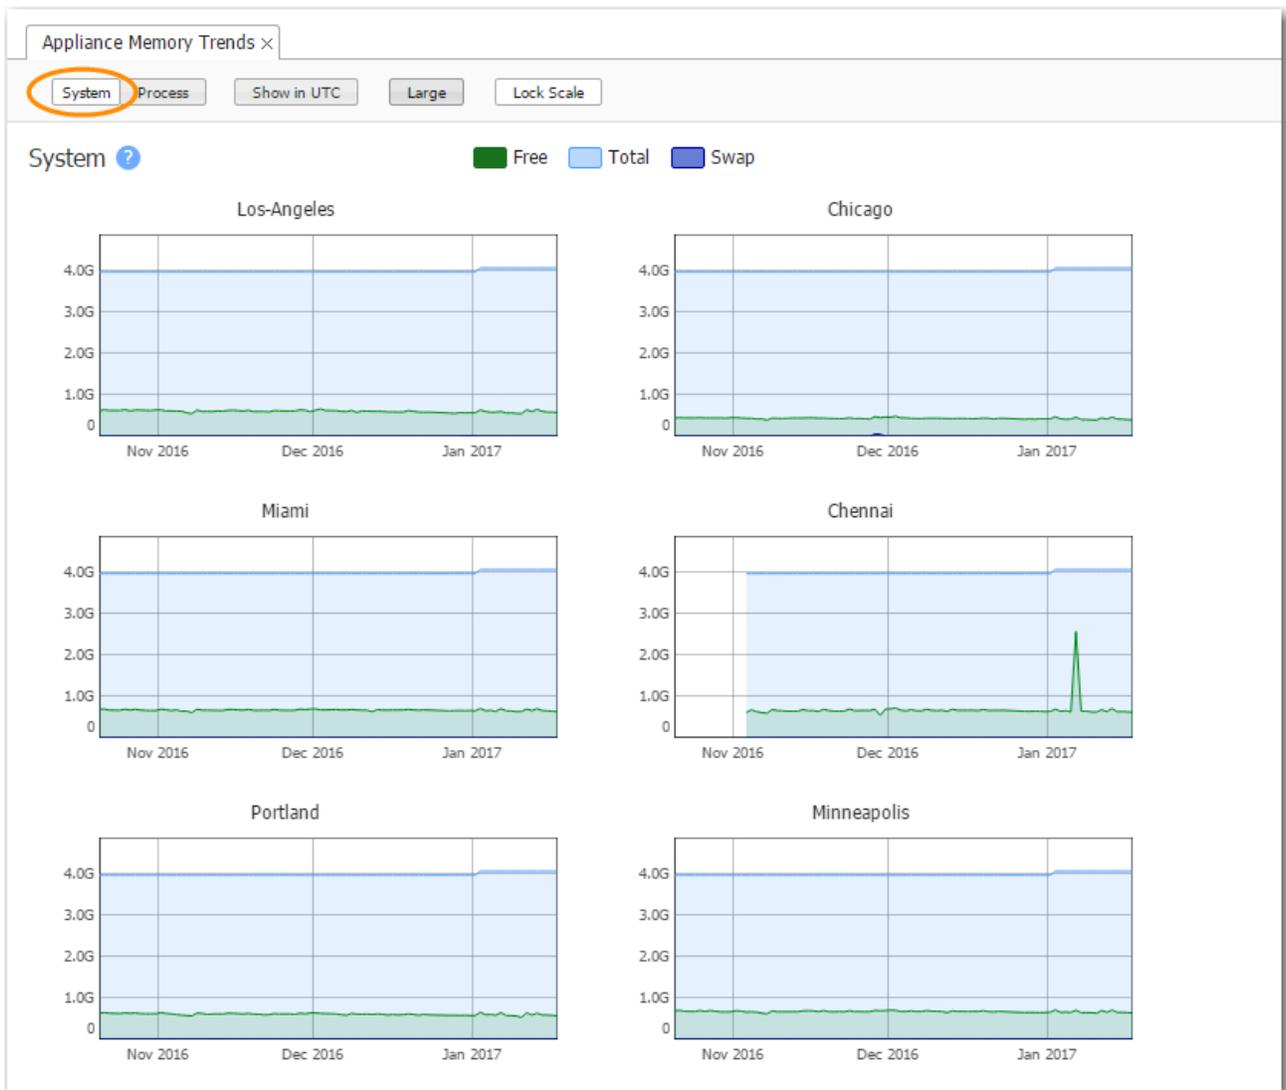
This tab shows internal packet drops over time.



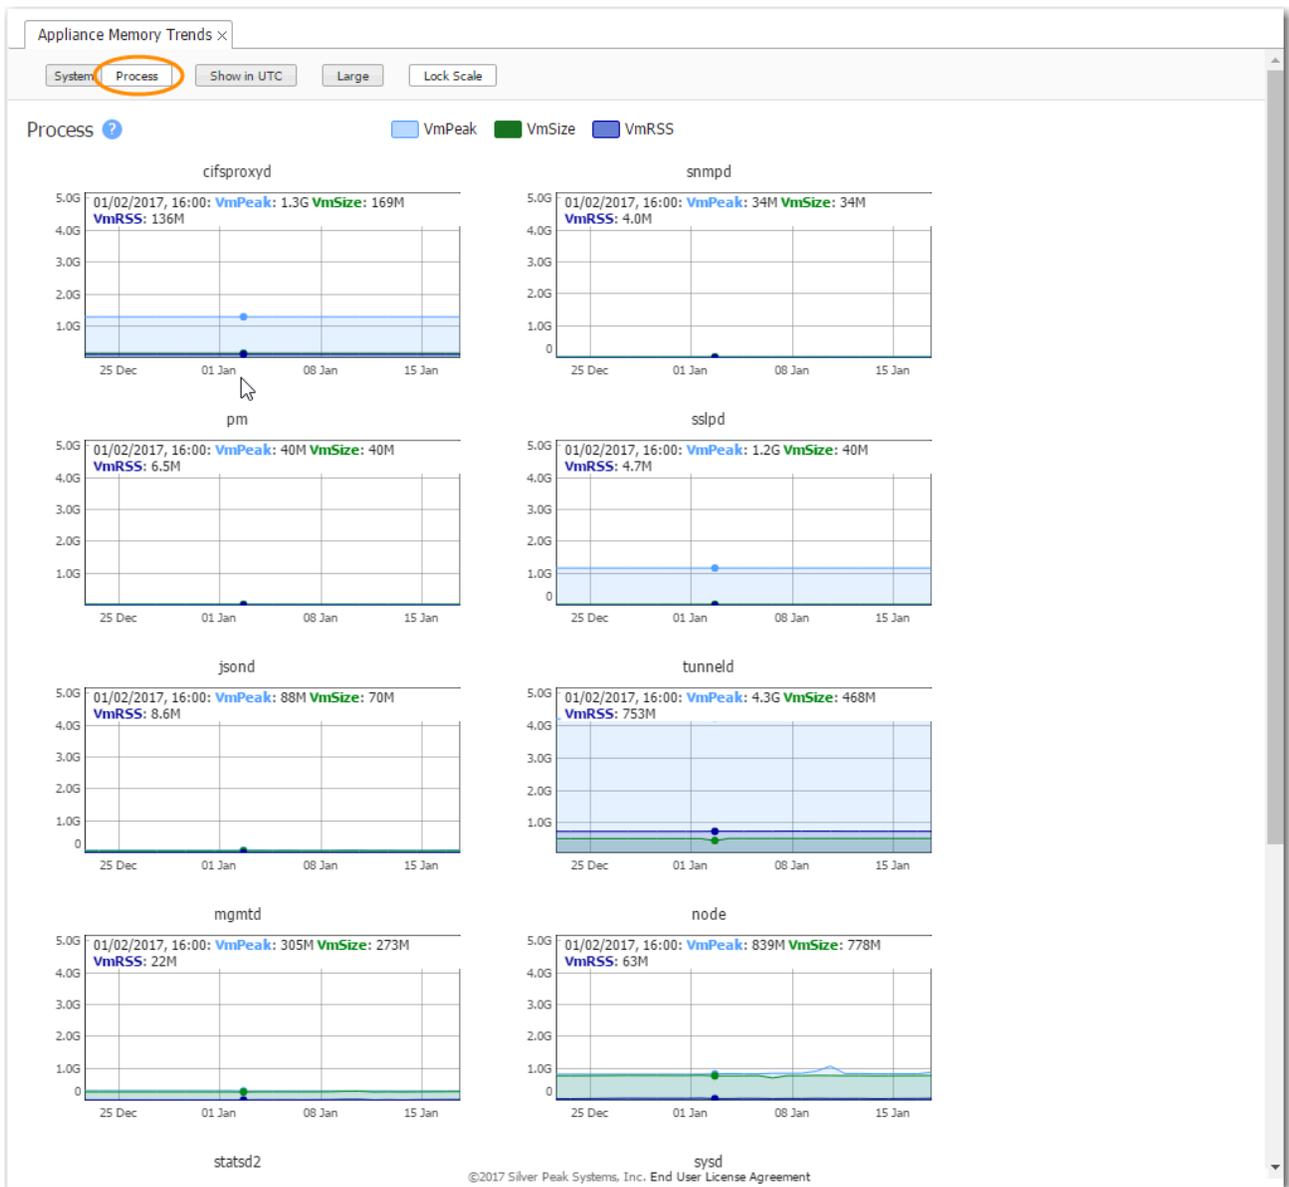
# Appliance Memory Trends

Support > [Reporting] Appliance Memory Trends

The **System** view shows appliance daily memory usage.



The **Process** view is for individual appliances.

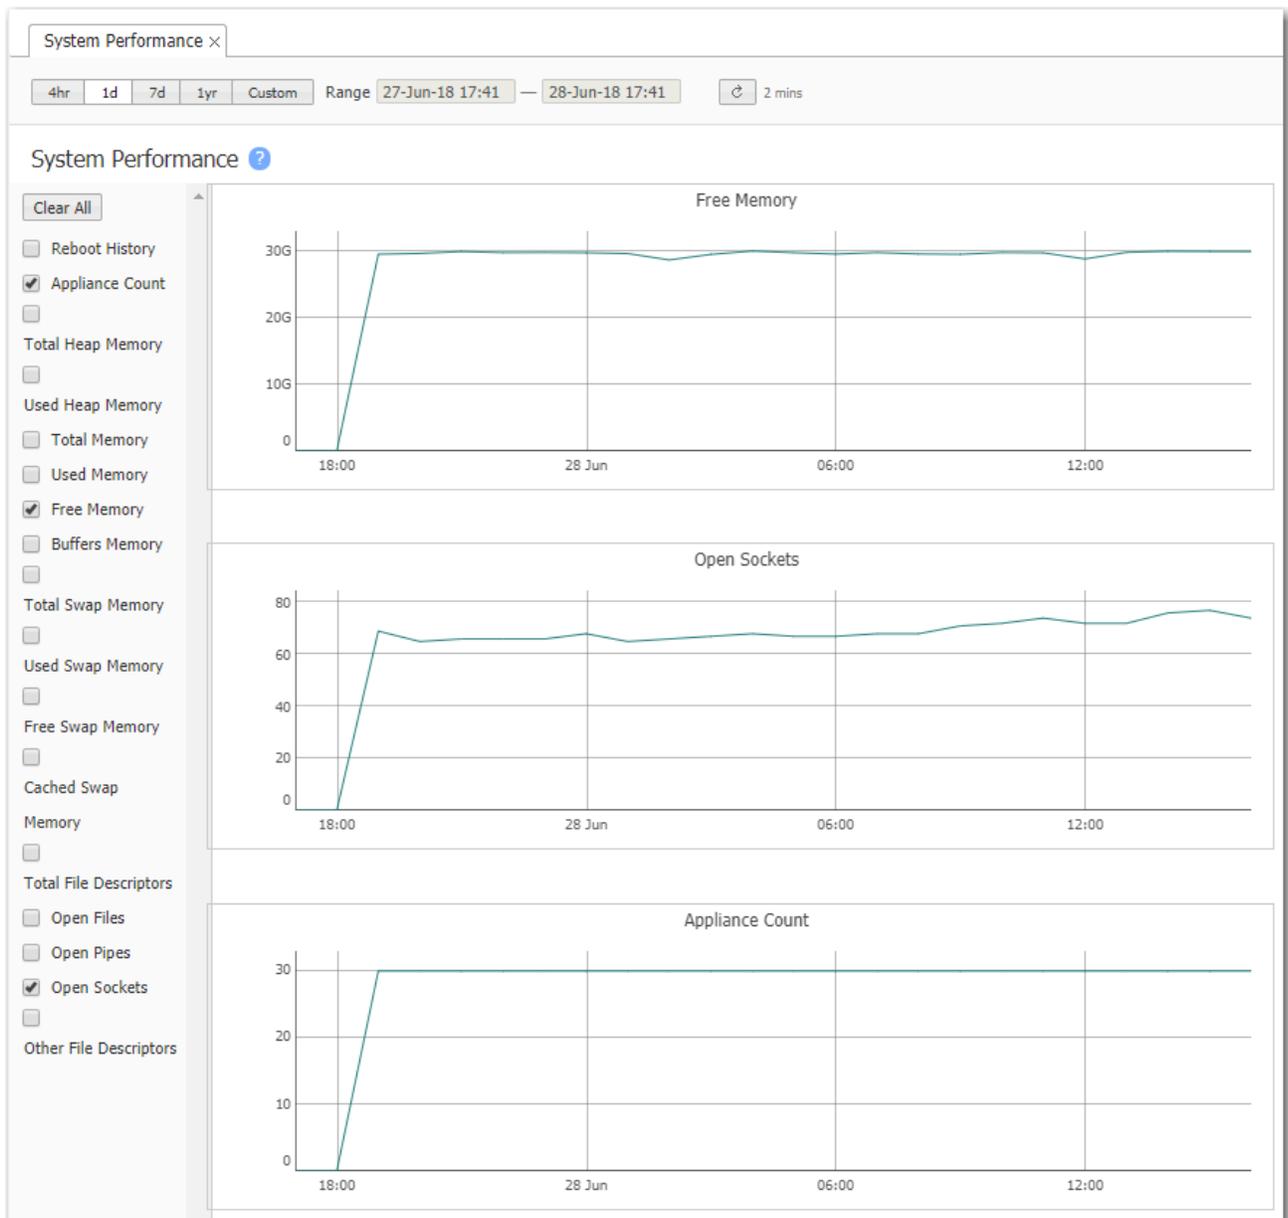


# System Performance

Support > [Reporting] System Performance

This tab shows Orchestrator metrics.

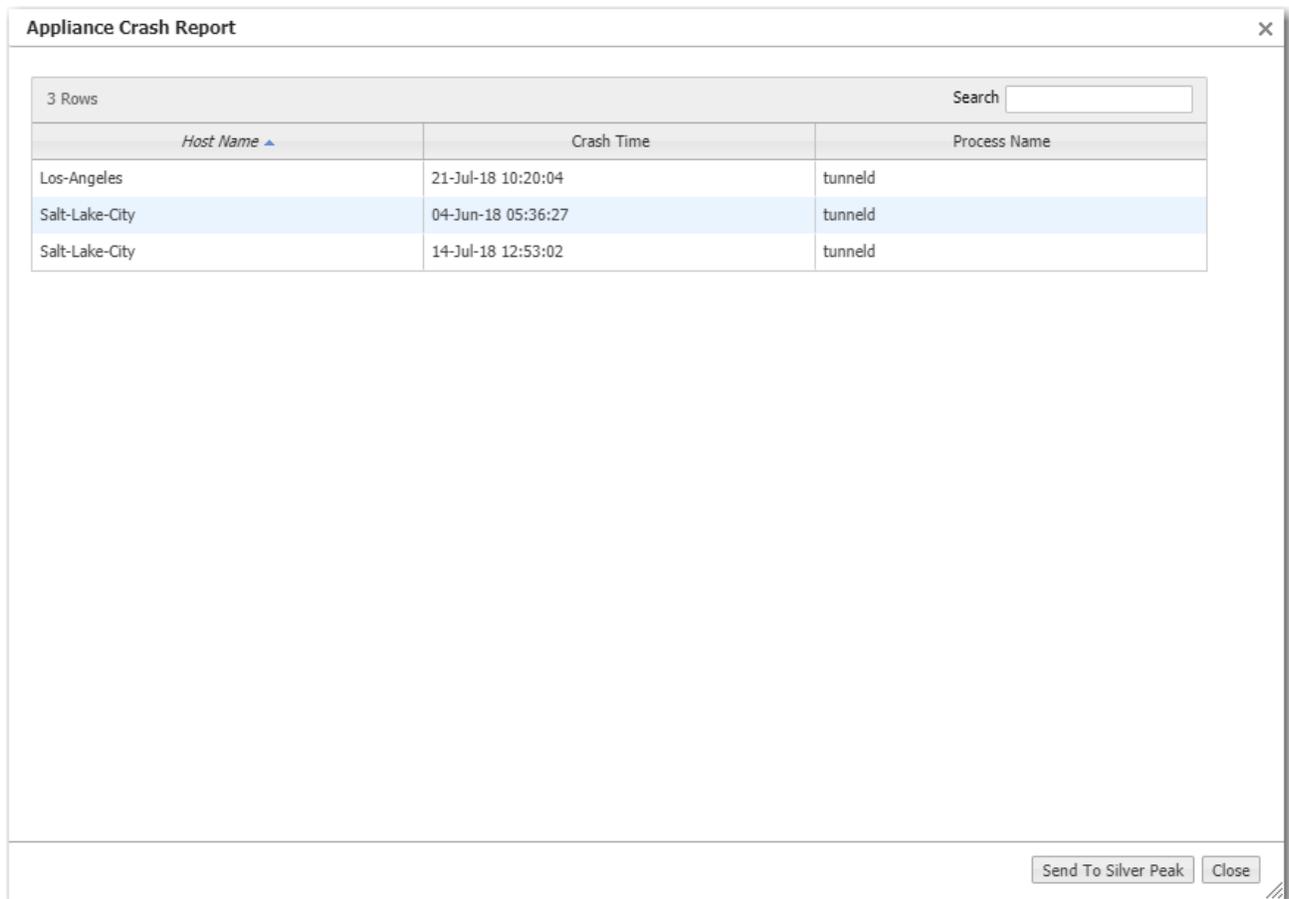
Orchestrators located in the cloud cannot display useful information about host memory, file descriptors, sockets, or pipes.



# Appliance Crash Report

Support > [Reporting] Appliance Crash Report

This report lists appliance crashes, which you can forward to Silver Peak.



The screenshot shows a window titled "Appliance Crash Report" with a close button (X) in the top right corner. The window contains a table with 3 rows and 3 columns. The columns are labeled "Host Name", "Crash Time", and "Process Name". The first row is "Los-Angeles" with a crash time of "21-Jul-18 10:20:04" and process name "tunneld". The second row is "Salt-Lake-City" with a crash time of "04-Jun-18 05:36:27" and process name "tunneld". The third row is "Salt-Lake-City" with a crash time of "14-Jul-18 12:53:02" and process name "tunneld". There is a search bar in the top right of the table area and two buttons, "Send To Silver Peak" and "Close", at the bottom right of the window.

Host Name ▲	Crash Time	Process Name
Los-Angeles	21-Jul-18 10:20:04	tunneld
Salt-Lake-City	04-Jun-18 05:36:27	tunneld
Salt-Lake-City	14-Jul-18 12:53:02	tunneld

# Orchestrator Debug

Support > [Reporting] Orchestrator Debug

This screen contains the various debugging tools available to Support for troubleshooting and debugging issues with Orchestrator.

Statistics Information									
Appliance Info		Appliance Polling		MySQL Tables		MySQL Partitions		Polling Stats	
Reachability Stats		Stack Dump		Quartz Jobs		Websockets		Overlay Cache Stats	
Overlay Manager Stats		Sync Stats		REST Request Time Stats		Orchestration Task		Orchestration Progress	
30 Rows <span style="float: right;">Search <input type="text"/></span>									
ID	Hostname	Site	Mgmt IP	Discovered Fro...	Software Versi...	UUID	Portal Object ID	Dynamic UUID	
0.NE	Chennai		10.0.185.29	PORTAL	0.0.0.0_71872	da6c7c1e-489d-47ae-b5dd-8f3f4deb962b	59e632fc26fcf525dea4dbd0	e0672c9a-a8ef-4d01-8c71-0235d7c1a158	
1.NE	San-Jose		10.0.185.47	PORTAL	0.0.0.0_71872	c3aa83d6-8a73-4ca9-a6c4-e7c98779f8c7	59e6344526fcf525dea4dbf4	c212c6c4-e02c-40f3-ac1f-cfa52bfd5d3f	
10.NE	Minneapolis		10.0.185.28	PORTAL	0.0.0.0_71872	59e88c6e-681a-4c49-a874-5ae31f7e50af	59e632e526fcf525dea4dbcc	f0b9043a-fd29-4a88-9520-8e23e46a3a9a	
11.NE	Mumbai		10.0.185.44	PORTAL	0.0.0.0_71872	282c3a7a-c364-4524-8542-b244882b96e3	59e6343826fcf525dea4dbf2	6ed5bdc3-0fc3-4876-a955-91794d28ca14	
12.NE	Los-Angeles		10.0.185.25	PORTAL	0.0.0.0_71872	3082126c-fd60-49da-aba4-25b7e540253e	59e6326f26fcf525dea4dbcc2	70540b55-903e-41e6-a207-6e130e886ed5	
13.NE	Portland		10.0.185.26	PORTAL	0.0.0.0_71872	6c83b38e-0218-4cfc-aa0a-ed93da713fa8	59e632b726fcf525dea4dbcc6	6fac4146-7cce-4fac-b5a8-3d662039abf1	
14.NE	Salt-Lake-City		10.0.185.27	PORTAL	0.0.0.0_71872	e6f6eff7-795c-4c9e-829a-ac1001c730fb	59e632af26fcf525dea4dbcc4	8c843b37-a6a7-4537-af4a-e690de617f4a	
15.NE	Singapore		10.0.185.30	PORTAL	0.0.0.0_71872	542eefcf-4117-4dd7-bf50-77cabeb3837e	59e632db26fcf525dea4dbca	669c74ac-8b1b-4b41-bcda-8e979d2880ac	
16.NE	Milan		10.0.185.31	PORTAL	0.0.0.0_71872	6db04d99-0a5f-467d-b2fe-7d19ef6d096e	59e632f826fcf525dea4dbce	d4b6146e-5997-4f80-9e4f-e14138b8ab69	
17.NE	Paris		10.0.185.33	PORTAL	0.0.0.0_71872	97a2b896-5e60-4403-a35a-205cae11d18b	59e6331526fcf525dea4dbd5	147a7381-c133-4a5a-8038-ac6418db1c1f	
18.NE	Tokyo		10.0.185.34	PORTAL	0.0.0.0_71872	52901e40-41a6-43cd-84f2-877c69f298f	59e6332d26fcf525dea4dbd7	3f317cf1-60e9-41bf-84e1-52514312133f	
19.NE	London		10.0.185.37	PORTAL	0.0.0.0_71872	fd25bb96-8d0b-479f-8f8a-c0ce36ce632d	59e6337c26fcf525dea4dbde	cb8d724e-3480-4a9c-b836-e54eac883a10	
2.NE	Geneva		10.0.185.45	PORTAL	0.0.0.0_71872	ec3ce87e-c2e6-4edf-bcfa-c186e59805fc	59e6344f26fcf525dea4dbf6	3b734ddf-b1f6-4805-bccb-b21848263f35	
20.NE	Frankfurt		10.0.185.38	PORTAL	0.0.0.0_71872	40e4333f-f872-4ae3-9ef2-c25ed38c2f59	59e6339326fcf525dea4dbe0	1f6c1c70-f235-497a-85fc-14d68e17f236	
21.NE	Barcelona		10.0.185.40	PORTAL	0.0.0.0_71872	1988e63f-a999-4e58-a883-3179989d1d0	59e633bd26fcf525dea4dbe8	58d0cff8-9ee5-484f-b082-7392eda8541d	
22.NE	Seoul		10.0.185.41	PORTAL	0.0.0.0_71872	6034357a-3014-4458-bf1d-96792a1c7c35	59e633d726fcf525dea4dbea	4f1b3d8a-f195-4285-82d9-adf8fe84f39f	
23.NE	Osaka		10.0.185.42	PORTAL	0.0.0.0_71872	a5f28a69-1663-438a-aa23-f2d6615eb332	59e633ee26fcf525dea4dbec	e69168f7-683d-438d-a3ec-4dc56a5e6b46	
24.NE	San-Antonio		10.0.185.46	PORTAL	0.0.0.0_71872	752eafa6-c34f-43e4-8c3d-00270a5ad2c	59e6346726fcf525dea4dbfa	318091a7-fbd1-4b5c-ba2d-2dd8c9141e1b	
25.NE	New-Orleans		10.0.185.48	PORTAL	0.0.0.0_71872	0a6c65b3-e013-4473-bb7-9e0d5d2b9fb7	59e6345c26fcf525dea4dbf8	3d9650b9-e9ee-4e20-b864-2214d8d9fd2c	
26.NE	Toronto		10.0.185.52	PORTAL	0.0.0.0_71872	237a7b0d-9173-4c3b-a3bc-6871f5c3881d	59e634d826fcf525dea4dc04	00ba0f0e-1735-44a2-944f-9aaf25df753f	
27.NE	New-York		10.0.185.53	PORTAL	0.0.0.0_71872	eb01d02f-14ab-430f-89bb-f6b0355498c4	59e634ee26fcf525dea4dc08	be23974e-8882-4d54-a15a-a129269a725a	
28.NE	Pittsburgh		10.0.185.49	PORTAL	0.0.0.0_71872	9bb8440b-5964-4ffd-b388-430cd02c7160	59e6347a26fcf525dea4dbfc	62df50d8-a56b-42c8-b302-7a680d9e4bae	
29.NE	Albuquerque		10.0.185.50	PORTAL	0.0.0.0_71872	f8e9a141-8b09-4585-a84f-1a574e31ae89	59e634a626fcf525dea4dbfe	b8d36639-2b1d-458b-a525-5d34df9ad225	
30.NE	Edinburgh		10.0.185.54	PORTAL	0.0.0.0_71872	13c4f99a-db70-dfb8-9a0b-468aa6fa0a06	59e634a826fcf525dea4dbf6	55916216-5f8f-4c11-43f0-97f4817a7c43	

## IPSec UDP Status

Support > [Reporting] IPSec UDP Status

This tab displays IPSec UDP key and activation status.

IPSec UDP Status x

IPSec UDP Status ? ↻

Seed Activation Time 02-Aug-18 10:20  
Previous Seed Activation Time 01-Aug-18 10:20

30 Rows Search

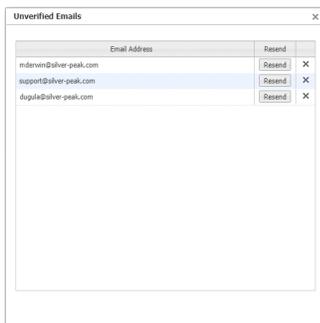
Host Name	Has Current Key	Activation Status
San-Jose	true	true
Edinburgh	true	true
Pittsburgh	true	true
Dallas	true	true
Portland	true	true
Singapore	true	true
Mumbai	true	true
Barcelona	true	true
Tokyo	true	true
Mexico-City	true	true
Boston	true	true
New-York	true	true
New-Orleans	true	true

## Unverified Emails

*Support > [Reporting] Unverified Emails*

When you add an email address to either the Alarms or the Reports distribution, Orchestrator sends the recipient an email containing a link, asking them to click to provide verification.

If Orchestrator doesn't receive a verification, then either the recipient hasn't responded or the email address is invalid.



Email Address	Resend	
mldevin@silver-peak.com	Resend	X
support@silver-peak.com	Resend	X
dupla@silver-peak.com	Resend	X

- An unverified email address remains inactive and doesn't generate an alarm.
- You can retest an address with **Resend**.
- You can only correct an email address in the Alarm or Reports email distribution list.