



Silver Peak Unity Orchestrator User Guide

Orchestrator 9.1.0
Last updated on September 24, 2021
200095-003
Revision B

Copyright and Trademarks

Copyright © 2021 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 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

We are dedicated to continually improving our products and documentation. If you have suggestions or feedback for our documentation, send an e-mail to techpubs@silver-peak.com.

Contents

What's New	14
Orchestrator 9.1.0	14
Intrusion Detection System (IDS)	14
Aruba ClearPass Policy Manager Integration	14
Remote Statistics Collection	14
Support for ACL Group Objects	14
Support for Non-routing Hub (Stub Hub)	14
OSPF Template	14
Separation of Active and Historical Alarms	15
Zone Name in Routes Dialog	15
Orchestrator 9.0.5	15
Improvements in Denied Devices List	15
One-click Cloud EC-V	15
Third-Party Service Orchestration	15
Getting Started	16
Supported Browsers	17
Guidelines for Creating Passwords	18
Overview of SD-WAN Prerequisites	19
Monitor Status and Performance	21
Dashboard	21
Topology Settings and Legend	21
View Tunnels in the Topology Map	24
Health Map	27
Alarms Tab	29
Disable Alarms	30
Customize Alarms	30
Alarm Severity	30
Alarm Recipients	31
Additional Alarm Indications	31
Schedule and Run Reports	31
View Reports	33
Sample Report	34
Scheduled and Historical Jobs	34
Overlay-Interface-Transport	35
Interface Bandwidth Trends	36
Interface Summary	38
Application Bandwidth	38
Application Pie Charts	39
Application Trends	40
Top Talkers	41
Domains	43
Countries	43
Ports	44
Traffic Behavior	45

Appliance Bandwidth	46
Appliance Max Bandwidth	47
Appliance Bandwidth Utilization	48
Appliance Bandwidth Trends	48
Appliance Packet Counts	48
Tunnels Bandwidth	49
Show Underlays	49
Traceroute	50
Live View	50
Tunnels Pie Charts	51
Tunnel Bandwidth Trends	52
Tunnel Packet Counts	53
DRC Bandwidth Trends	54
Dynamic Rate Control	55
Flows - Active and Recent	55
Reset or Reclassify Flows	57
Additional Information about Flows	58
Appliance Flow Counts	59
Appliance Flow Trends	60
Tunnel Flow Counts	60
DSCP Bandwidth	61
DSCP Pie Charts	61
DSCP Trends	61
Traffic Class Bandwidth	62
Traffic Class Pie Charts	62
QoS (Shaper) Trends	63
Shaper Summary	64
Boost Tab	65
Boost Trends	66
Change Boost Configuration	66
Firewall Drops	67
Live View	68
Loss	68
Loss Trends	69
Jitter Summary	70
Jitter Trends	71
Latency	72
Latency Trends	73
Out of Order Packets	74
Mean Opinion Score (MOS) Summary	77
Mean Opinion Score (MOS) Trends	77
Tunnels Summary	78
Orchestrator Configuration	80
Unity Overlays	81
Business Intent Overlays	81
Overview	81
SD-WAN Traffic to Internal Subnets	81
Breakout Traffic to Internet and Cloud Services	84
Apply Overlays	85

Interface Labels	85
Manage Labels	86
Hubs	87
Deployment Profiles	88
Map Labels to Interfaces	89
LAN-side Configuration: DHCP	89
WAN-side Configuration	89
Descriptions	91
A More Comprehensive Guide to Basic Deployments	92
Bridge Mode	92
Router Mode	93
Server Mode	97
Deployment - EdgeConnect HA	98
Enable EdgeConnect HA Mode	98
IPSec over UDP Tunnel Configuration	99
VRRP Configuration	99
LAN-side Monitoring	99
Firewall Zones	99
Internet Traffic	100
IPSec Pre-Shared Key Rotation	101
Failure Handling and Orchestrator Reachability	101
Schedule IPSec Key Rotation Dialog Box	102
Intrusion Detection System (IDS)	102
Prerequisites	103
Enable or Disable IDS on Appliances	103
Enable or Disable Rules with the IDS Allow List	104
Specify Traffic to Be Inspected	105
Advanced Reporting and Analytics	106
SSL Certificates Tab	107
SSL CA Certificates Tab	108
SSL for SaaS Tab	109
Discovered Appliances	111
Preconfigure Appliances	112
Appliance Configuration Wizard	113
Licenses	116
Assign a License to an Appliance	117
Cloud Portal	117
Network Configuration Tabs	119
Deployment Tab	119
Interfaces Tab	120
Terminology	122
NAT	122
NAT Rules and Pools	123
VRRP Tab	124
WCCP Tab	124
PPPoE Tab	125
Loopback Interfaces	128
Loopback Orchestration	128
Virtual Tunnel Interface	129
VTI	129

DHCP Server Defaults	130
DHCP Settings	131
DHCP Leases	132
DHCP Failover	132
DHCP Failover State	133
Link Aggregation	134
View Aggregation Details	134
Modify Link Aggregation	135
Regions	136
Routing Segmentation	138
Management Services	143
Management Services Dialog Box	144
Inter-segment DNAT Exceptions	144
Inter-segment SNAT Exceptions	145
BGP Tab	145
BGP Information	147
Routes Tab	151
Route Maps	151
Edit or Add Routes	155
Import Subnets	156
OSPF Tab	157
OSPF Edit Row	157
Add Interface	157
OSPF Route Redistribution Maps	158
Multicast	160
Multicast Dialog Box	160
Peer Priority Tab	162
Peer Priority Edit Row	162
Admin Distance Tab	163
Management Routes Tab	164
Tunnels Tab	164
Troubleshooting	167
Use Passthrough Tunnels	167
Tunnel Exception	167
Schedule Auto MTU Discovery	168
Policy Configuration Tabs	169
DNS Proxy Policies	169
Configure DNS Proxy Policies	169
Route Policies Tab	170
Priority	171
Match Criteria	171
Source or Destination	171
Wildcard-based Prefix Matching	171
QoS Policies Tab	172
Handle and Mark DSCP Packets	173
Priority	177
Match Criteria	177
Source or Destination	177
Wildcard-based Prefix Matching	177
Schedule QoS Map Activation	178

Optimization Policies Tab	178
Priority	179
Match Criteria	179
Source or Destination	180
Wildcard-based Prefix Matching	180
Set Actions	181
TCP Acceleration Options	182
NAT Policies Tab	185
Advanced Settings	187
Inbound Port Forwarding	189
Security Policies Tab	190
Wildcard-based Prefix Matching	191
Access Lists Tab	192
Match Criteria	193
Wildcard-based Prefix Matching	193
Address Groups	193
Add an Address Group	194
Add a Rule to an Address Group	195
Delete an Address Group	196
Export Address Groups	196
Import Address Groups	197
View a Single Address Group	198
Edit or Delete a Rule	198
Using Address Groups in Match Criteria	198
Address Group Formats	199
Service Groups	199
Add a Service Group	200
Add a Rule to a Service Group	202
Delete a Service Group	203
Export Service Groups	203
Import Service Groups	204
View a Single Service Group	205
Edit or Delete a Rule	205
Using Service Groups in Match Criteria	206
Shaper Tab	206
SaaS Optimization Tab	209
Configuration Tab	209
Application Definitions	209
Application Groups Tab	211
Threshold Crossing Alerts Tab	211
IP SLA Tab	214
IP SLA Monitor Use Cases	214
Configuration Templates	222
Templates Overview	222
Template Groups	222
System Template	222
SNMP Template	226
Flow Export Template	228
DNS Template	228
Logging Template	229

Minimum Severity Levels	229
Configure Remote Logging	230
Banner Messages Template	230
Date/Time Setting	231
Data Collection	232
HTTPS Certificate Template	232
User Management Template	233
Default User Accounts	234
Command Line Interface Privileges	234
SSL Certificates Template	234
SSL CA Certificates Template	235
SSL for SaaS Template	236
Auth/Radius/TACACS+ Template	238
Authentication and Authorization	238
Appliance-based User Database	238
RADIUS	238
TACACS+	239
What Silver Peak Recommends	239
DNS Proxy Policies	239
Tunnels Template	240
VRRP Template	242
Peer Priority Template	244
Admin Distance Template	245
Route Redistribution Template	246
Shaper Template	247
Dynamic Rate Control	247
QoS Policies Template	249
Priority	250
Match Criteria	250
Source or Destination	251
Wildcard-based Prefix Matching	251
Handle and Mark DSCP Packets	251
Routes Template	255
BGP Template	256
OSPF Template	257
Optimization Policies Template	258
Priority	259
Match Criteria	259
Source or Destination	259
Wildcard-based Prefix Matching	259
Set Actions Fields	260
Route Policies Template	261
Why?	262
Priority	262
Match Criteria	262
Source or Destination	263
Wildcard-based Prefix Matching	263
Set Actions Fields	263
NAT Policies Template	264
When to NAT	264

Advanced Settings	266
Threshold Crossing Alerts Template	268
TCA Metrics	270
SaaS Optimization Template	271
TIPS	272
Security Policies Template	273
Wildcard-based Prefix Matching	274
CLI Template	274
Session Management Template	275
Management Services Template	277
Apply Template Groups	277
Cloud Services	278
AWS Transit Gateway Network Manager	278
Orchestrator Configuration	281
Check Point CloudGuard Connect	284
Import and Export Subnets	286
Microsoft Azure Virtual WAN	286
Azure Prerequisites	287
Orchestrator Prerequisites	287
Orchestrator Configuration	288
Verification	289
Works with Office 365	289
Zscaler Internet Access	290
Enable Zscaler	293
Verification	293
Service Orchestration	294
Prerequisites	295
Remote Endpoint Configuration	295
Bulk Edits	297
Interface Labels	297
Tunnel Settings	297
IP SLA Settings	297
Pause Orchestration (Optional)	298
+BIO Breakout	298
Remote Endpoint Association	299
Add Tunnel Local Identifiers to Netskope	299
Verification	299
Set Up a New Service	299
Deploy EC-V in Cloud	300
Deploy EC-V in Cloud Tab	301
Cloud Deployment Accounts	302
AWS Account Configuration	303
EC-V Deployment Configuration	304
Appliance Administration	306
Appliance User Accounts Tab	306
Auth/RADIUS/TACACS+ Tab	307
Authentication and Authorization	307
RADIUS and TACACS+	308
Date/Time Tab	308

DNS (Domain Name Servers) Tab	309
SNMP Tab	310
SNMP Overview	310
Modify SNMP Configuration	310
Flow Export Tab	312
Silver Peak Custom Information Elements	312
Logging Tab	316
Severity Levels	317
Remote Logging	317
Banners Tab	317
HTTPS Certificate Tab	318
Orchestrator Reachability Tab	319
Custom Appliance Tags	320
System Information	320
Software Versions	326
Upgrade Appliance Software	326
Appliance Configuration Backup	327
View Configuration History	329
Restore a Backup to an Appliance	330
Remove Appliance from Orchestrator	330
Remove Appliance from Orchestrator and Account	331
Synchronize Appliance Configuration	332
Put the Appliance in System Bypass Mode	333
Broadcast CLI Commands	334
Link Integrity Test	334
TCPPERF Version 1.4.8	335
Disk Management	339
Erase Network Memory	340
Reboot or Shut Down an Appliance	341
Behavior During Reboot	342
Schedule an Appliance Reboot	342
Behavior During Reboot	343
Active Sessions Tab	345
Orchestrator Administration	347
Role Based Access Control	347
Roles	347
Appliance Access	349
Assign Roles and Appliance Access	350
View Orchestrator Server Information	350
Restart, Reboot, or Shutdown	351
Manage Orchestrator Users	351
Add a User	352
Multi-Factor Authentication	352
Modify User	354
API Key	355
Remote Authentication	356
Configure a RADIUS or TACACS+ Server	358
Configure an OAuth Server	358
Configure a JWT Server	361

Configure a SAML Server	363
Cloud Portal	365
Audit Logs	366
Orchestration Settings	367
Tunnel Settings Tab	368
Orchestrator Blueprint Export	372
Brand Customization	373
Maintenance Mode	373
Upgrade Orchestrator Software	374
Check for Orchestrator and Appliance Software Updates	375
Back Up on Demand	376
Schedule Orchestrator Backup	377
Schedule Stats Collector Backup	378
SMTP Server Settings	380
Proxy Configuration	380
Orchestrator's HTTPS Certificate	381
Timezone for Scheduled Jobs	382
Orchestrator Advanced Properties	382
Change the Orchestrator Log Level	383
Minimum Severity Levels	384
IP Allow List	384
Orchestrator's Getting Started Wizard	385
Statistics Retention	387
Stats Collector Configuration	387
Prerequisites	388
Before You Begin	388
Configure the New Stats Collector Feature	389
Notification Banner	392
ClearPass Policy Manager	393
Manage ClearPass Policy Manager Accounts	394
Pause ClearPass Policy Manager Integration	395
Customer and Technical Support	396
Tech Support - Appliances	396
Tech Support - Orchestrator	396
Take Action with Files	397
Log In to the Support Portal	397
Monitor Transfer Progress	398
Packet Capture	399
Upload Local Files	399
Create a Support Case	400
Remote Access	401
Partition Management	402
Remote Log Receivers	402
Routing Peers Table	405
RMA Wizard	406
Run the RMA Wizard	407
Add a Backup Appliance	408
Upgrade and Downgrade	408
Built-in Policies	408

Realtime Charts	409
Historical Charts	410
Appliance Charts	411
Internal Drop Trends	412
Appliance Memory Trends	414
System Performance	415
Appliance CPU Usage	416
Appliance Crash Report	417
Orchestrator Debug	418
IPSec UDP Status	419
Unverified Emails	419

What's New

This page provides a brief description and links to additional information about new features in recent Orchestrator releases.

Orchestrator 9.1.0

The following features were introduced in Orchestrator 9.1.0:

Intrusion Detection System (IDS)

The Intrusion Detection System (IDS) that can monitor traffic for potential threats and malicious activity and generate threat events based on preconfigured rules. Packets are copied and inspected against signatures downloaded to Orchestrator from Cloud Portal. Traffic is designated for inspection using matching rules enabled in the zone-based firewall. For more information, see [Intrusion Detection System \(IDS\)](#).

Aruba ClearPass Policy Manager Integration

Orchestrator now supports association with ClearPass Policy Manager, which provides role-based and secure network access for devices. This integration provides user and role information for an IP address, which you can view on the Flows and Top Talkers tabs of Orchestrator. For more information, see [ClearPass Policy Manager](#).

Remote Statistics Collection

To improve Orchestrator performance, a new remote stats collector feature eliminates the use of Orchestrator resources for monitoring appliances. This new architecture allows you to scale your network with greater performance. For more information, see [Stats Collector Configuration](#).

Support for ACL Group Objects

This release includes two new features related to ACLs: Address Groups and Service Groups. An address group is a logical collection of IP hosts or subnets, and a service group is a logical collection of protocols and ports. Both can be referenced in source or destination matching criteria in the zone-based firewall and security policies. For more information, see [Address Groups](#) and [Service Groups](#).

Support for Non-routing Hub (Stub Hub)

This release adds support for designating a non-routing hub or stub hub by configuring it to not re-advertise spoke-learned routes to other hubs in the region. For more information, see [Hubs](#).

OSPF Template

This release of Orchestrator includes a new template for configuring OSPF. For more information, see [OSPF Template](#).

Separation of Active and Historical Alarms

This release separates active and historical alarms into different database tables. This update will help to address potential deadlock issues in the alarms database and provides support for displaying alarms in the user's own time zone. On the Alarms tab in Orchestrator, you can toggle the alarm view between Active, History, and All.

Zone Name in Routes Dialog

In this release, users can configure a firewall zone when configuring a user-defined route.

Orchestrator 9.0.5

The following features were introduced in Orchestrator 9.0.5:

Improvements in Denied Devices List

Virtual appliances are no longer displayed on the Denied Devices list, and users now have the option to permanently delete one or more appliances from the list and from Orchestrator.

One-click Cloud EC-V

This feature enables users to quickly deploy one or more EdgeConnect Virtual (EC-V) appliances in supported public cloud providers. In this release, Silver Peak supports one-click EC-Vs in AWS. After creating an AWS Identity and Access Management (IAM) account with required permissions for Orchestrator and an EC2 key pair, users can quickly deploy one or more new EC-Vs by providing some basic configuration and deployment details. For more information, see [Deploy EC-V in Cloud](#).

Third-Party Service Orchestration

Service Orchestration automates the integration of third-party services without an API. Service Orchestration automates the creation and deployment of IPSec tunnels and IP SLA probes and manages the lifecycle of the tunnels and probes. For more information, see [Service Orchestration](#).

Getting Started

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

In this section:

Supported Browsers	17
Guidelines for Creating Passwords	18
Overview of SD-WAN Prerequisites	19

Supported Browsers

Orchestrator and the Appliance Web user interfaces support the following browsers:

- Google Chrome (recommended)
- Microsoft Edge
- Mozilla Firefox
- Opera
- Safari

We recommend that you use the latest version available for your browser.

Guidelines for Creating Passwords

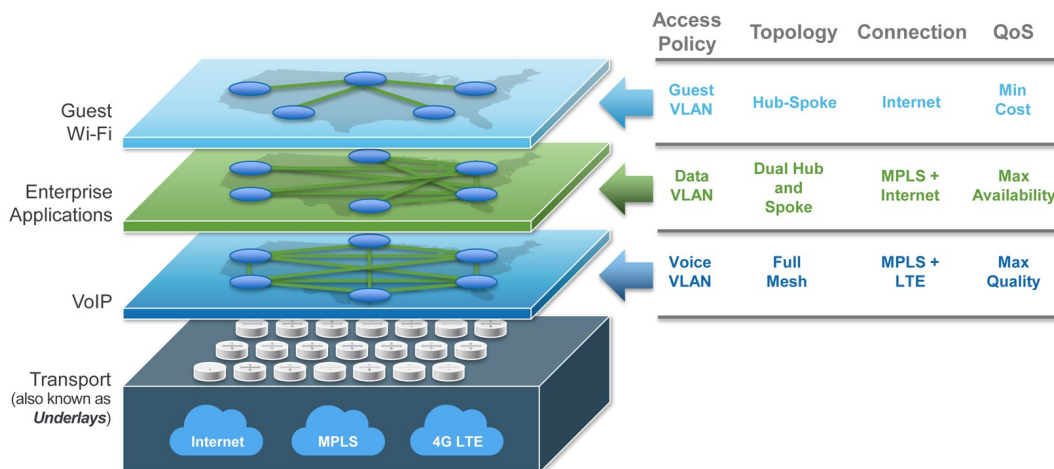
- Passwords should be a minimum of eight characters.
- There should be at least one lower case letter and one upper case letter.
- There should be at least one digit.
- There should be at least one special character.
- Consecutive letters in the password should not be dictionary words.

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.

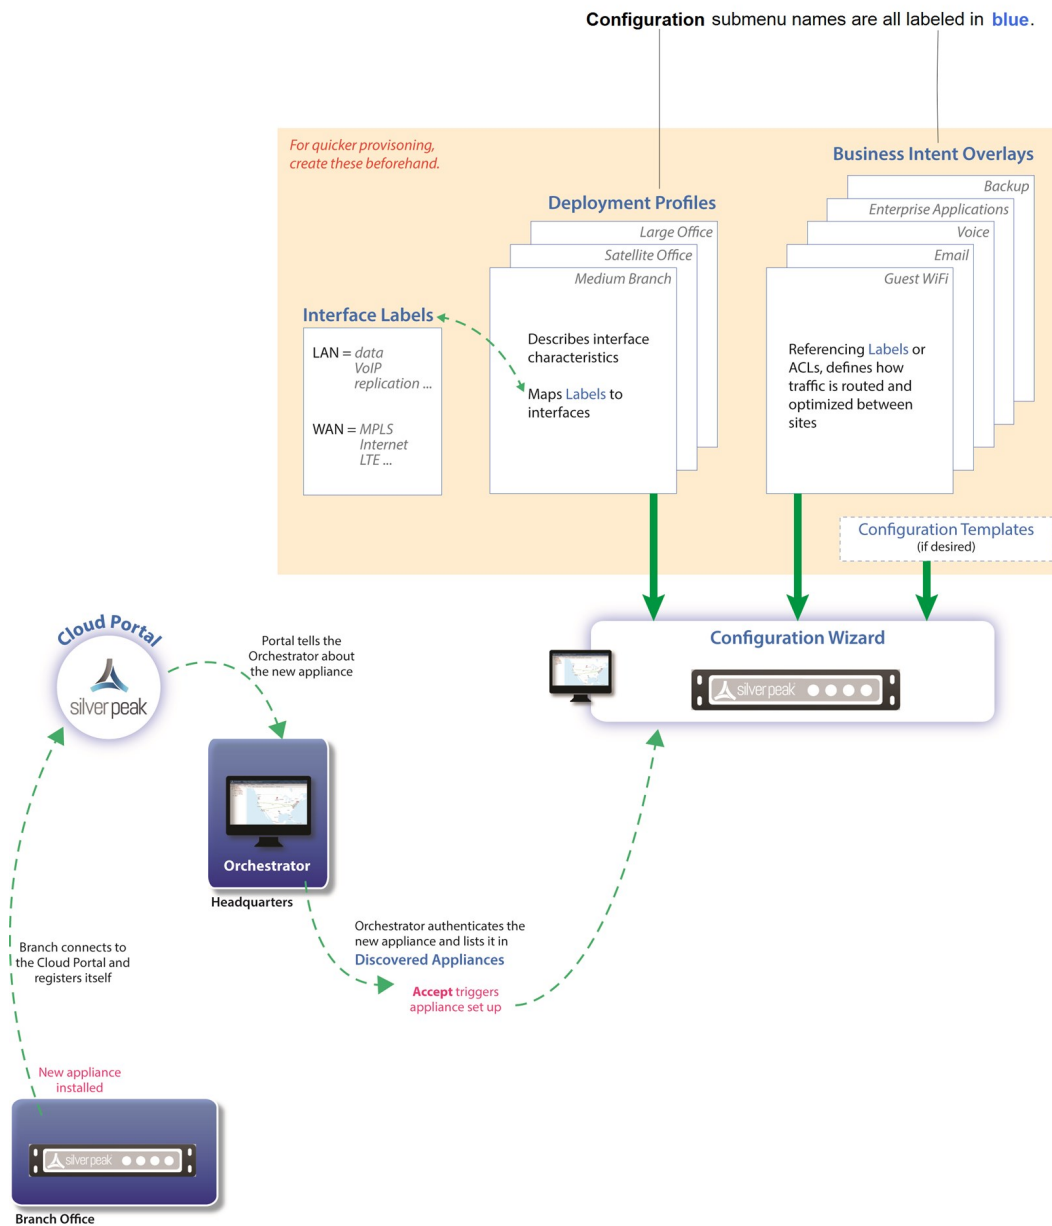
This diagram shows the basic architecture and capabilities of **Overlays**.



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.** Because the wizard prompts you to select profiles, it is easier to create these ahead of time.

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




Monitor Status and Performance

These topics focus on reports related to performance, traffic, and appliance status. Additionally, [Threshold Crossing Alerts](#) are helpful in monitoring your network.

Dashboard

Monitoring > Summary > Dashboard

The Dashboard integrates information from multiple components—or widgets—into a unified display for monitoring your network. It displays appliance license information, topology, health map, top talkers, top domains, and so forth, on one tab. The collection of widgets are customizable and persist for each user account.

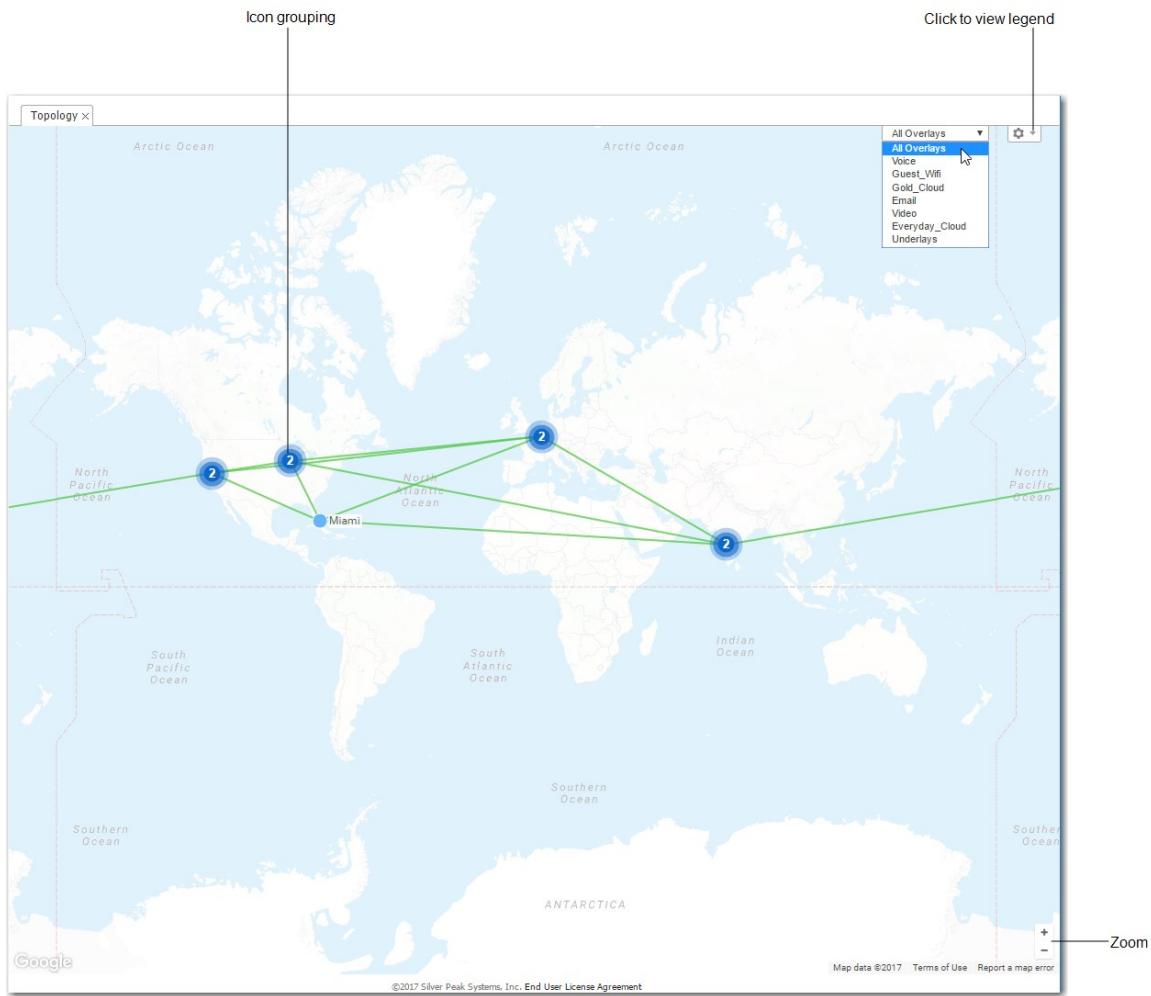
- Click **Settings** [] to select the widgets you want to show or hide.
- To move widgets, drag them by title.
- To access more detail in its corresponding tab, click a widget's title.
- To filter on various widgets, select **Src** or **Dest**, **Overlay** or **Underlay**, or **Inbound** or **Outbound**. The filter varies depending on the widget you are selecting.
- You can choose and change the grouping variable for Overlay-Transport and Overlay-Interface by clicking **Flip**.
- The **Appliance Licenses** widget displays an inventory by appliance model, as well as license type, availability, and usage.
- To search for appliances in the tree, enter an appliance name and the tag will be displayed above the tree.
- To filter collections of appliances, select **Show Tags** and select from among the tag options.

Topology Settings and Legend

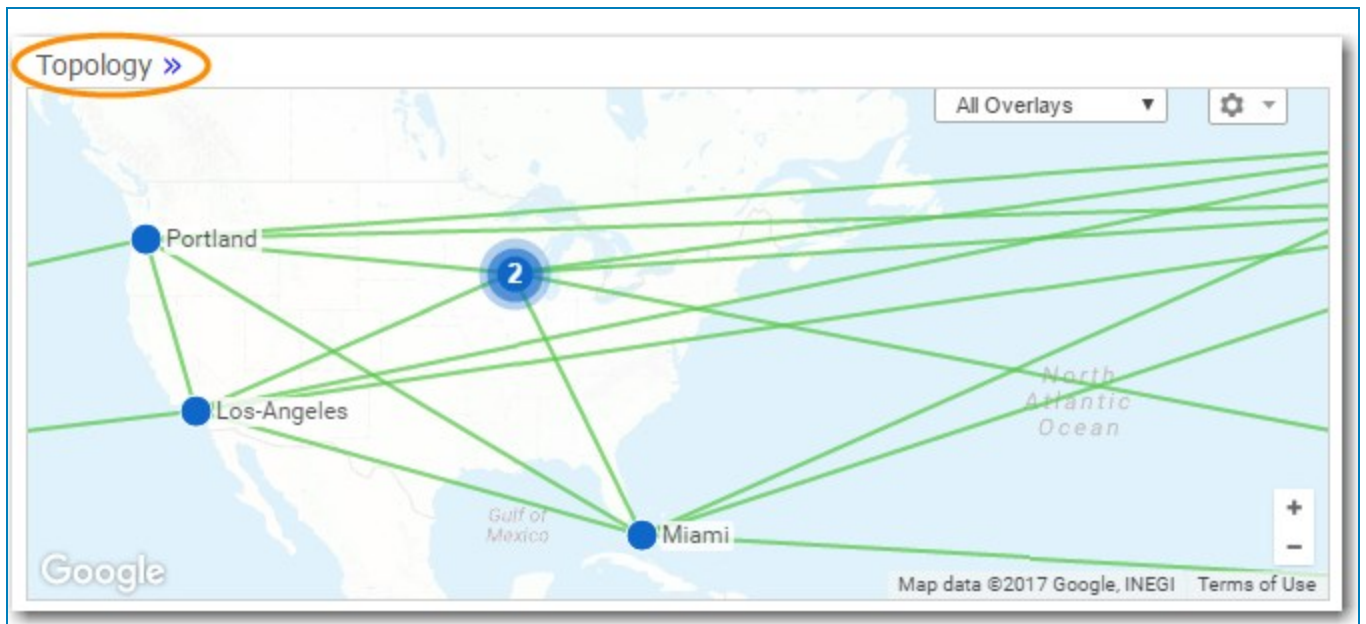
Monitoring > Summary > Topology

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

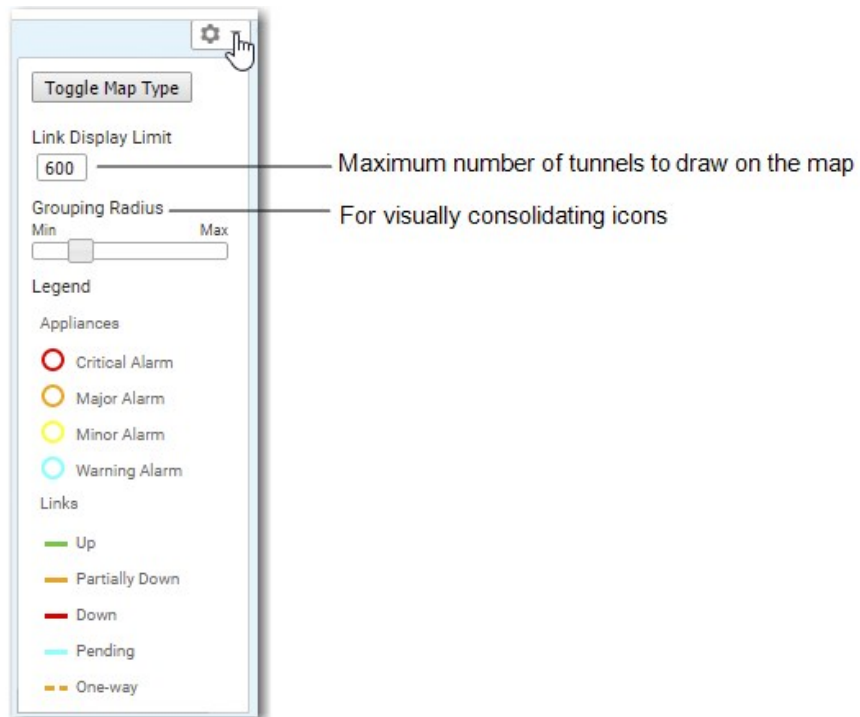
When configuring a software-defined WAN (**SD-WAN**), you can view **All Overlays**, individual **Business Intent Overlays** (BIOs), or the single and bonded **Underlay** tunnels that support them.



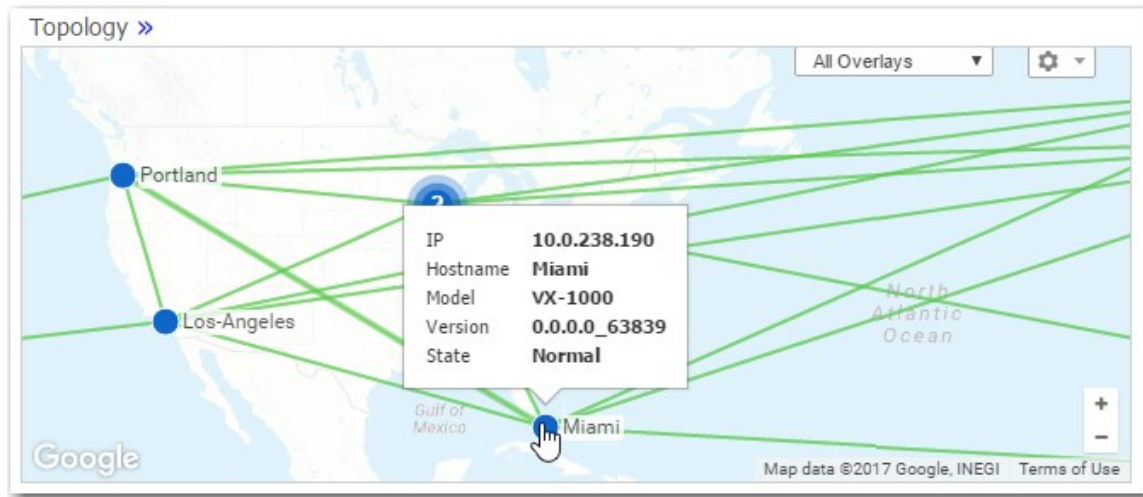
You can access it under **Monitoring** in the menu bar, or by clicking the widget title on the **Dashboard** tab.

Topology widget on Dashboard tab

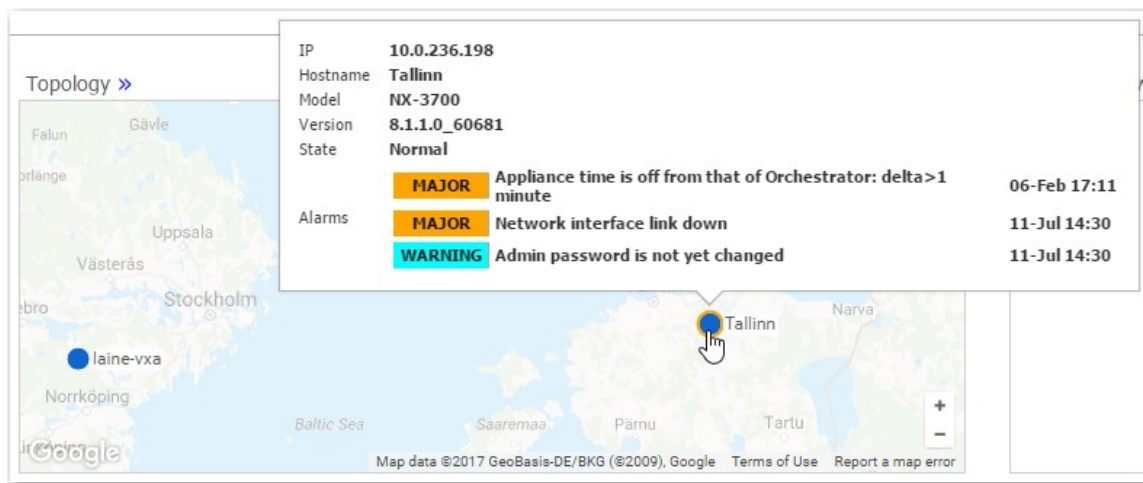
- The Legend details the appliances' management and operational states.



- 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.

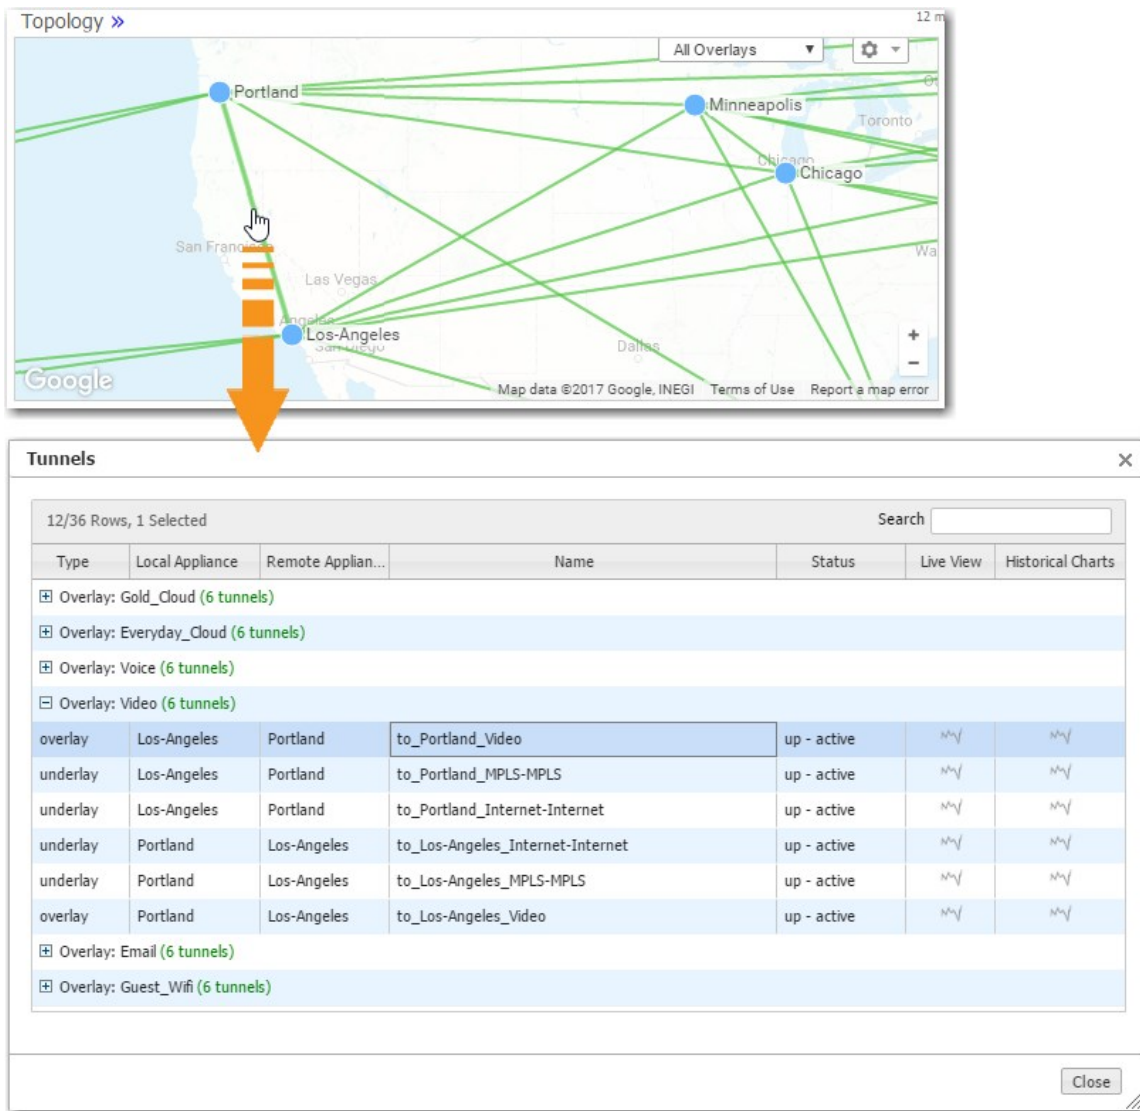


When the icon is encircled by a ring, indicating an alarm, those also display.



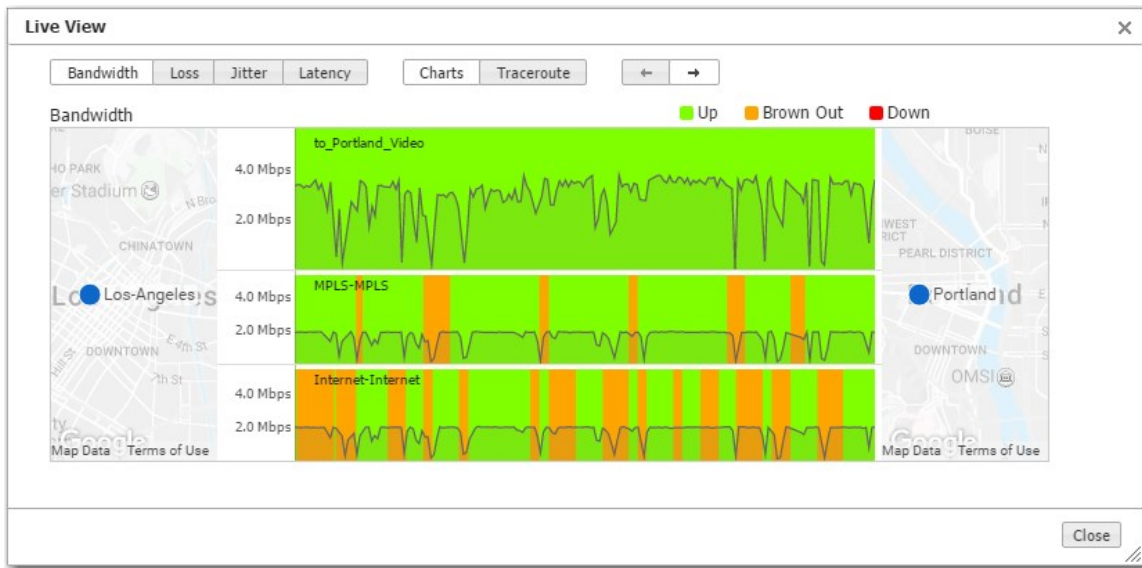
View Tunnels in the Topology Map

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



Live View

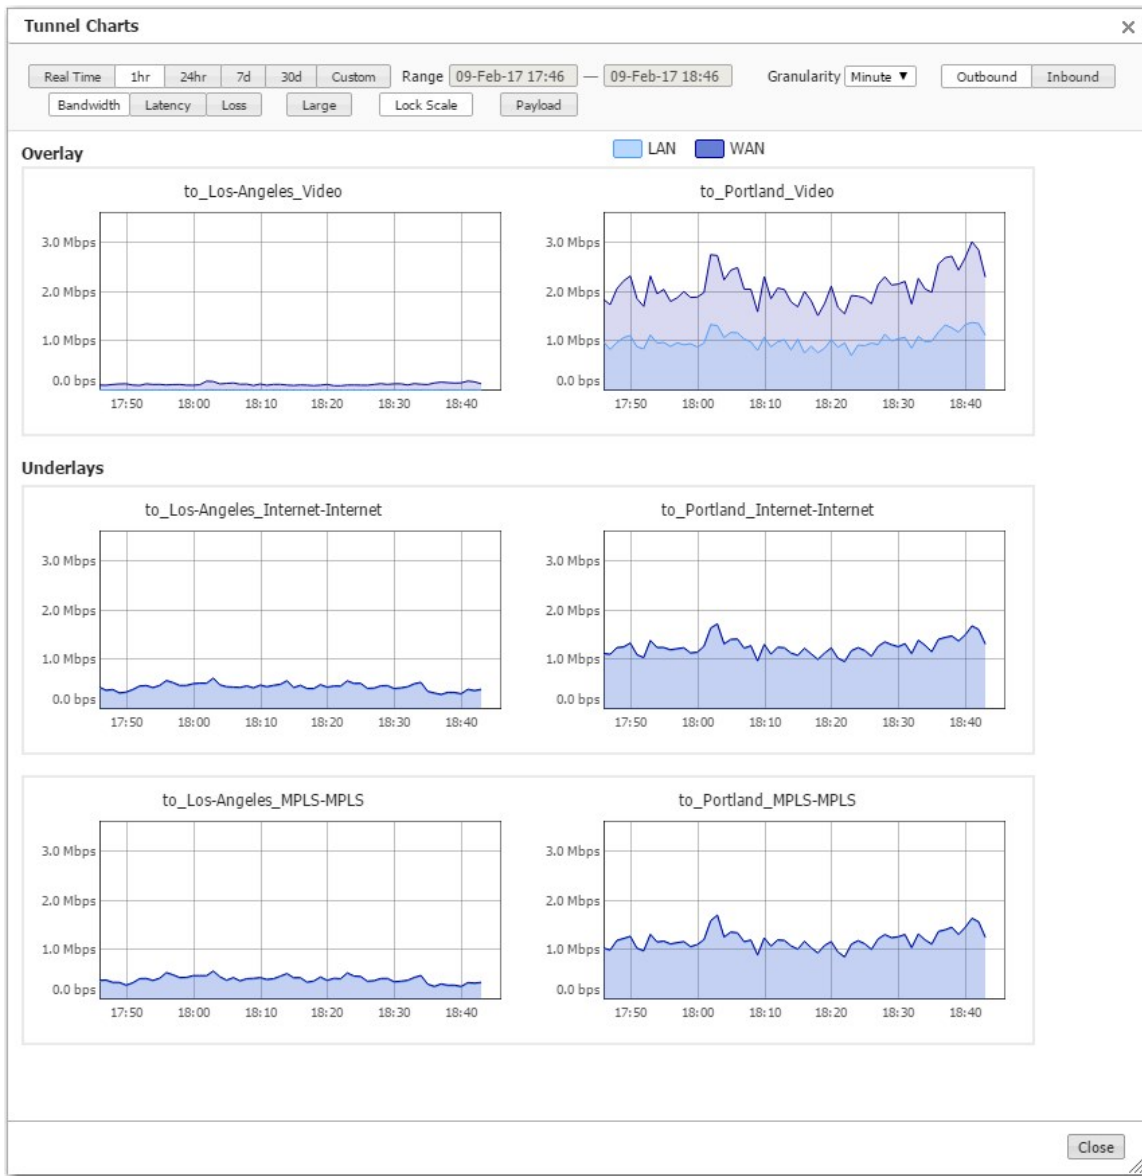
From the table, you can access the link's **Live View** graph.



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

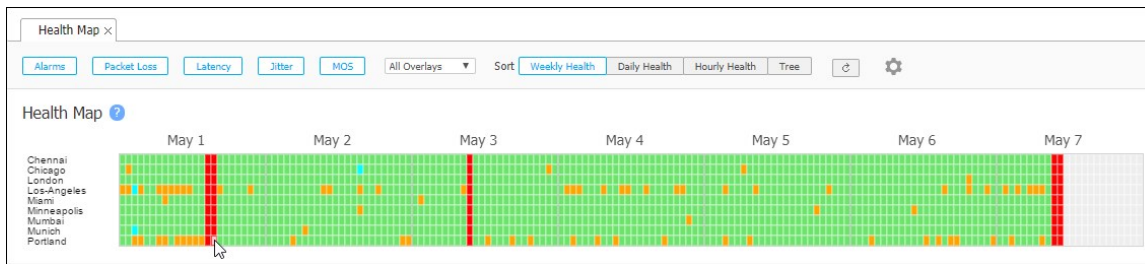
These charts enable you to selectively view the tunnel's components and behavior.




Health Map

Monitoring > Summary > Health Map

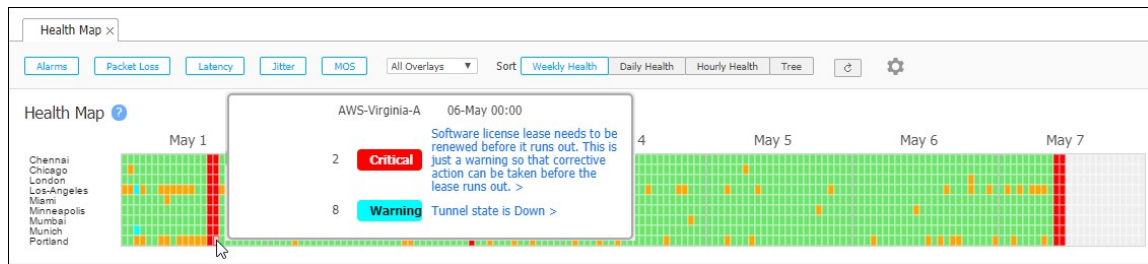
The **Health Map** provides a high-level view of your network's health, based on real-time measurements of network conditions between appliances.



- View filters are available for alarms, packet loss, latency, jitter, MOS (mean opinion score), and Business Intent Overlay.
- The health map can be sorted by weekly, daily, hourly health, or tree (by group, and then alphabetical by hostname).
- Each block represents one hour and uses color coding to display the most severe event among the selected filters. Color codes correspond to alarm severity and thresholds.
 - **Green** – Normal operation.
 - **Red** – Critical. Steps must be taken immediately in order to restore the affected service.
 - **Orange** – Major. Steps must be taken as soon as possible because the affected service has degraded drastically.
 - **Yellow** – Minor. A problem that does not yet affect service, but could if the problem is not corrected.
 - **Aqua** – Warning. A potential problem that could affect service.
 - **Grey** – No data available.
- Thresholds can be configured by clicking on the gear icon .



- Clicking a color 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.



- While filter and sort order customizations persist for each user account, threshold settings apply globally.
- Threshold settings are not retroactive. Setting new thresholds does not redisplay historical data based on newly edited values.
- Deleting an appliance deletes its data.
- If you are using overlays, note the following:
 - You can view each overlay's health individually.
 - If you remove an individual overlay, its data is not recoverable. However, its historical data remains included in **All Overlays**.

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** enables 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 about either Appliance alarms or Orchestrator alarms. Select **Add Recipient** in the **Alarm Recipients** window. Select the **alarm type** and select the check boxes that you want to receive emails for. Click **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.
- **Ack, Acked By, and Ack Time:** These columns indicate whether an acknowledgment has been received between devices.
 - **Acked By:** The name of the appliance that did the acknowledgment.
 - **Acked Time:** The time when the acknowledgment was received by the appliance.

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 could 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 on 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](#).

Schedule and Run Reports

Monitoring > Reporting > Schedule & Run Reports

Use the Schedule & Run Reports tab to create, configure, run, schedule, and distribute reports. You can specify what you want to include in your report based on appliances, the time range of the report, traffic type, and the types of charts to include. You also can specify email recipients for the report.

Reports and statistics help you bracket a problem, question, or analysis. Orchestrator reports fall into two broad categories:

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

Configure the following in this tab:

- **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**.
- **Name** of the report.
- **Email Recipients** – Enter the email address to which to send the report.
 - To send a test email or to configure another SMTP server instead, navigate to **Orchestrator > Software & Setup > Setup > SMTP Server Settings**.
 - If a test email does not arrive within minutes, check your firewall.
- Default range of reports – **Daily** = 14 days, **Hourly** = 24 hours. Increasing the scope uses additional memory.
- A Scheduled or Single Report.

Additionally, you can specify the following for a generated report:

- Appliances in Report – Fill in the box or click **Use Tree Selection** to display appliances.
- Amount of Top Reports (10, 25, 50, 100, 1000).
- Traffic Type.
- Select the check boxes next to the following charts to be included in the report:
 - Application Charts
 - Tunnel Charts
 - Appliance Charts
- Lock Scales for Local Trends – Automatically scales graphs for specified scheduled reports.

TIP To specify the timezone for scheduled jobs and reports, navigate to **Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**.

View Reports

Monitoring > Reporting > View Reports

Use this tab 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 also can be emailed depending on the configuration set on 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 and 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 ×

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

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...

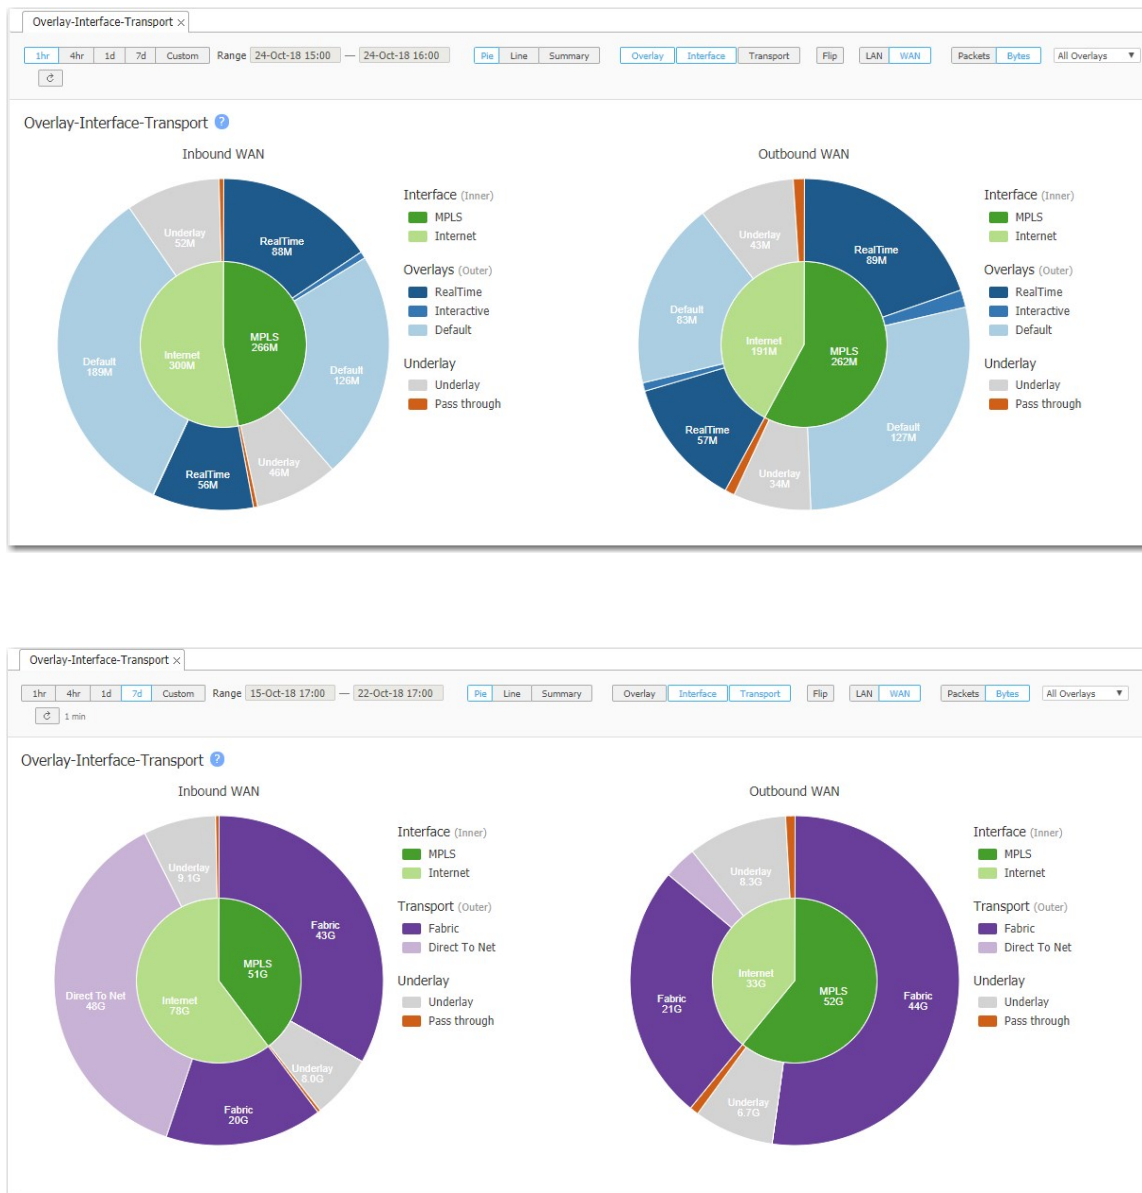
Overlay-Interface-Transport

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.

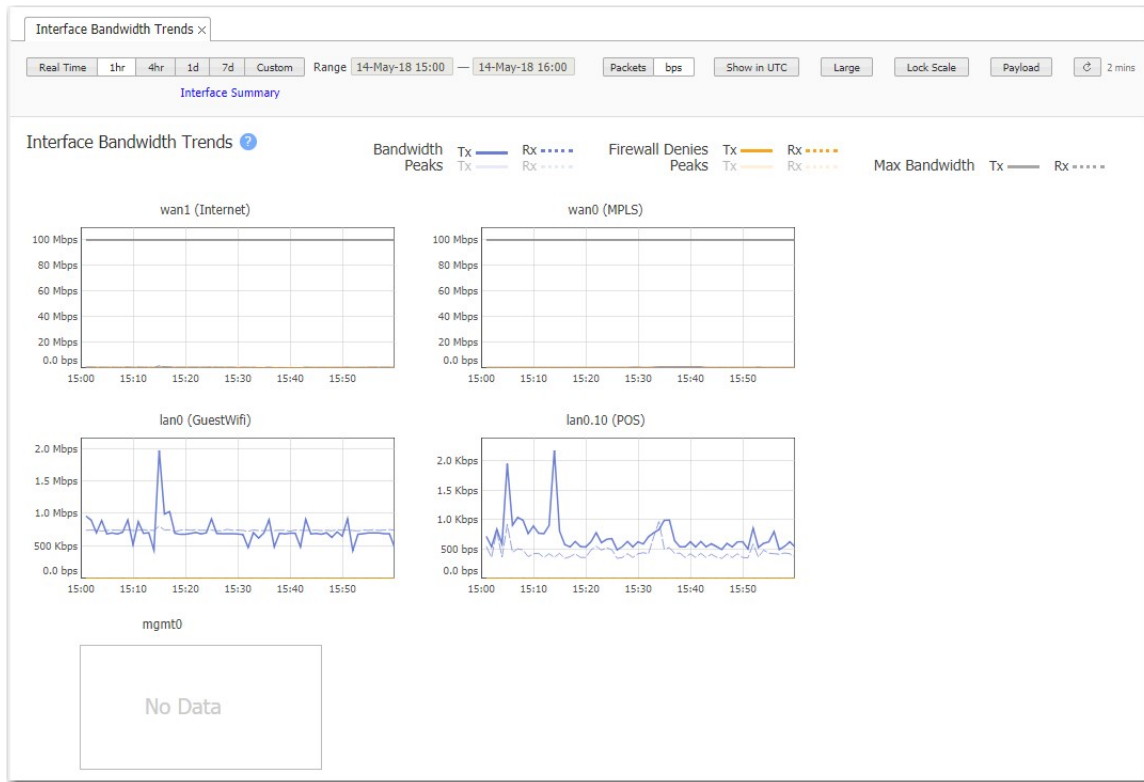
You also can view how much traffic is transported from one EdgeConnect appliance to another on the SD-WAN fabric (Overlays), versus how much is broken out locally, direct to the internet. The Underlay legend displays non-overlay traffic.



Interface Bandwidth Trends

Monitoring > Bandwidth > Overlays & Interfaces > Interface Trends

The Interface Bandwidth Trends tab shows interface statistics for a single selected appliance in real time or for a specific period. Real time charts show the past five minutes of usage and refresh every second. By default, charts display transmit and receive statistics for bandwidth and firewall denies. You can toggle peak statistics or maximum bandwidth statistics on or off by clicking the sample indicator line next to each statistic name.



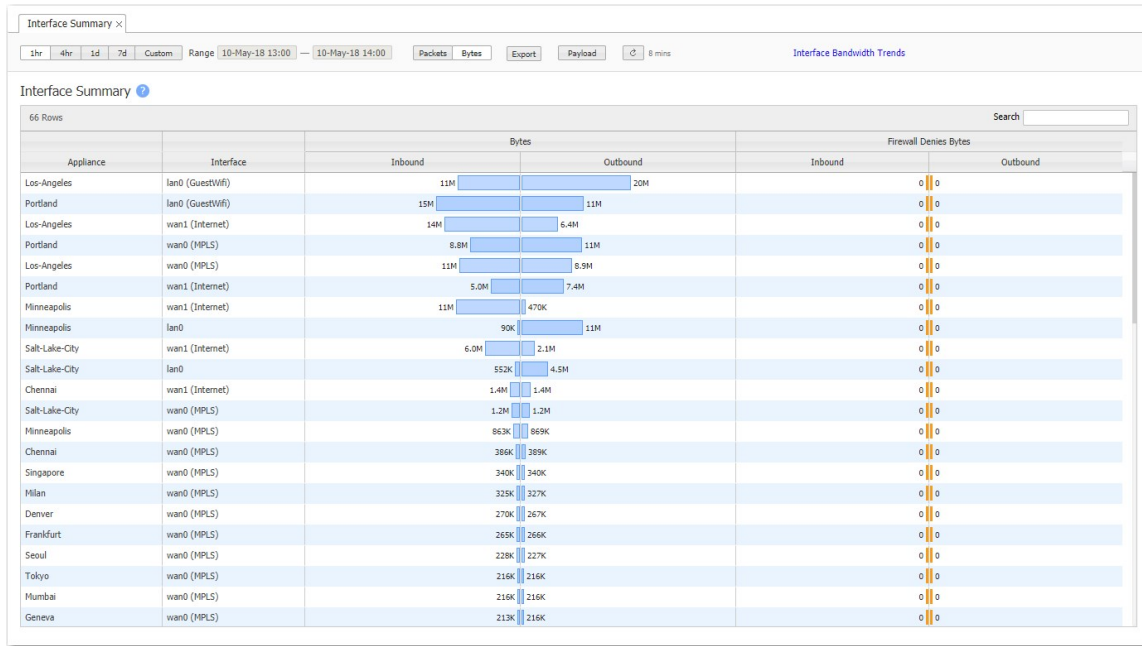
You can customize the chart settings using the controls at the top of the tab, as follows:

Option	Description
Time period	<ul style="list-style-type: none"> Click Real Time to enable live statistics for all available interfaces. Click a predefined time period (1h, 4h, 1d, 7d) to display statistics over the last hour, four hours, day, or seven days. Click Custom and set your own custom time range to display statistics for that time period.
Packets/bps	<ul style="list-style-type: none"> Click Packets to display statistics according to the number of packets sent and received. Click bps to display statistics for bits per second sent and received.
Show in UTC	Click this option to toggle chart times between local appliance time or UTC.
Large	Click this option to toggle the size of the charts between smaller (default) and large.
Lock Scale	By default, each chart uses its own scale that is relative to the data displayed. Click this option to apply and lock the same scale to each chart.
Payload	By default, charts show complete bandwidth usage statistics—payload plus all SD-WAN overhead (headers, FEC, and so forth). To see bandwidth usage for payload only, click to enable the Payload button.

Interface Summary

Monitoring > Bandwidth > Overlays & Interfaces > Interface Summary

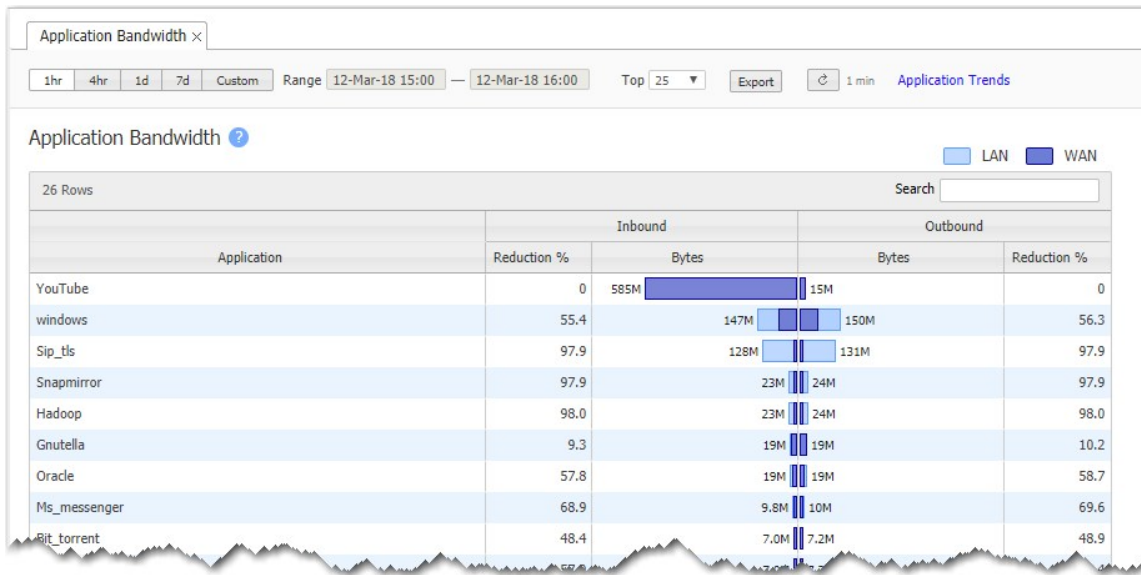
This tab shows interface summary stats, including inbound and outbound Packets or Bytes per interface, as well as Firewall Denies (Drops). The stats are summarized for the selected time period.



Application Bandwidth

Monitoring > Bandwidth > Applications > Summary

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

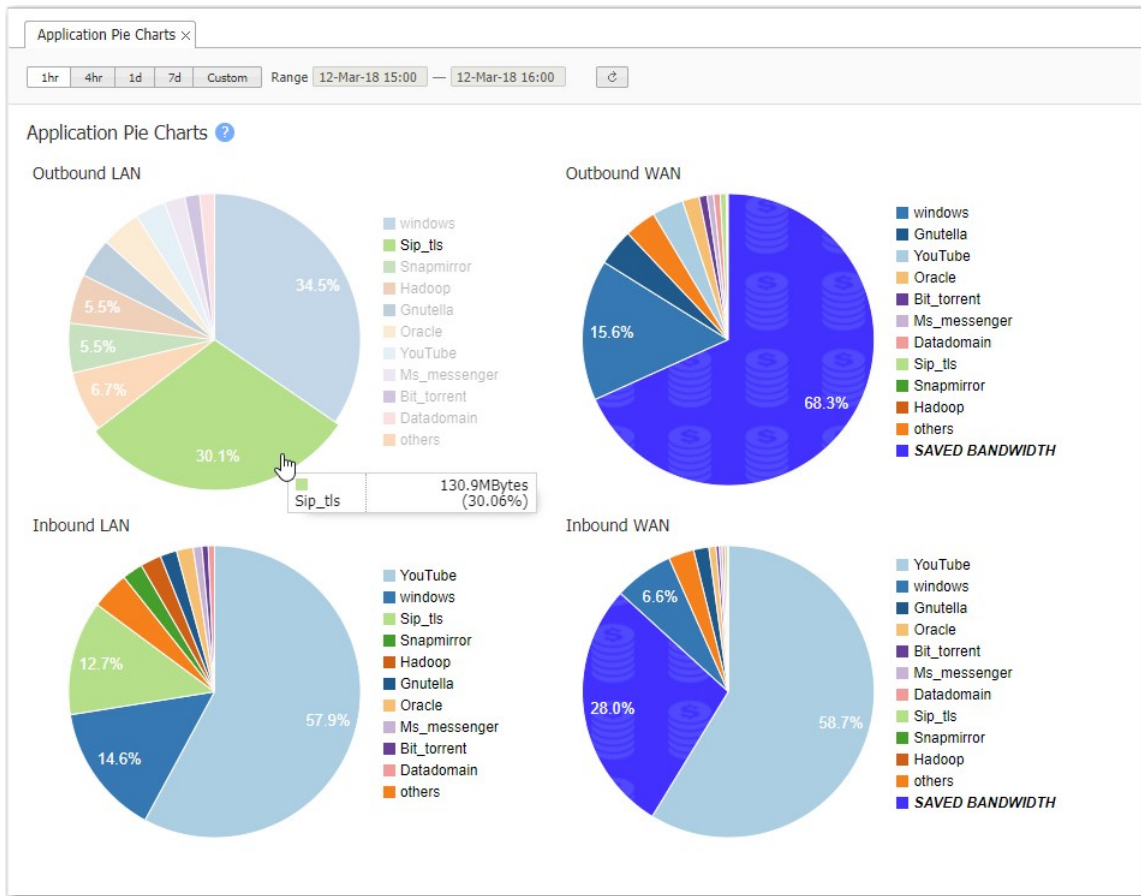


Application Pie Charts

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.

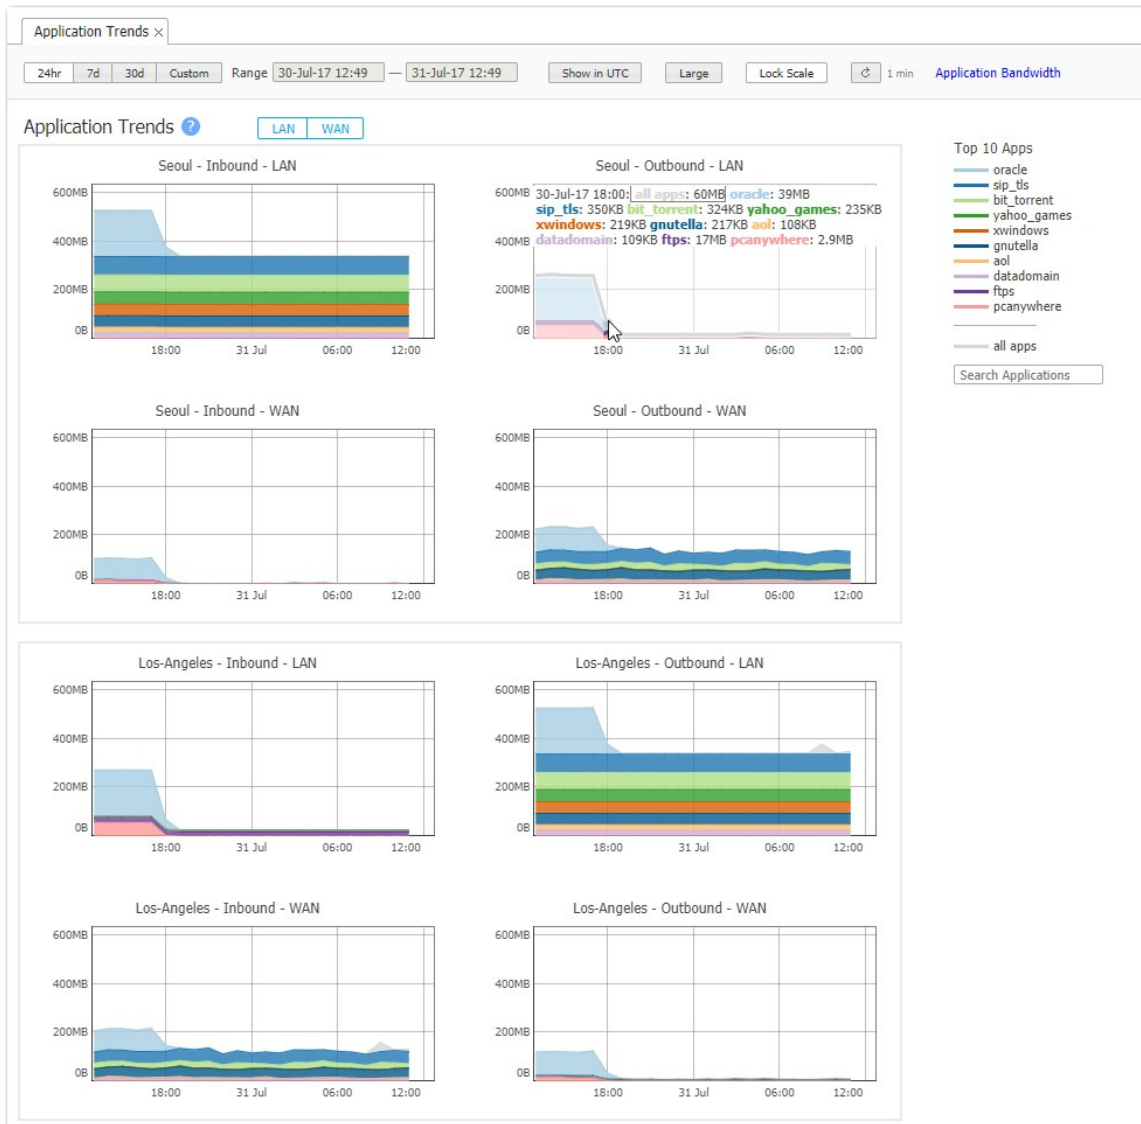
- Mousing over the charts and the legends reveals additional information.
- The WAN charts identify what percentage of the bandwidth the EdgeConnect appliance saved by optimizing the traffic.



Application Trends

Monitoring > Bandwidth > Applications > Trends

This tab shows application trends over time.



Top Talkers

Monitoring > Bandwidth > Identifiers > Top Talkers

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

Top Talkers ×

1hr4hr1d7dCustom

Range12-Mar-18 15:00 — 12-Mar-18 16:00

Top 25 ▾

Export

4 mins

Top Talkers ?

LAN

WAN

25 Rows

Search

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

You also can 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-n4v7knll.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-o097znll.googlevideo.com (74.125.166.168)	10M	192K	4	4
quote.cnbc.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 x

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

Domains ?

14 Rows Search

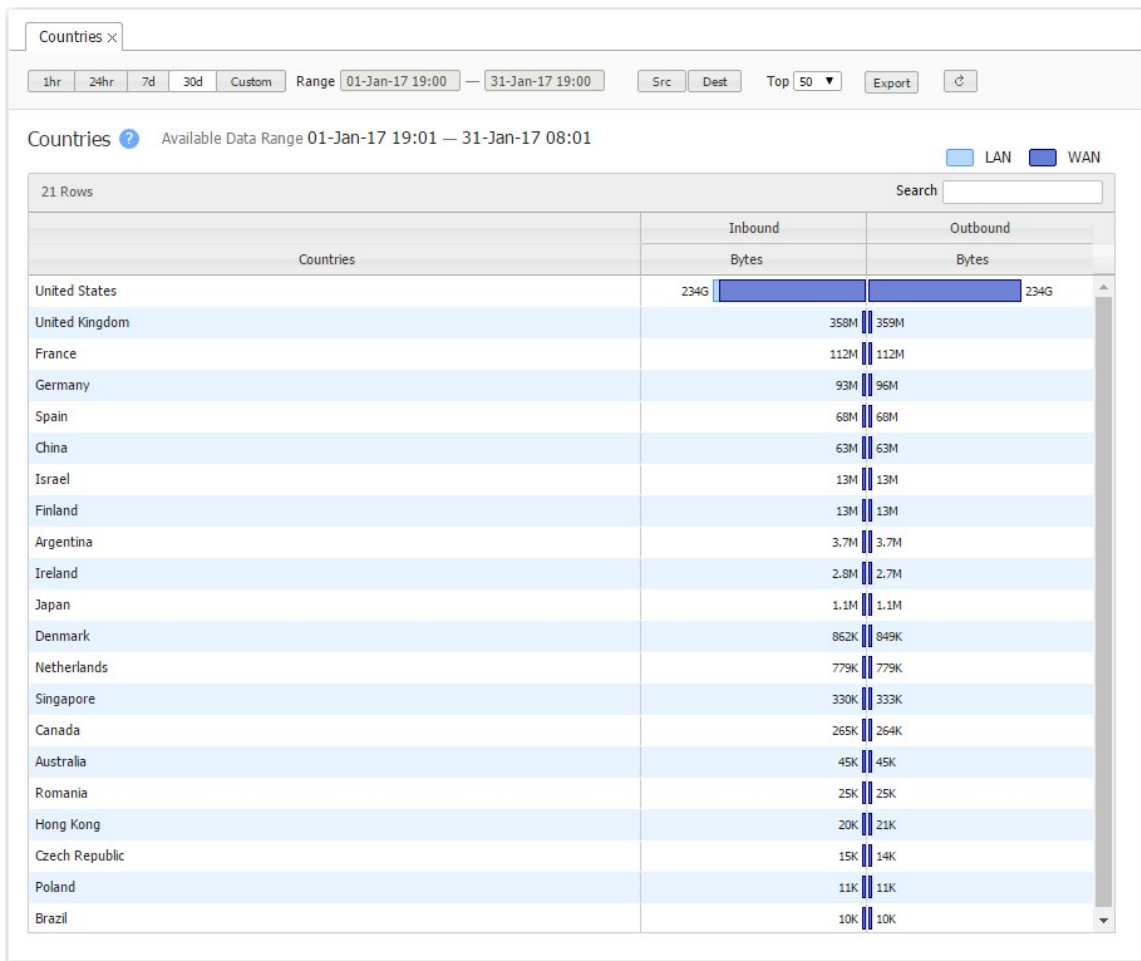
Domains	Inbound		Outbound	
	Reduction %	Bytes	Bytes	Reduction %
*googlevideo.com	0	145M	3.3M	0
*cnbc.com	0	1.7M	311K	0
*youtube.com	0	173K	197K	0
*nytimes.com	0	192K	147K	0
*mozilla.net	0	252K	12K	0
*nyt.com	0	97K	31K	0
*nr-data.net	0	32K	65K	0
*yimg.com	0	74K	5.1K	0
*doubleclick.net	0	28K	21K	0
*googleapis.com	0	24K	16K	0
*googlesyndication.com	0	13K	13K	0
*mozilla.com	0	12K	5.5K	0
*ggpht.com	0	5.8K	2.1K	0
*google.com	0	2.6K	2.8K	0

LAN WAN

Countries

Monitoring > Bandwidth > Identifiers > Countries

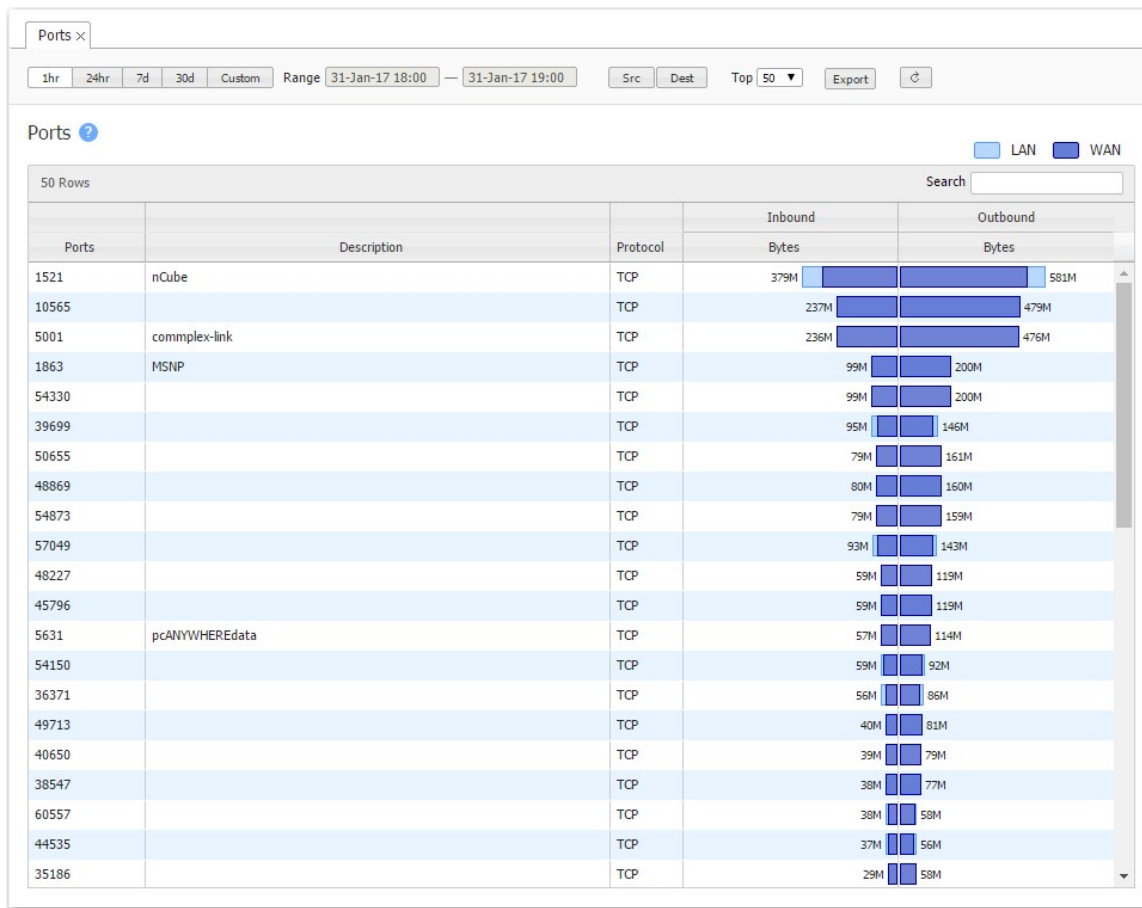
This tab lists the countries that use the most bandwidth.



Ports

Monitoring > Bandwidth > Identifiers > Ports

This tab lists the ports that use the most bandwidth.



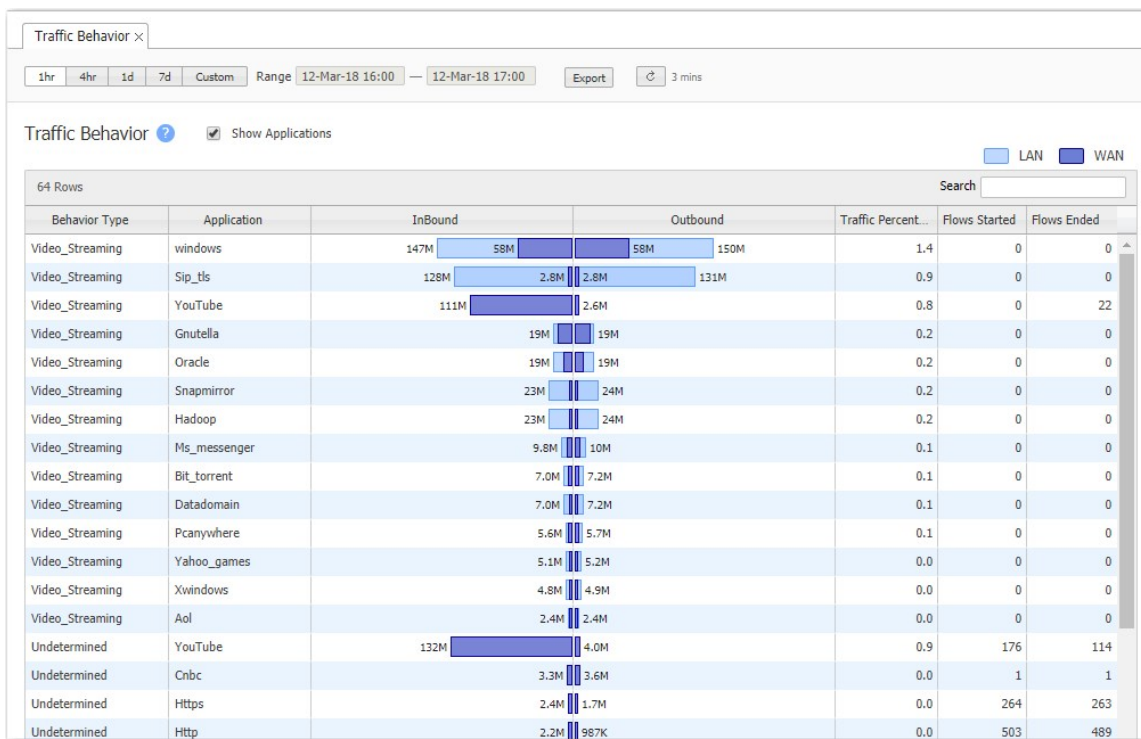
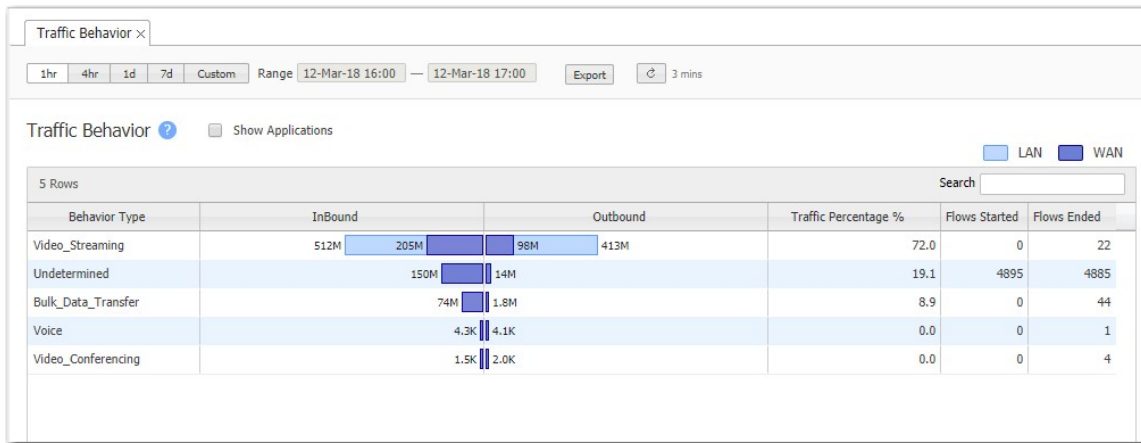
Traffic Behavior

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

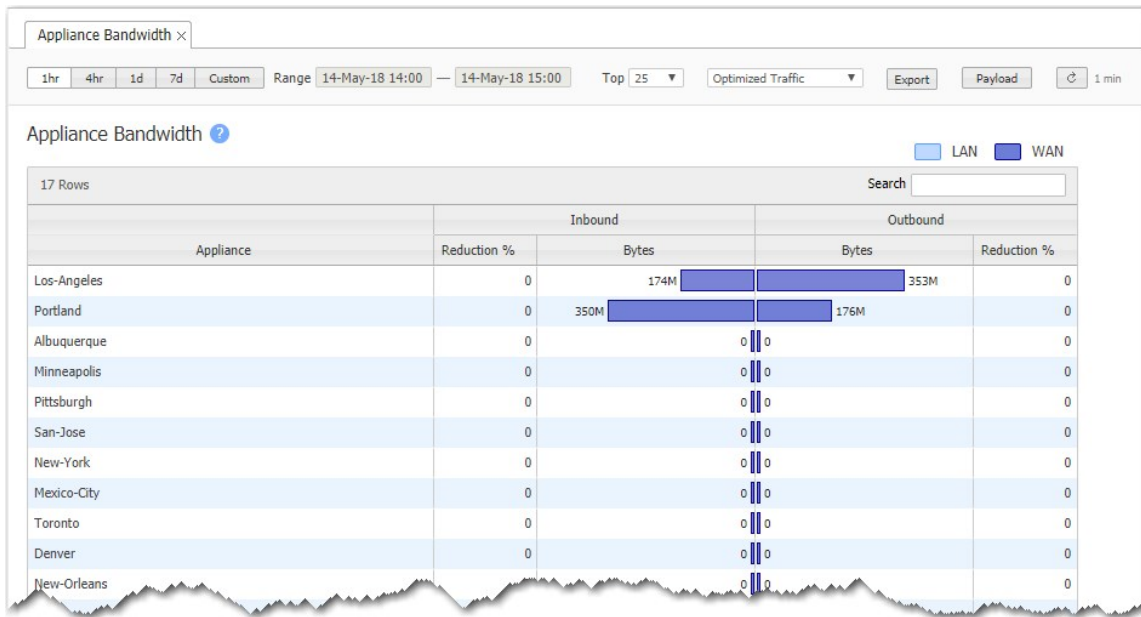
You also can specify these categories as match criteria when creating policies or ACLs (Access Control Lists).



Appliance Bandwidth

Monitoring > Bandwidth > Appliances > Summary

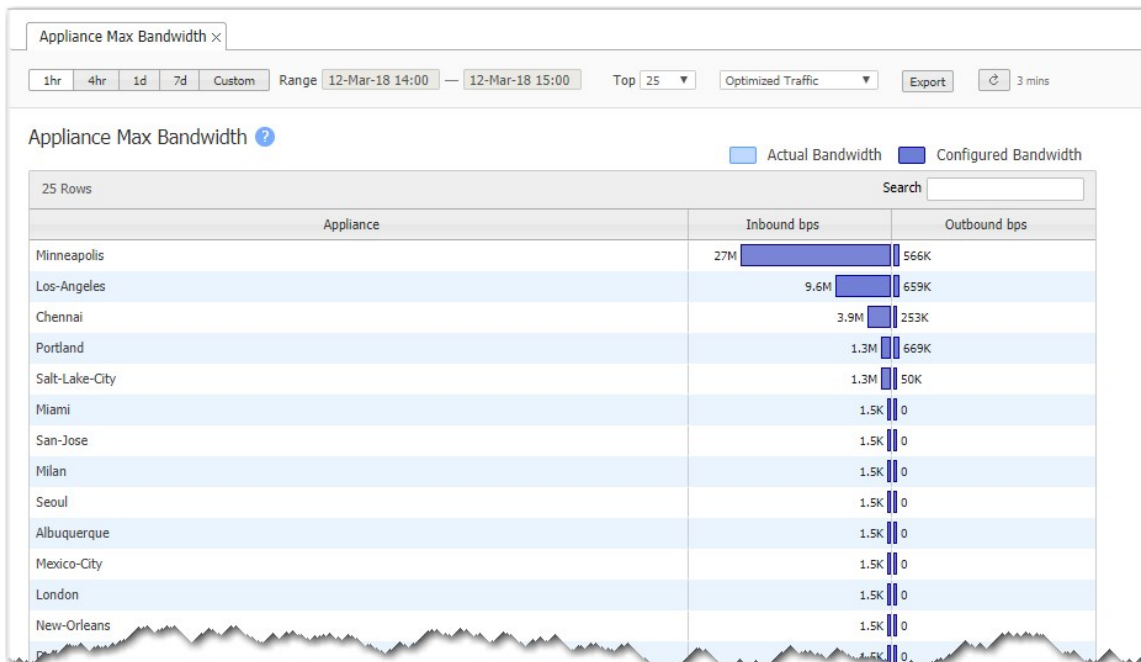
The **Appliance Bandwidth** chart lists the top appliances based on the total volume of inbound and outbound traffic before reduction. It shows how many bytes the EdgeConnect appliance saved when transferring data, aggregated over a selectable time period.



Appliance Max Bandwidth

Monitoring > Bandwidth > Appliances > Max

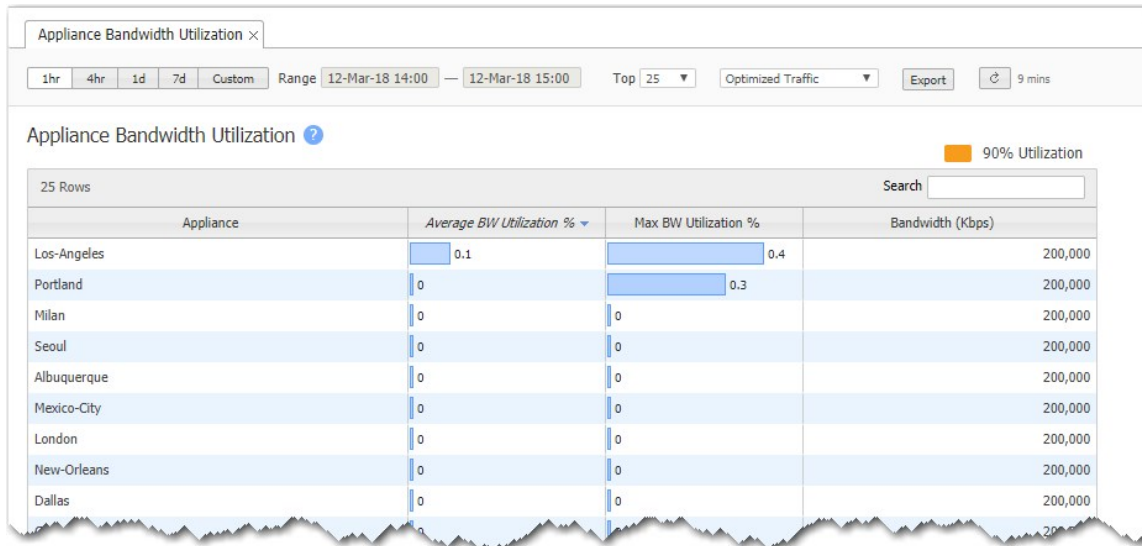
The **Appliance Max Bandwidth** chart lists the top appliances by the peak throughput (in either direction) within a selected time period. It compares the system bandwidth of the appliance to the effective bandwidth it is providing.



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 determine whether 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

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-Juan	0			0		262

Tunnels Bandwidth

Monitoring > Bandwidth > Tunnels > Summary

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

Tunnel Bandwidth

1hr 24hr 7d 30d Custom Range 01-Aug-17 15:18 — 01-Aug-17 16:18 Top 10 All Overlays Filter Tunnels Granularity Minute Export Payload 1 min

Tunnel Bandwidth ?

7 Rows Search

Appliance	Tunnel	Reduction %	Bytes		Reduction %	Destination Tunnel	Destination Appliance	Show Underlays	Live ...
			LAN	WAN					
Seoul	to_Los-Angeles_Default	96.6	492M	21M	60.3	to_Seoul_Default	Los-Angeles		✓
Los-Angeles	to_Seoul_Interactive	67.3	327M	122M	61.2	to_Los-Angeles_Interactive	Seoul		✓
Seoul	to_Los-Angeles_RealTime	87.7	415M	102M	32.2	to_Seoul_RealTime	Los-Angeles		✓
Singapore	to_Denver_Default	0	0	0	0	to_Singapore_Default	Denver		✓
Mexico-City	to_Chicago_RealTime	0	0	0	0	to_Mexico-City_RealTime	Chicago		✓
Minneapolis	to_Singapore_Interactive	0	0	0	0	to_Minneapolis_Interactive	Singapore		✓
Denver	to_Tokyo_RealTime	0	0	0	0	to_Denver_RealTime	Tokyo		✓

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).

Underlay Tunnels of to_Los-Angeles_Voice

2 Rows

Search

Appliance	Tunnel	Reduction %	Bytes	Bytes	Reduction %	Bandwidth (Kbps)	Remote Tunnel	Remote Appliance	Traceroute	
Portland	to_Los-Angeles_MPLS-MPLS	81.9	1.4G	584M	245M	0	4,000 (Auto)	to_Portland_MPLS-MPLS	Los-Angeles	why/
Portland	to_Los-Angeles_Internet-Internet	39.0		393M	240M	0	4,000 (Auto)	to_Portland_Internet-Int...	Los-Angeles	why/

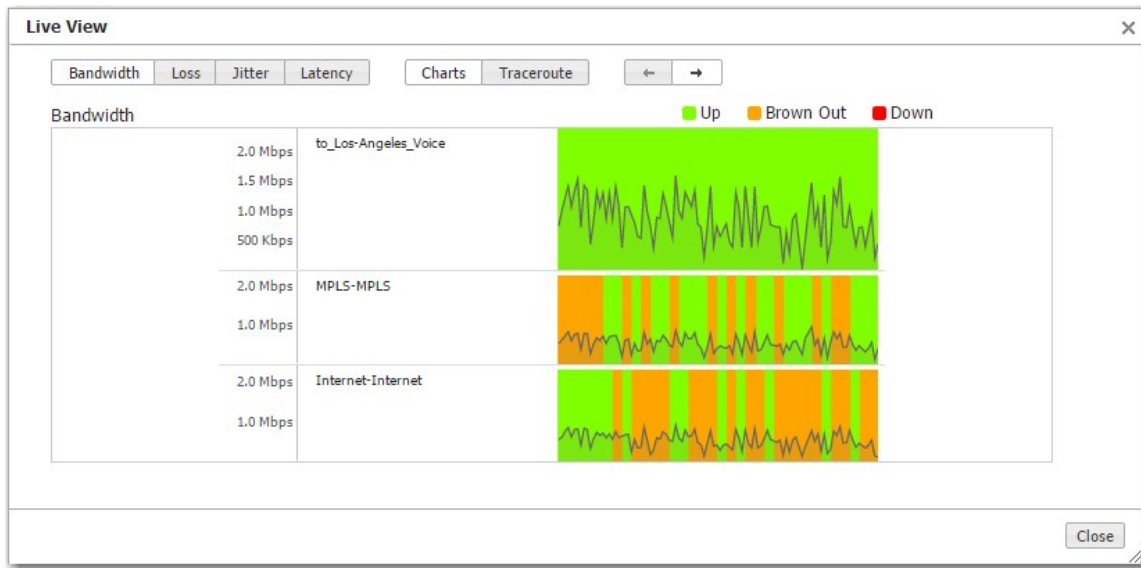
Traceroute

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**.



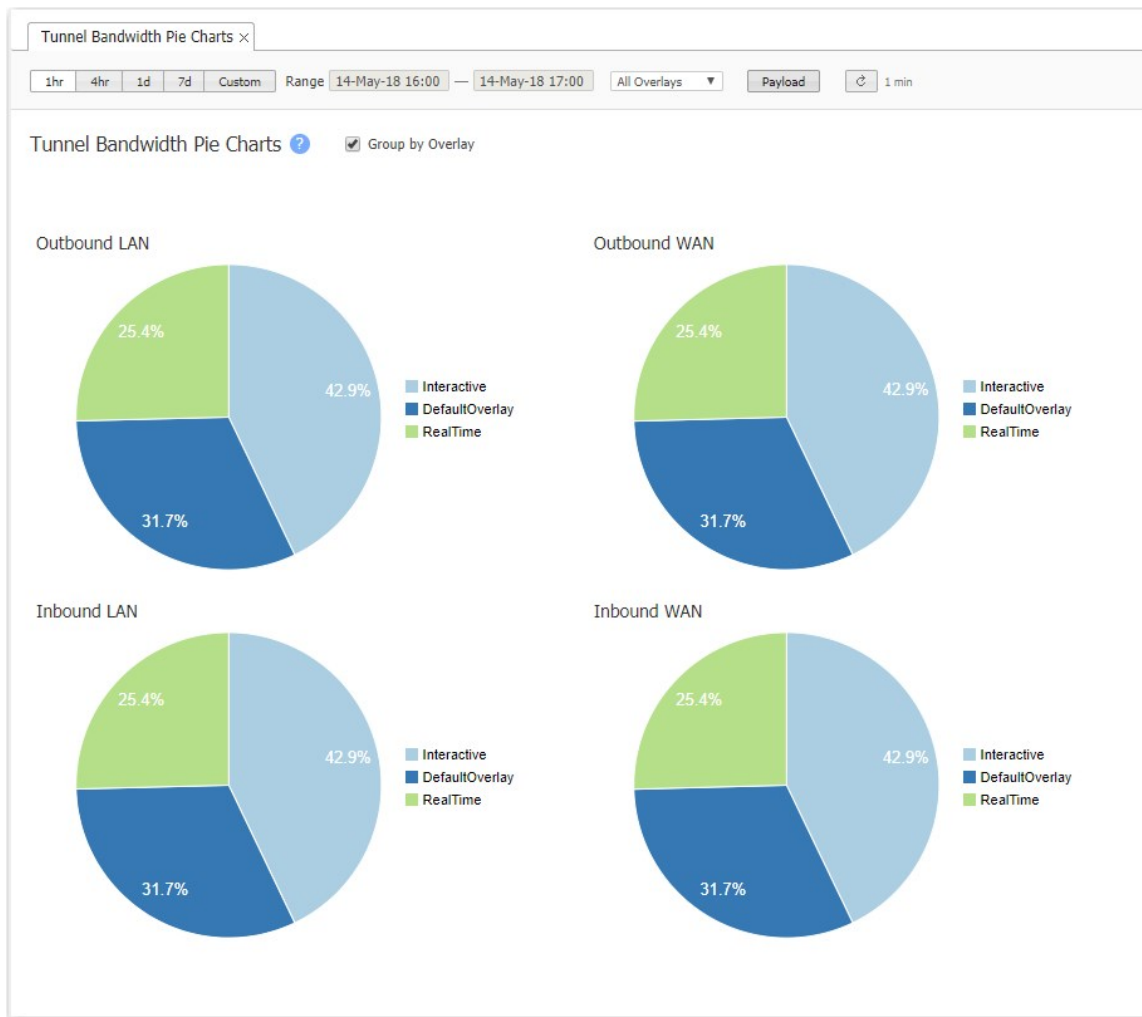
In real-time, LiveView shows how Silver Peak creates synergy to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and is delivering consistent application performance while both underlays are in persistent brown-out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

Tunnels Pie Charts

Monitoring > Bandwidth > Tunnels > Pie Charts

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

- Hovering over the charts and the legends reveals additional information.
- The WAN charts identify the 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 specified Link Bonding Policy 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 the tunnels that sent the most packets.

Tunnel Packet Counts

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

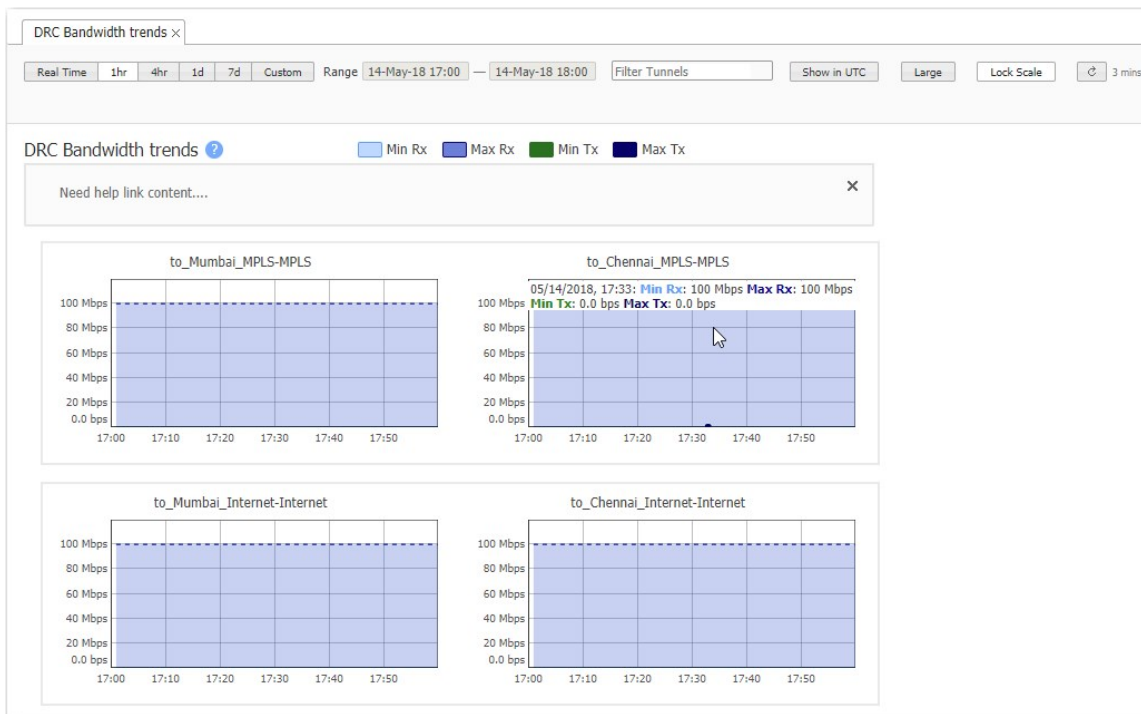
Appliance	Tunnel	Inbound			Outbound		
		LAN Packets	LAN Max pps	WAN Packets	LAN Packets	LAN Max pps	WAN Packets
Portland	to_Los-Angeles_Interactive	205,482	130	180,494	206,522	142	185,660
Los-Angeles	to_Portland_Interactive	204,394	144	183,705	207,239	134	182,053
Portland	to_Los-Angeles_DefaultOv...	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_Sao-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 Minimum 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** 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 and Recent

Monitoring > Bandwidth > Flows > Active & Recent Flows

The Flows tab enables 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	Description
Application	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.

Field	Description
App Group	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.
Domain	Includes the domain you can specify to filter your flow. Use the format <i>*.domain.*</i> or <i>*.domain.[com, info, edu, org, net, and so forth.]</i> Select the text field and a list displays. Choose the domain you want to apply.
Protocol	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.
IP/Subnet	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 Enter without selecting the Apply button if you want to do so.
Port	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.
Segment	Displays flows originating in the specified segment. Click the double arrow icon to enable both fields and filter by destination segments as well.
Zone	You can filter flows to the desired firewall zone. Select the text field and a list displays. If the From:To check box is not enabled, flows are filtered from and to the specified zone. If the check box is enabled, the flows are filtered from both the filtered From:To zones.
VLAN	Identifies the Virtual Local Area Network of a packet. Enter the VLAN ID you want to apply to your flow in the text field.
DSCP	Select the desired DSCP from the list. You can choose any or a specified DSCP from the list.
Overlay	The overlay the flow are applied. Overlays are defined on the Business Intent Overlay tab.
Transport	Select any of the three transport types: SD-WAN , Breakout , and Underlay . You also can apply a third-party service in this column if you have one configured.
Flow Characteristics	<p>You can apply any of the following flow characteristics to your flow: Boosted, Directly Attached, Pass-Through, Stale, Route Dropped, Firewall Dropped, Asymmetric, and Slow Devices.</p> <hr/> <p>NOTE You can select only one flow characteristic at a time.</p> <hr/>
Include EdgeHA	If not selected, Edge HA flows are excluded (default). If selected, the flows between Edge HA will be included.

Field	Description
Include Built-In	Includes the built-in policy flows. If not selected, they are excluded (default). If selected, they will be included.
Active/Ended	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 Custom is selected from the date widgets will be enabled to specify an exact time frame.
Duration	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).
Bytes	You can specify whether you want to filter flows that have transferred their total bytes or within the last five minutes.
Filter	This list has all the saved filters. When selected, the filter configurations are loaded. See more information below about the Filter 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. When 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 cannot 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 enable you to hide or unhide any selected fields.
- You also can select, drag, and drop any of the columns in the table to the order you want.

Additional Information about Flows

Note the following version specific and general information about flows:

ECOS 9.1 Behavior Changes

All flows in drop state are reset at flow reclassify time, overriding intervals described below.

ICMP/UDP Flows

- For any non-TCP connection (such as icmp, UDP), a flow is deleted only from inactivity.
- The inactivity timeout is three minutes for this type of flow. For example, after a ping connection is stopped, the flow still appears in the "Current Flows" for three minutes. This setting can be modified by using the system template.

TCP Non Accelerated Flows

- For a TCP connection, a flow is deleted under different timeouts. A half-open (single SYN) connection stays for two minutes if the connection does not establish correctly. A half-close (single FIN) or unclean-close (RST) deletes the connection after two minutes. A normal close (FIN-FIN) deletes the connection almost immediately.
- A TCP connection also has an inactivity timeout. If no activity is detected on an established TCP connection for 30 minutes (by default), the flow is deleted. This setting can be modified by using the system template.

TCP Accelerated Flows

- Timeout is determined by the configured Keep Alive Timers.
 - A heartbeat ACK is sent to idle endpoints after ten minutes.
 - If the endpoints have closed, an RST is returned and the connection is deleted after two more minutes due to the unclean-close.
- The timers can be modified per sequence number by using the Optimization Template.
 - Idle Timeout: The period of time that a TCP connection has to be idle before a keep-alive is sent. (Default 600 seconds)
 - Probe Interval: The time in seconds between each keep-alive probe. (Default 30 seconds)
 - Probe Count: The number of times TCP probes the connection to determine whether it is alive after the keep-alive option has been activated. The connection is assumed to be lost after sending this number of keep-alive probes. (Default 8)
- **Auto Reset Flows** - Enables or disables the auto-reset of TCP flows. If a connection is seen by an appliance but after the handshake already completed, the connection would normally remain but without TCP Acceleration. If this feature is enabled, and a connection is reclassified in the Flows report, around 30 seconds later, it will be reset. When the endpoints re-establish the flow, it now will be subject to the optimization and route policies it matches. This feature is disabled by default. It can be enabled per sequence number by using the Optimization Template.

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.

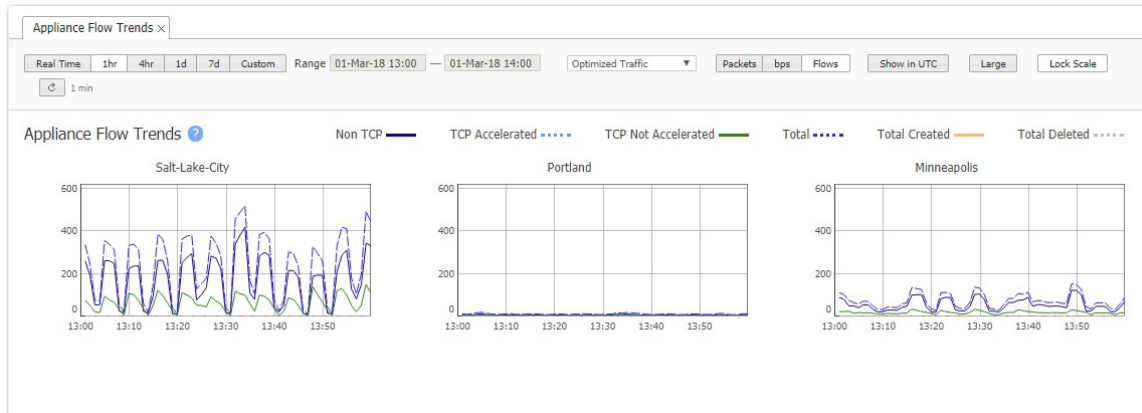
When you filter on **All Traffic**, the **Created** and **Deleted** columns display the number of new and ended flows for that same time period. The **Max** column value is from a one-minute window within the time range.

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													
Appliance	TCP Accelerated				TCP Unaccelerated				Non TCP				Search
	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** chart 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

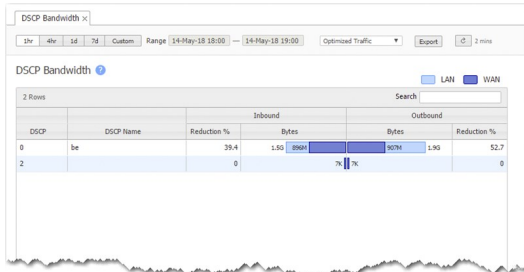
The **Tunnel Flow Counts** chart lists the tunnels with the most flows on average. It differentiates flows into TCP (accelerated and unaccelerated) and non-TCP, and also shows peak values.

Tunnel Flow Counts								
<div> 1hr 4hr 1d 7d Custom Range 14-May-18 17:00 — 14-May-18 18:00 Top 25 All Overlays Export </div>								
Tunnel Flow Counts								
25 Rows								
Appliance	Tunnel	TCP Accelerated		TCP Unaccelerated		Non TCP		Search
		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

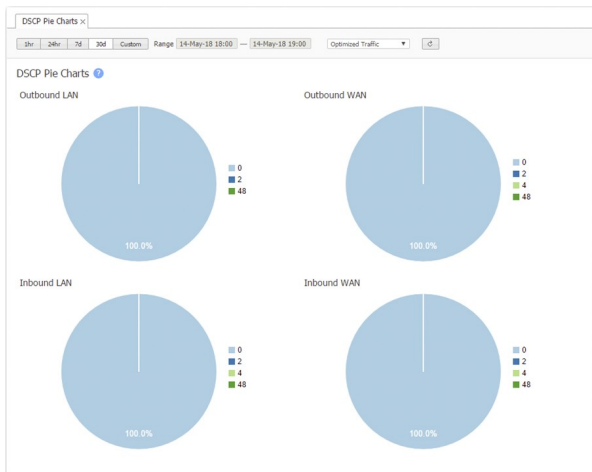
The **DSCP Bandwidth** chart shows the DSCP classes that are sending the most data.



DSCP Pie Charts

Monitoring > Bandwidth > DSCP > Pie Charts

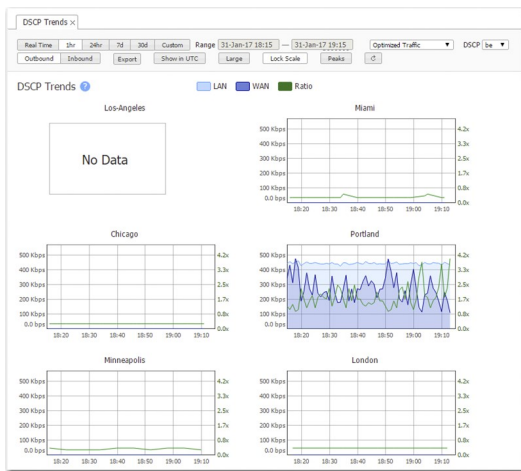
The **DSCP Pie Charts** show the proportion of traffic in each DSCP class. Hovering over the charts and the legends reveals additional information.



DSCP Trends

Monitoring > Bandwidth > DSCP > Trends

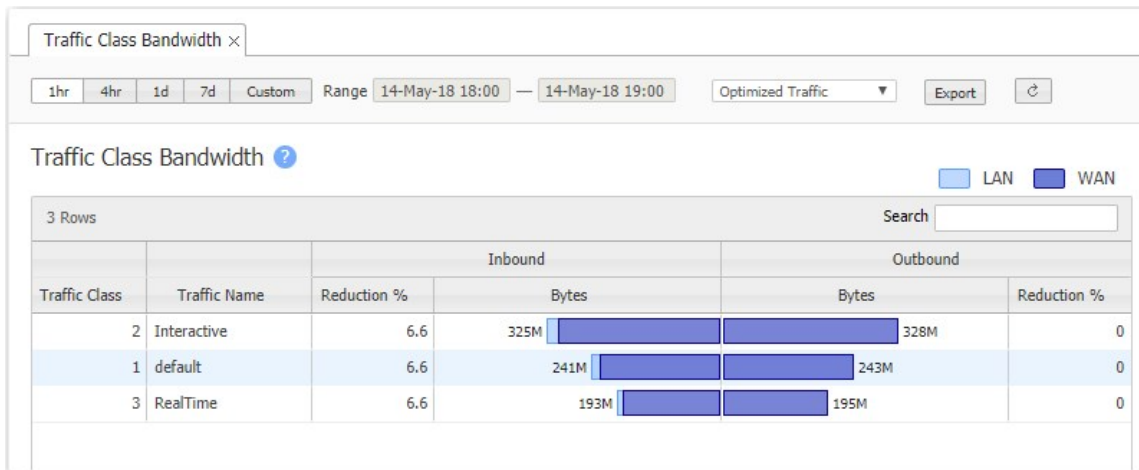
This tab shows DSCP usage over time.



Traffic Class Bandwidth

Monitoring > Bandwidth > QoS > Summary


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



Traffic Class Pie Charts

Monitoring > Bandwidth > QoS > Pie Charts

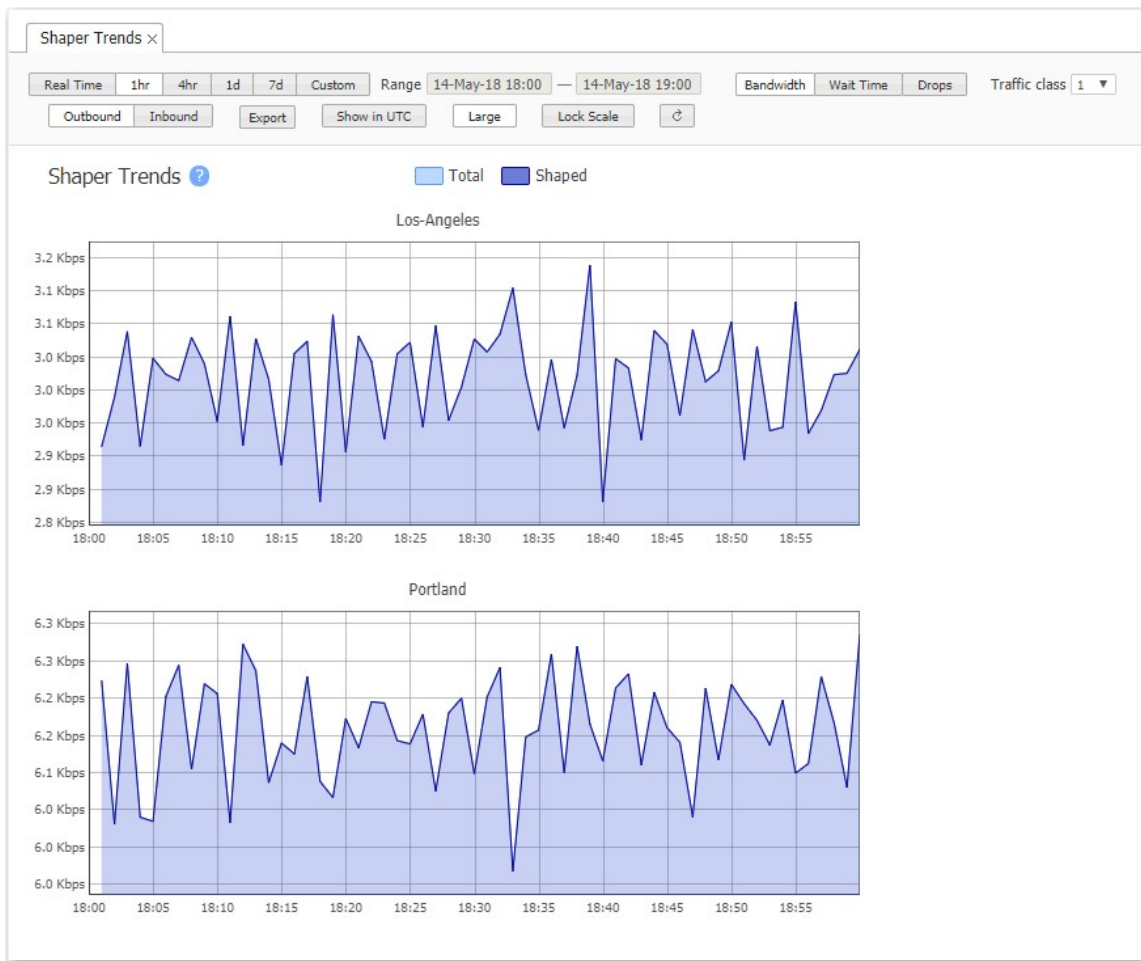
The **Traffic Class Pie Charts** show the proportion of traffic in each Traffic class. Hovering over the charts and the legends reveals additional information.



QoS (Shaper) Trends

Monitoring > Bandwidth > QoS > Trends

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



Shaper Summary

Use this tab to view the Shaper Summary for all traffic classes on selected appliances. The Shaper delays certain packet types to optimize overall network performance. For more information about shaping, see [Shaper Tab](#) and [Shaper Template](#).

- Use the controls above the table to specify how much data—time and date range—you want to see in the summary.
- Use the **Top X** filter to limit data according to top applications by total traffic bytes. You can include the top 10, 25, 50, 100, or 1000 applications.
- Click **Outbound** or **Inbound** to change the summary by traffic direction.

The following information is included in the Shaper Summary:

Field	Description
Appliance	Name of the appliance that is shaping traffic to generate the Shaper Summary.

Field	Description
Traffic Class	Traffic classes defined by Shaper parameters. The following four are pre-configured by Orchestrator: Real-time, Interactive, Default, and Best Effort. The user can configured the remaining six classes.
Total Bytes	Total amount of bytes being shaped.
Shaped Bytes	Amount of bytes used for shaping.
Shaped Packets	Amount of packets used for shaping.
Average Wait Time (ms)	Specified amount of time Orchestrator waits until packets are dropped while shaping is in progress.
Drop Packets	Amount of packets that have been reported as dropped due to expiration in the Shaper queue.
Other Drops	Refers to all other drops besides the expired drop packets.
Trends	Click the graph icon to see the Shaper Bandwidth Trends charts, which show Inbound and Outbound traffic trends in graphs.

Boost Tab

Monitoring > Bandwidth > Boost > Summary

This tab provides a summary of the Boost configuration and usage for selected appliances. You can change the time period for which Boost statistics are displayed by using the **1hr**, **4hr**, **1d**, and **7d** buttons at the top of the tab, or click **Custom** to specify a custom date range and granularity.

1hr	4hr	1d	7d	Custom	Range	30-Jun-21 09:23	—	07-Jul-21 09:23	Export	↻	4 mins
Boost ? EC Boost 7,010,000 Kbps / 10,000,000 Kbps Used 2,990,000 Kbps Remaining Configure Boost											
2 Rows Search <input type="text"/>											
Appliance	Configured Boost (Kbps)	% Time Insufficient Boost	Minutes Insufficient Boost	Total Boost Bytes	Trends						
Remediated Proxy	100000	0	0	3.1M							
Remediated Proxy	50000	0	0	19.1M							

This tab provides the following details about your Boost configuration:

Field	Description
Appliance	Name of the appliance.
Configured Boost (Kbps)	Boost bandwidth configured on the appliance.

Field	Description
% Time Insufficient Boost	Percentage of time when Boost bandwidth was not available for use.
Minutes Insufficient Boost	Amount of time (in minutes) when Boost bandwidth was not available for use.
Total Boost Bytes	Total amount of Boost bandwidth used over the specified time range.
Trends	Graph displaying detailed Boost trends for the specified appliance.


The total Boost bandwidth available to your network is controlled by your license. If necessary, you can purchase additional Boost bandwidth.

If a Boost license is available, you can assign Boost to appliances on the Licenses tab or on an appliance's Deployment page. You also can configure Boost allocation using Business Intent Overlays.

NOTE Your network uses a single queue for Boost across all appliances. When that queue is completely utilized, appliances will have insufficient Boost for any additional demands.

Boost Trends

To view Boost trends for a specific appliance, click the graph icon in the Trends column. The Boost Trends graph displays Configured Boost, Boost, and Minutes Insufficient Boost over the time period specified on the Boost tab.



Change Boost Configuration

To change the Boost configuration of one or more appliances selected in the table, click **Configure Boost**. You can increase or decrease Boost bandwidth by 20%, or set the bandwidth to a specific value in Kbps. Click **Apply** to save and apply your changes, or click **Close** to cancel.

Update Boost Bandwidth

Increase 20%

Decrease 20%

Set to this Value
Kbps

These changes may take a few moments

Apply

Close

Firewall Drops

Monitoring > Bandwidth > Firewall Drops > Summary

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.

- You can select a range of time (in hours and days) to view the firewall drops. You also can select 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

2hr

4hr

1d

7d

Custom

Range 11-Dec-18 08:00 — 11-Dec-18 09:00

Matrix View

Table View

Export

11 mins


Firewall Drops

30 Rows

Search

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

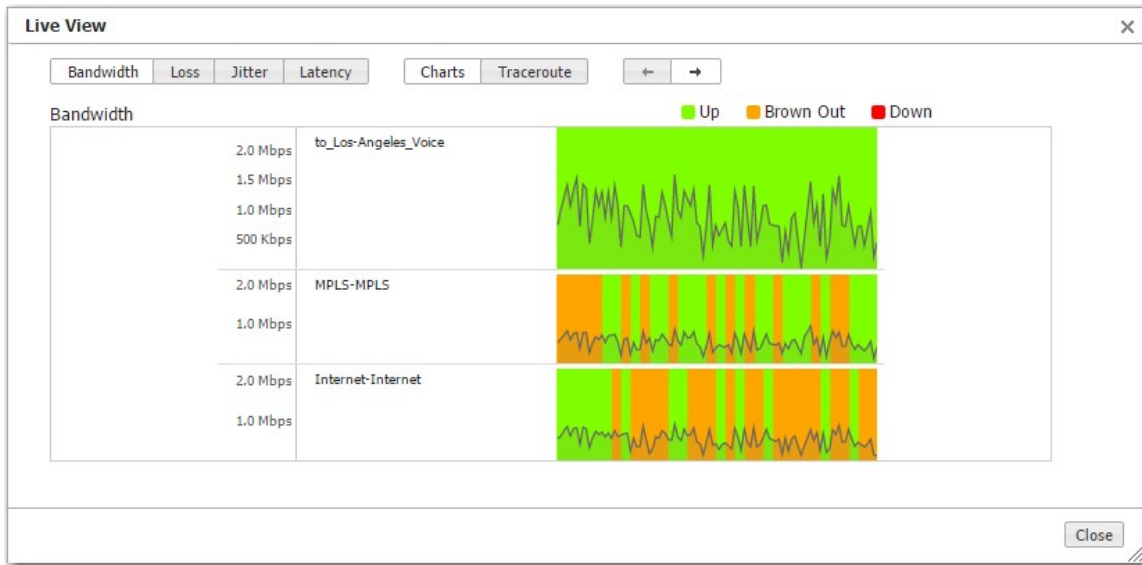
- If segmentation is enabled, you can specify the **Source Segment** and the **Destination Segment** to search for the flows, packets, and firewall drops in that segment.
- In the charts column, you can select the chart icon.
 - In this pop-up, you can see packets, and bytes dropped or allowed by a zone-based firewall for a given time range.



Live View

Monitoring > Tunnel Health > Live View

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

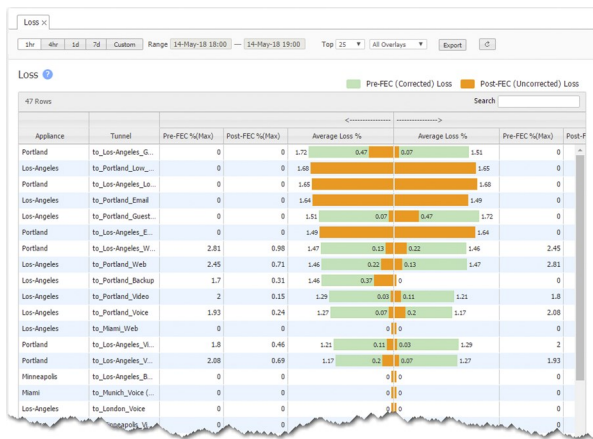


In real-time, LiveView shows how Silver Peak creates synergy to maintain coverage. The real-time chart shows the SD-WAN overlay at the top and the underlay networks at the bottom. The overlay is green and delivering consistent application performance while both underlays are in persistent brown-out state.

Loss

Monitoring > Tunnel Health > Loss > Summary

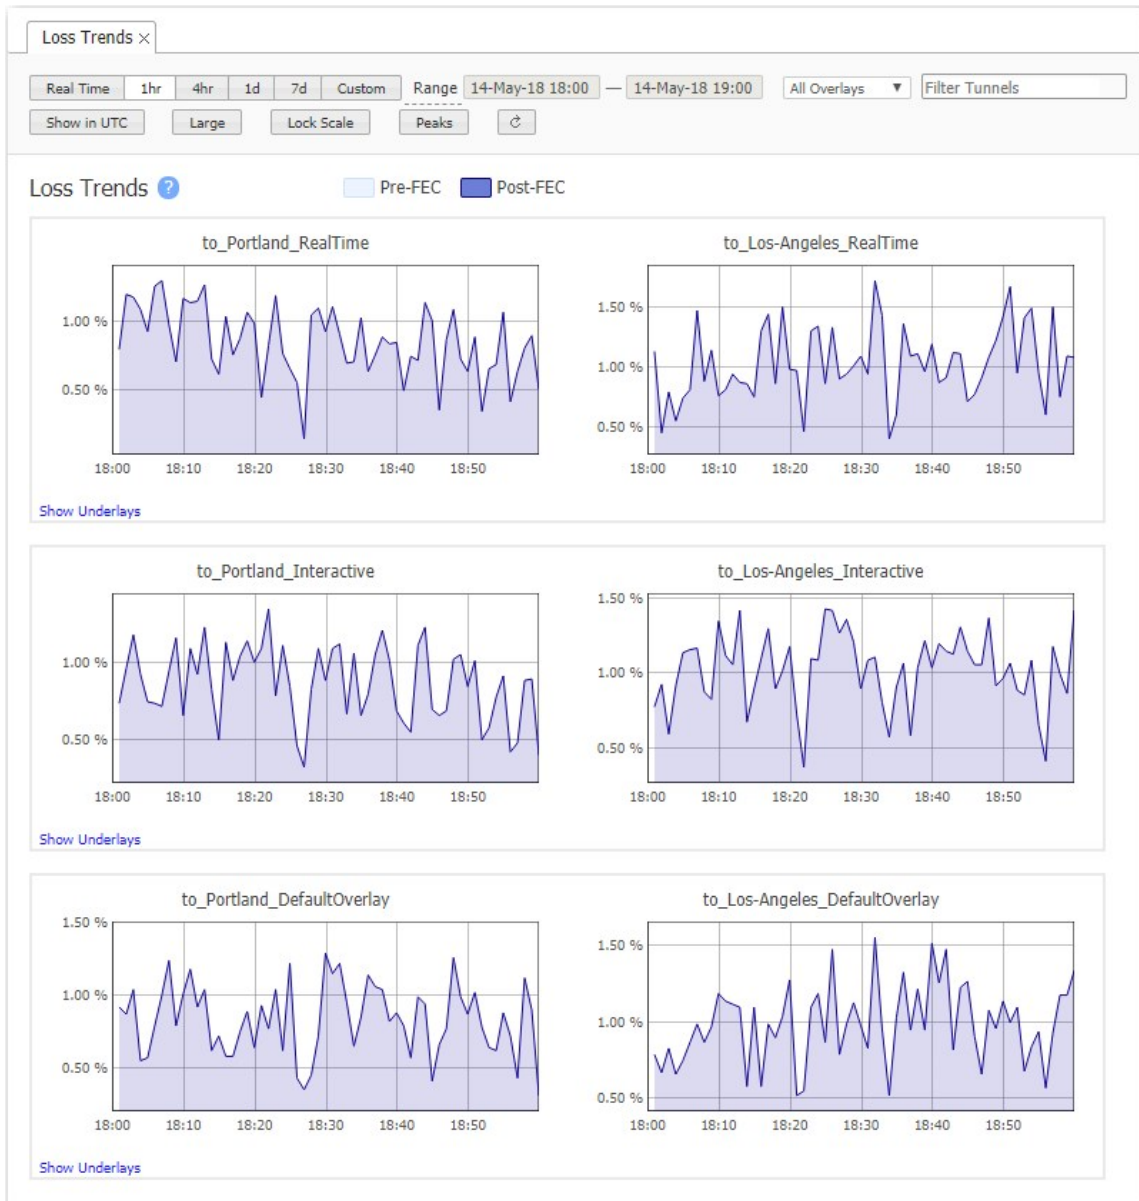
The **Loss** chart shows the tunnels that have the most dropped packets.



Loss Trends

Monitoring > Tunnel Health > Loss > Trends

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

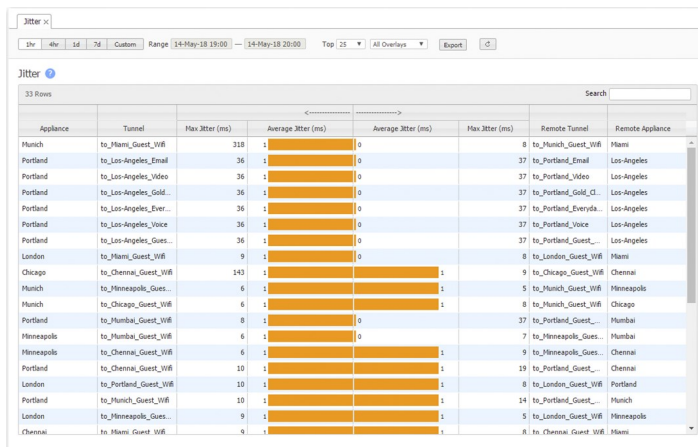


NOTE Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Jitter Summary

Monitoring > Tunnel Health > Jitter > Summary

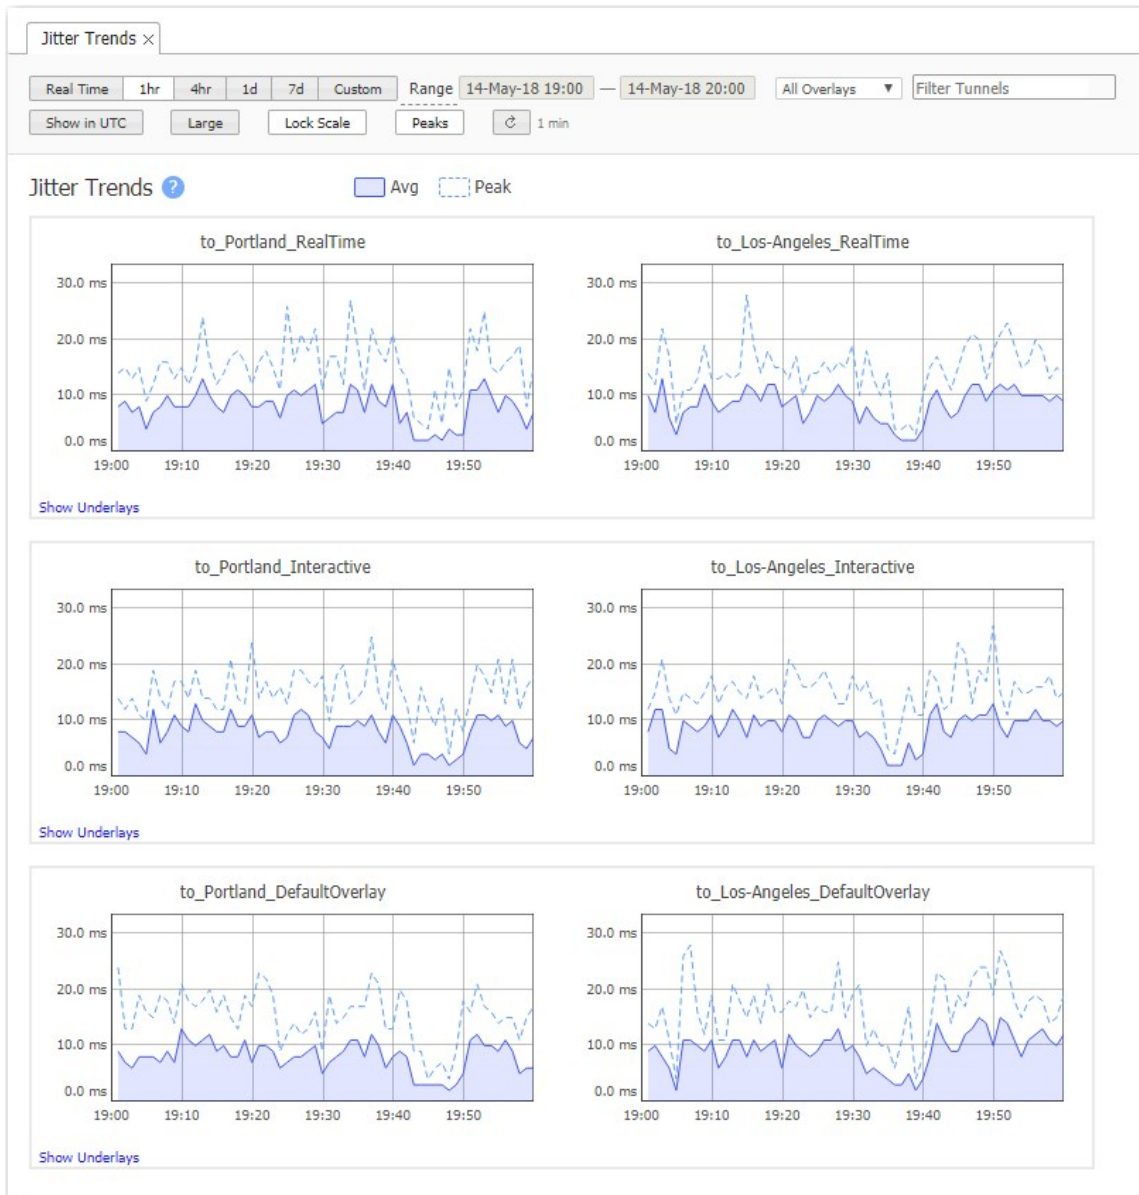
The **Jitter** chart shows the tunnels that have the most jitter. Jitter can be caused by congestion in the LAN, firewall routers, bottleneck access links, load sharing, route flapping, routing table updates, and timing drifts.



Jitter Trends

Monitoring > Tunnel Health > Jitter > Trends

This tab shows tunnel jitter time.

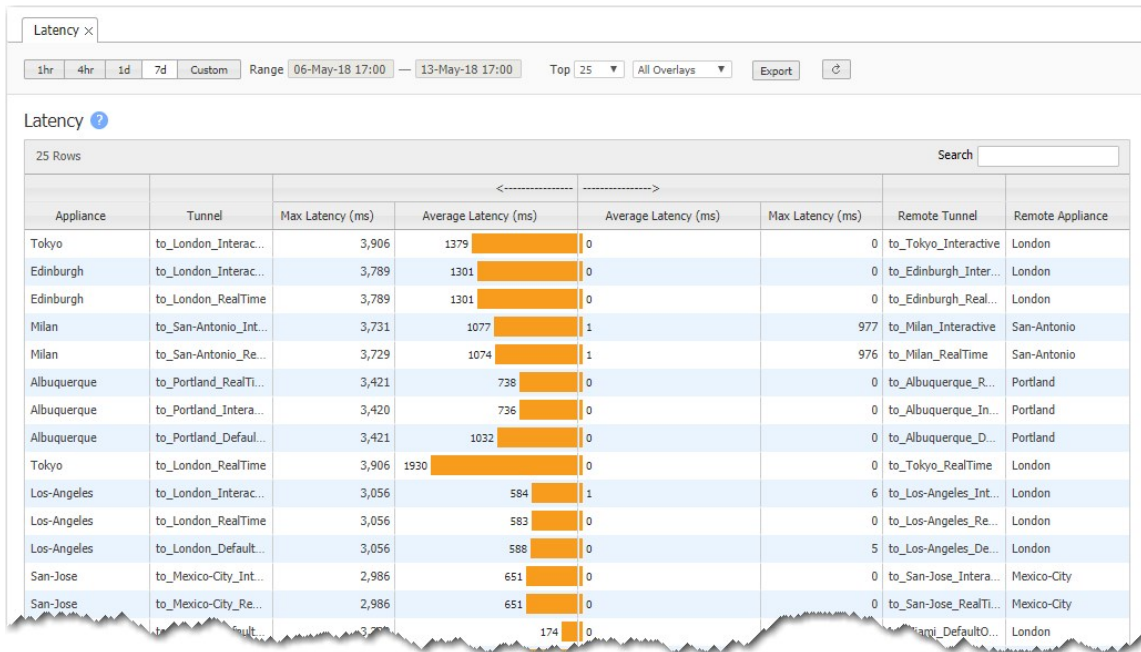


NOTE Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Latency

Monitoring > Tunnel Health > Latency > Summary

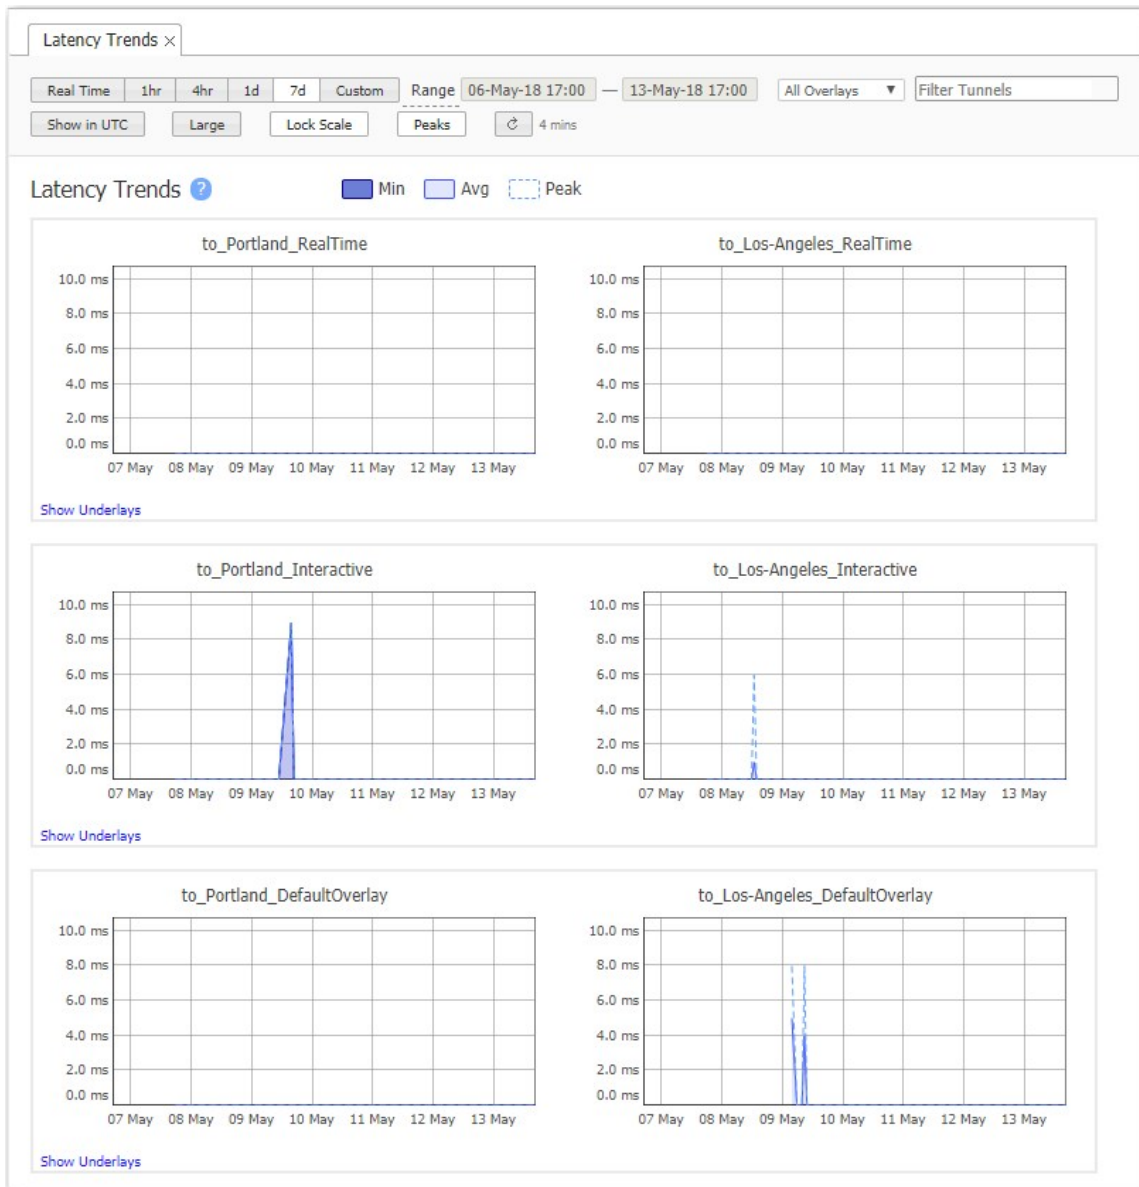
The **Latency** chart shows the tunnels that have the most transmission delay, generally as a result of congestion.



Latency Trends

Monitoring > Tunnel Health > Latency > Trends

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

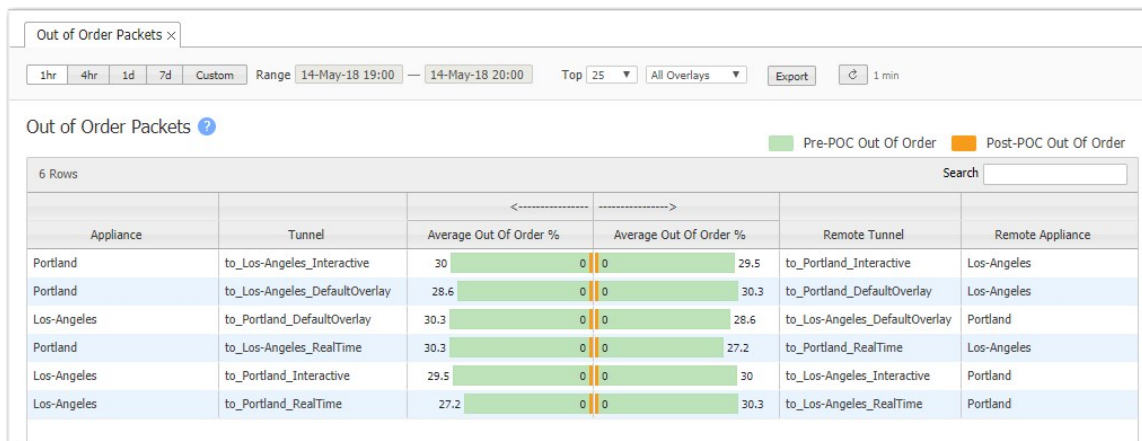


NOTE Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

Out of Order Packets

Monitoring > Tunnel Health > Out of Order Packets > Summary

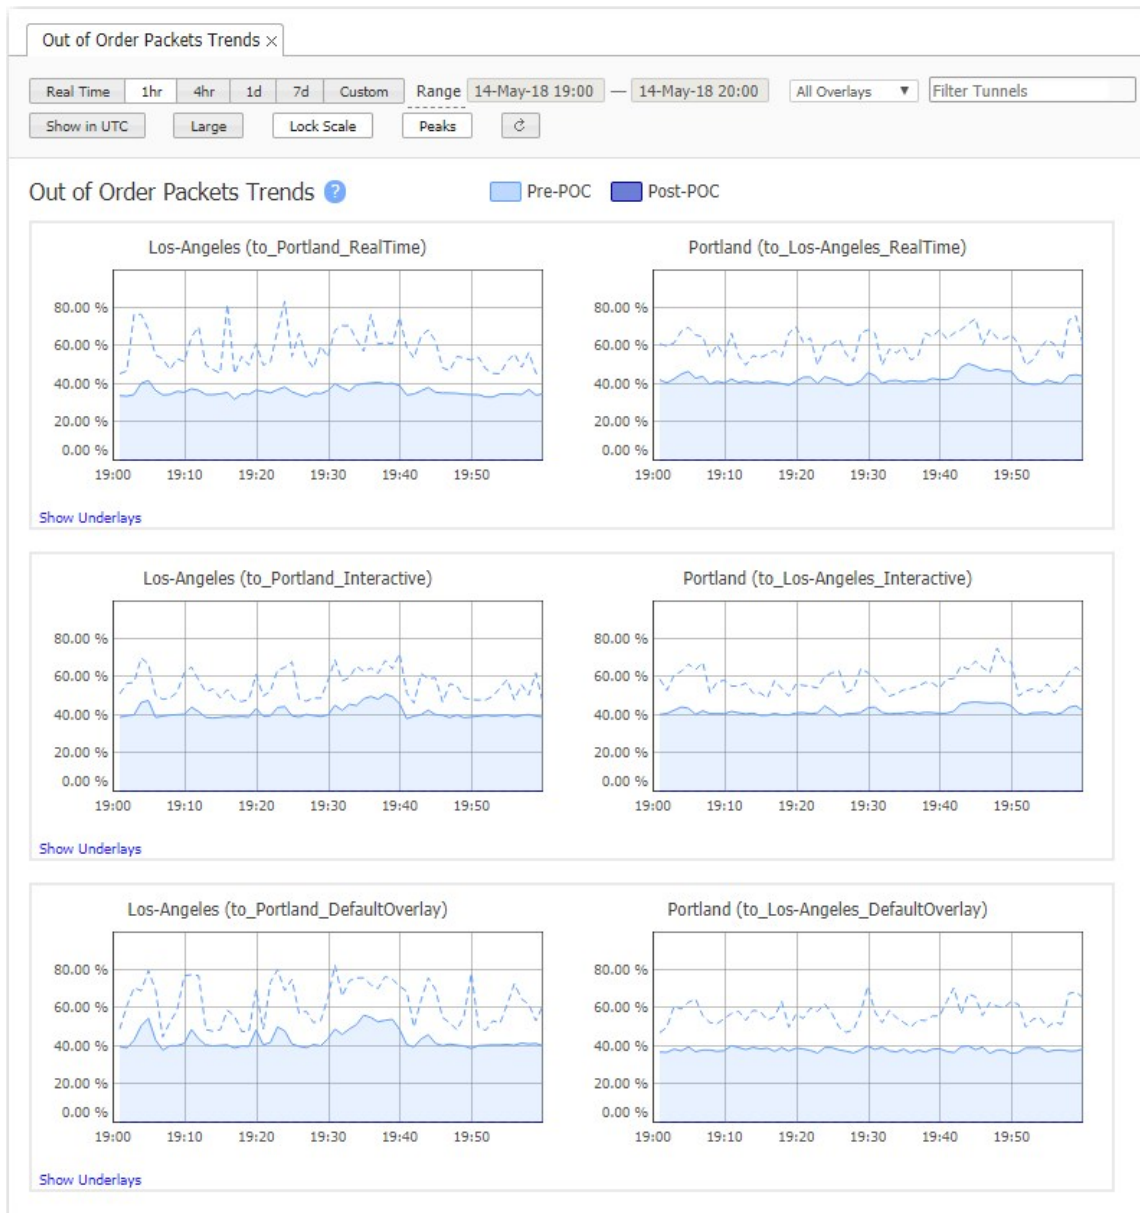
The **Out of Order Packets** chart shows the tunnels that receive the most packets out of sequence relative to how they were sent.



Out of Order Packets Trends

Monitoring > Tunnel Health > Out of Order Packets > Trends

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

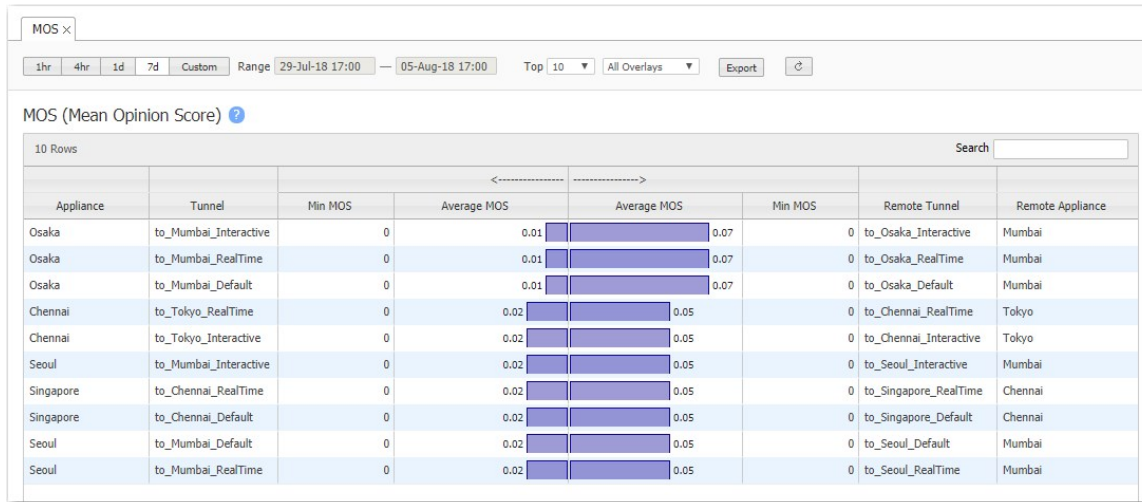


NOTE Underlay tunnels are a shared resource among overlays. Therefore, underlay charts display aggregated data.

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.

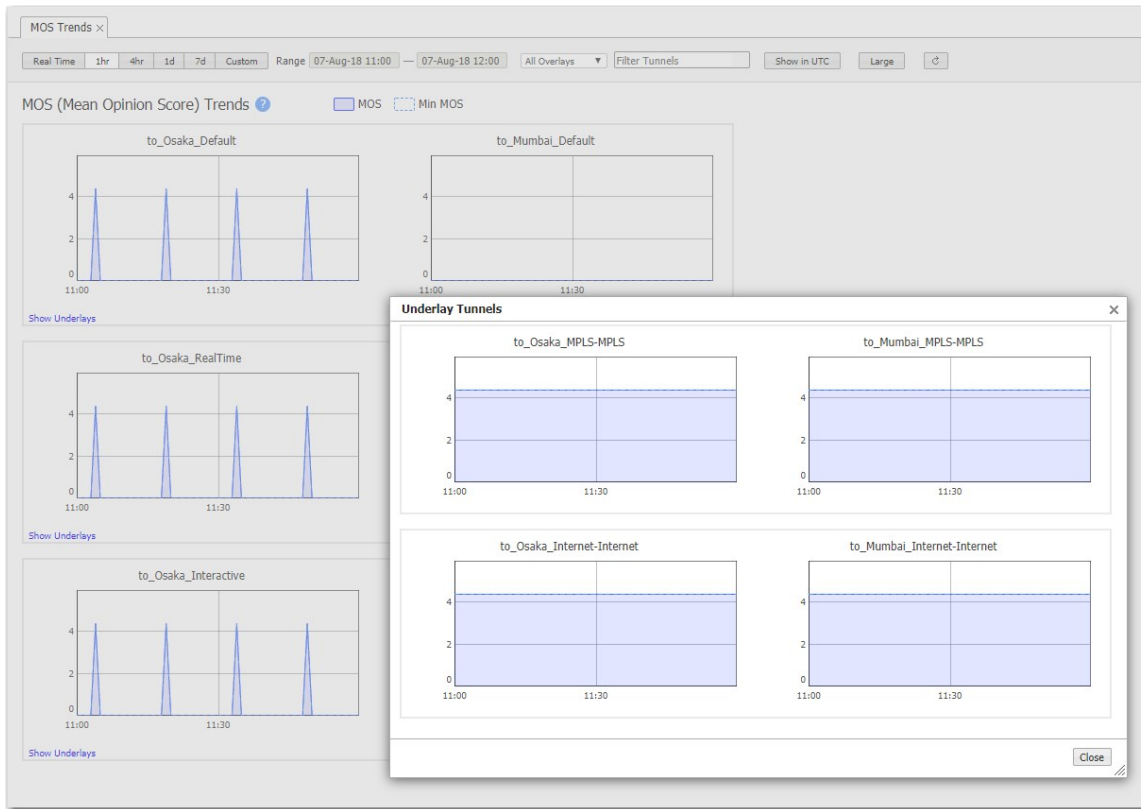


The value can be affected by loss, latency, and jitter. In practice, a value of **4.4** is considered an excellent quality target.

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 > Tunnels Summary

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

Tunnels Summary ×

1hr1d7dCustom

Range14-May-18 19:00--14-May-18 20:00

Top 25 | All Overlays

ExportPayload

7 min

Tunnels Summary ⓘ

50 Rows

Search

Tunnel	Status	Inbound						Outbound						Packets Loss %		Jitter (ms)		Latency (ms)	
		LAN	WAN	Reduction %	LAN Through	WAN Through	LAN W	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 Kbps	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	0.00	26 Mbps	17 Mbps	0	0	72.00	32.00	50.02	0.00		
Portland: to_Hon...	up - active	0	0	0.00	0	0	0	0	0.00	0	0	0	0	53.00	9.00	99.93	0.00		
Honolulu: to_Los...	up - active	0	0	0.00	0	0	0	0	0.00	0	0	0	0	6.00	3.00	49.98	0.00		
Portland: to_Ch...	up - active	0	0	0.00	0	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	0.00	0	0	0	0	5.00	3.00	50.00	0.00		
Portland: to_Ha...	up - active	0	0	0.00	0	0	0	0	0.00	0	0	0	0	51.00	9.00	99.93	0.00		
Honolulu: to_Los...	up - active	0	0	0.00	0	0	0	0	0.00	0	0	0	0	58.00	7.00	99.63	0.00		
Portland: to_Los...	up - active	0	0	0.00	0	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	0.00	0	0	0	0	7.00	17.00	49.97	0.00		
Portland: to_Ch...	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		
Honolulu: to_Los...	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_Ha...	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_Hon...	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_Hon...	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 Kbps	0	0	72.00	32.00	50.02	0.00		
Honolulu: to_Los...	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		
Mumbai: 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_Hon...	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_Hon...	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 specified Link Bonding Policy 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.

Orchestrator Configuration

These topics focus on how to configure Orchestrator. The options available under this menu are organized as follows:

- [Overlays & Security](#)
- [Networking](#)
- [Templates & Policies](#)
- [Cloud Services](#)

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

Configuration > Overlays & Security > 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 box launches. You also can 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 box 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 ACLs 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 about [Regions](#), refer to the help on 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 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.
- **+Secondary:** Click **+Secondary** to enable secondary interfaces. You can specify when you choose Orchestrator to go to Secondary by selecting **Down** or **Not Meeting Service Levels**.

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 can 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. You also can select custom bonding, which enables you to customize link prioritization and traffic steering policies based on multiple criteria.

Field	Description
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, and so forth.
High Efficiency	For everything else. This option sends load balance information on multiple links, with no FEC or overhead.
Custom	Specify the following: <ul style="list-style-type: none"> ■ FEC Wait Time (in milliseconds) ■ Exclude links: Overlay or Underlay brownout ■ Link Reorder Frequency: Aggressive, Moderate, Conservative ■ Path Conditioning (in percentage) ■ Packet Reorder Wait Time (in milliseconds) ■ Link Selection: Waterfall or Balanced

QoS, Security, and Optimization

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

Field	Description
FW Zone	Select the firewall zone you want to restrict traffic to from an overlay.
Boost	Select True or False 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: Use MPLS, Use Internet, Use LTE, Use Best Route, Drop.
Traffic Class	Channels traffic to the desired queue based on the applied service. Select Best Route 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 and 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. Click **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 also can 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.

Break Out Locally Using These Interfaces, Available Interfaces, and Link Selection

You can select the best internet breakout links by specifying the type of **Link Selection: Waterfall** or **Balance**. Drag and drop an available interface into **Primary** or **Backup** in the **Break Out Locally Using These Interfaces** and complete the following steps.

1. Select **Waterfall** or **Balanced** under **Link Selection**.
2. If waterfall is chosen, links are ranked on the selected threshold, from best to worst. The best link is chosen first and the next best link is chosen when the current, best link's bandwidth utilization is full. Select one of the following ways you want Orchestrator to first determine which link to use.

Field	Description
Auto	Default threshold if you do not specify the threshold for your links.
MOS	Measure of the voice connect quality.
Loss	Configured amount of loss the primary link is given.
Latency	Configured amount of time you assign to the primary link for latency.
Fixed Order	Always selects the primary interfaces set by the user. No detection of the best set link.

NOTE Backup links are used only when **all** primary links are down.

3. If **Balanced** is chosen, enter the amount for the three Performance Thresholds: **Loss**, **Latency**, and **Jitter**. Traffic is dispersed between one or more of the configured top or equally ranked links.

WARNING Random links are selected if no brownout thresholds for Loss, Latency, and Jitter have been set.

4. Click the edit icon next to **Break Out Locally Using** these interfaces and complete the dialog box if you choose to set IP SLA Rule destinations.

NOTE You can still enable Path Loading even if you do not select any primary links.

If you select **Exclude links That Are Below Performance Thresholds**, the selected policy order is applied.

Apply Overlays

Configuration > Overlays & Security > 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 also can view the status of the overlays if you select **View Status**.

Interface Labels

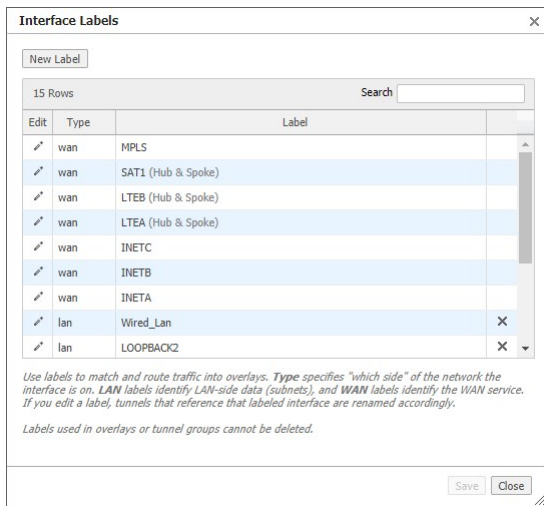
Configuration > Overlays & Security > Interface Labels

To make it easier to identify connections, you can create descriptive interface labels for each link type in your environment. Use labels to match and route traffic into overlays. The label 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, such as MPLS, Internet, or LTE. If you edit a label, tunnels that reference that labeled interface are renamed accordingly.

- LAN labels can be selected for a traffic access policy in a Business Intent Overlay (BIO), which in turn is applied to an appliance with those LAN labels. All traffic matching those interfaces is automatically processed by that BIO. If you use an ACL for a traffic access policy, the LAN label is ignored for that BIO.
- WAN labels are used by Orchestrator and BIOs to determine which interfaces on different appliances should be connected by tunnels built by Orchestrator. Orchestrator automatically pushes interface labels to appliances it manages.

Manage Labels

Use the Interface Labels dialog box to manage labels in Orchestrator, available under **Configuration > Overlays & Security > Interface Labels**.




From this dialog box, you can create, edit, or delete labels.

Create a Label

1. Click **New Label**.

The Interface Label Configuration dialog box opens.



2. Select **wan** or **lan** for the label type.
3. Enter a descriptive name in the **Label Name** field.

NOTE For WAN labels, if you want to allow Orchestrator to build tunnels using this label in any topology, leave the Topology selection set to **any**. If you want to override BIO settings and exclude this label in Full Mesh overlays, set Topology to **Hub & Spoke**.

4. Click **Done** to save your changes and close the dialog box. Otherwise, click **Close** to cancel and return to the list of interface labels.

Edit a Label

1. In the Interface Labels dialog box, click the edit icon to the right of an existing label.
2. Select **wan** or **lan** for the label type—you cannot change the label type if the label is currently in use.
3. If you want to change the label name, modify it in the **Label Name** field.

NOTE For WAN labels, if you want to allow Orchestrator to build tunnels using this label in any topology, leave the Topology selection set to **any**. If you want to override BIO settings and exclude this label in Full Mesh overlays, set Topology to **Hub & Spoke**.

4. Click **Done** to save your changes and close the dialog box. Otherwise, click **Close** to cancel and return to the list of interface labels.

Delete a Label

1. In the Interface Labels dialog box, click the **X** icon to the left of a label you want to delete.

NOTE Labels used in overlays cannot be deleted.

The label is deleted from the list but can be restored by closing the dialog box without saving.

2. To save your changes and permanently delete the label, click **Save**.

WARNING When deleting a label, a confirmation message warns you that deleted interface labels will be removed from all policies, interfaces, and deployment profiles that are currently using the label.

3. Click **Save** to confirm the removal. Otherwise, click **Cancel** to return to the Interface Labels dialog box.

Hubs

Configuration > Overlays & Security > Hubs

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

You can specify whether a hub will re-advertise routes that were previously received from a spoke in the hub's region or a hub in another region.

NOTE This feature requires appliance software version 9.1.0 or later.

You also can access the Regions tab and Business Intent Overlay tabs by clicking the links at the top of the page.

Complete the following steps to add a hub:

1. Start typing a name or select the appliance you want make a hub from the list.
2. Select one of the following:
 - **Re-Advertise Routes** – This hub will re-advertise its routes so that other appliances can learn them. This hub also will re-advertise routes learned from other EdgeConnects within its region.
 - **Do Not Re-Advertise Routes (Stub Hub)** – This hub will not re-advertise routes learned from other regions or spokes within the current region. All local routes (static, directly connected, BGP, and OSPF) will still be advertised. Hubs that do not re-advertise their routes are Stub hubs.
3. Click **Add Hub**.

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

NOTE You must remove all overlays before you can revert a hub back to a spoke.

Deployment Profiles

Configuration > Overlays & Security > 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 DHCP Server Defaults tab (**Configuration > Networking > DHCP Server Defaults**) before creating Deployment Profiles.

You can use Deployment Profiles to simplify provisioning, regardless of whether you choose to create and use **Business Intent Overlays**.

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

Map 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**.
 - Navigate to **Configuration > Overlays & Security > Interface Labels**.
- If you edit a label, the 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 > Networking > DHCP Server Defaults** and pre-populate this page. The other choices are **No DHCP** and having the appliance act as a **DHCP/BOOTP Relay**.
- Enter the LAN interface from the drop-down. Click **+IP** to add a specific IP address.
- Enter the IP address of the specific LAN interface above the **NO DHCP** link.
- The firewall zones you have already configured will be in the list under **FW Zone**. Select the Firewall Zone you want to apply to the LAN you are deploying.

NOTE You can only change the segment being applied for the LAN interfaces.

WAN-side Configuration

- Select the WAN-side label you want to apply to this deployment. Click the edit icon to add a new interface or delete a previously configured interface.
- **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 Firewall 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 an EdgeConnect appliance.
- For traffic outbound to the WAN, the appliance **only** allows IPSec tunnel packets and management traffic that terminate on an EdgeConnect 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 is the same as max system bandwidth.
- To enter values for shaping inbound traffic (optional), you must first select **Shape Inbound Traffic**.

EdgeConnect Licensing: Only visible on EdgeConnect appliances.

- For additional bandwidth, you can purchase **Plus**, and then select it here for this profile.
- If you have purchased a reserve of **Boost** for your network, you can allocate a portion of it in a Deployment Profile. You also can direct allocations to specific types of traffic in the Business Intent Overlays.
- To view how you have distributed **Plus** and **Boost**, navigate to the **Configuration > Overlays & Security > Licensing > Licenses** tab.
- Select the appropriate licensing you have applied to your EdgeConnect 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 EdgeConnect 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 to 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 also must be configured on the directly connected switch/router. Refer to the Silver Peak user documentation.

Descriptions

DHCP Server

Field	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.
DHCP failover	Enables DHCP failover. To set up DHCP failover, click the Failover Settings link.
DHCP Pool Subnet/Mask	Enter the DHCP pool subnet and mask IP addresses.
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.

Field	Description
NetBIOS node type	<p>NetBIOS node type of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types:</p> <ul style="list-style-type: none"> ■ B-node – 0x01 Broadcast ■ P-node – 0x02 Peer (WINS only) ■ M-node – 0x04 Mixed (broadcast, then WINS) ■ H-node – 0x08 Hybrid (WINS, then broadcast)
NTP server(s)	Specifies the associated Network Time Protocol server(s).
Subnet Mask	Mask that specifies the default number of IP addresses reserved for any subnet. For example, entering 24 reserves 256 IP addresses.

DHCP/BOOTP Relay

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

A More Comprehensive Guide to Basic Deployments

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

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

For detailed deployment examples, refer to the Silver Peak website for various deployment guides.

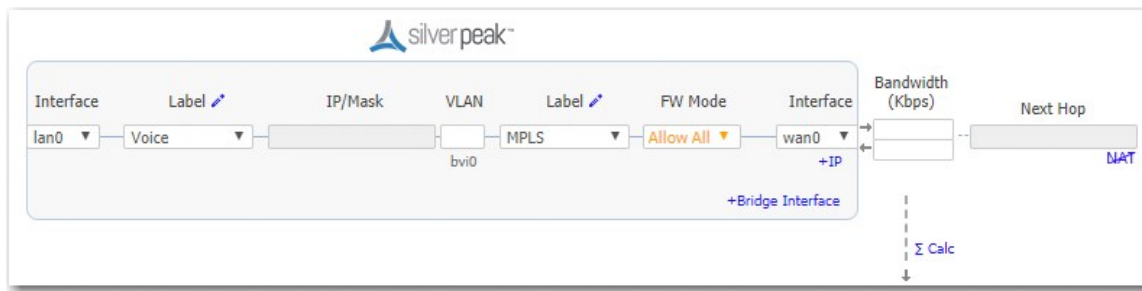
In Bridge Mode and in Router Mode, you can provide security on any WAN-side interface by **hardening the interface**. This means:

- 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

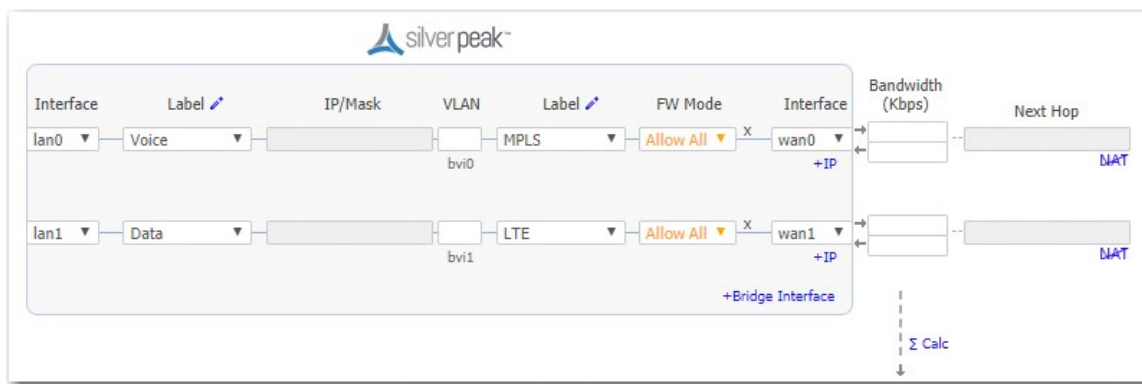
Single WAN-side Router

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, and so forth)

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, you need to configure VLANs on the EdgeConnect appliance 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 Silver Peak 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.

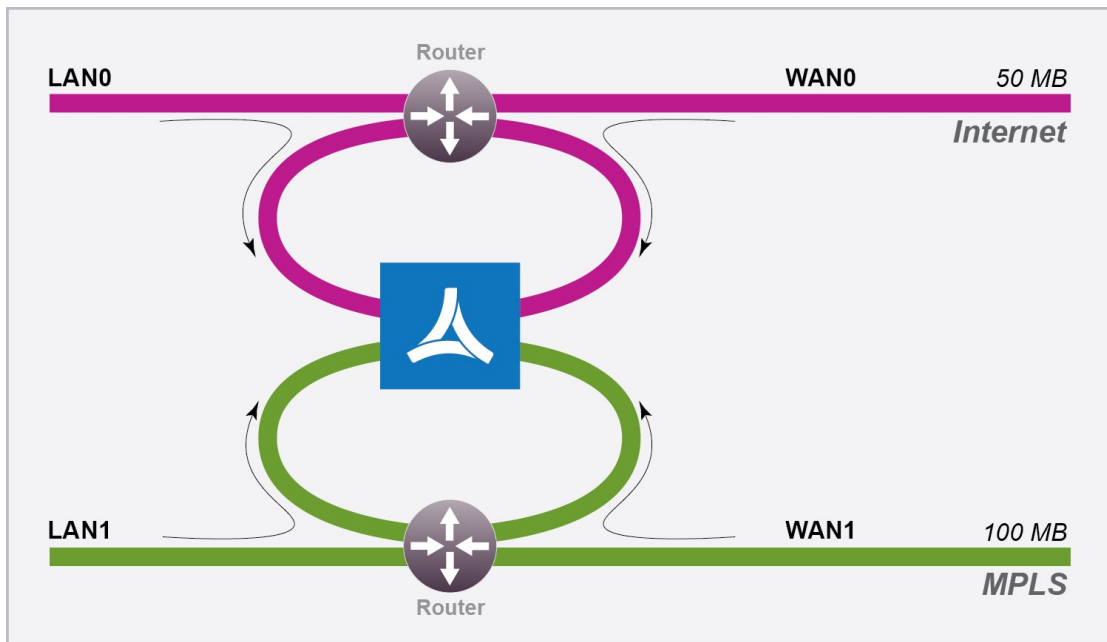
#2 - Dual LAN Interfaces & Dual WAN Interfaces



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

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


Out-of-path dual LAN and dual WAN interfaces



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 Silver Peak 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 EdgeConnect 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 EdgeConnect 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, and so forth).
- 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 might 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.

Examine 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, you also need to use VRRP (Virtual Router Redundancy Protocol).

- **WCCP** (Web Cache Communication Protocol) – Configured on both the router and the EdgeConnect appliance. You also can 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 EdgeConnect 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 EdgeConnect appliances) or **route policies** (which specify destination IP addresses), 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), 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**, the initial **TCP-based** or **IP-based** handshaking creates the tunnel. This 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 > Networking > Tunnels > Tunnels** page.

Server Mode

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

The screenshot shows a configuration form for a Silver Peak device. It includes the following fields and elements:

- IP/Mask:** An empty text input field.
- Label:** A dropdown menu currently set to "None".
- Interface:** A dropdown menu currently set to "mgmt0".
- Bandwidth (Kbps):** Two empty text input fields, one for the minimum and one for the maximum bandwidth.
- Next Hop:** An empty text input field.
- NAT:** A checkbox that is currently checked.
- Σ Calc:** A dashed arrow pointing from the Bandwidth fields to the Next Hop field, indicating a calculation.

ADD 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 High Availability (HA) 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 EdgeConnect devices in the EdgeConnect HA cluster to be able to build Business Intent Overlays using that transport resource.

Enable EdgeConnect HA Mode

1. In the appliance tree, 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** check box.
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, 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.

See Tunnel Settings in the Orchestrator (**Orchestrator > Orchestrator Server > Tools > Tunnel Settings**).

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.

For more information, refer to the IP SLA configuration guide.

Firewall Zones

Configuration > Overlays & Security > Security > Firewall Zones

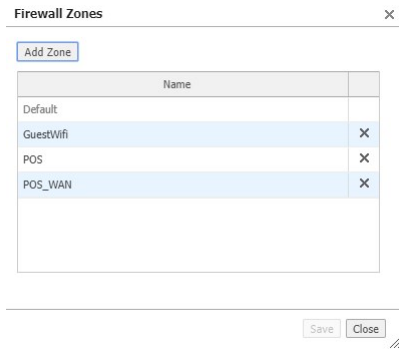
Zone-based firewalls are created on the Orchestrator.

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.



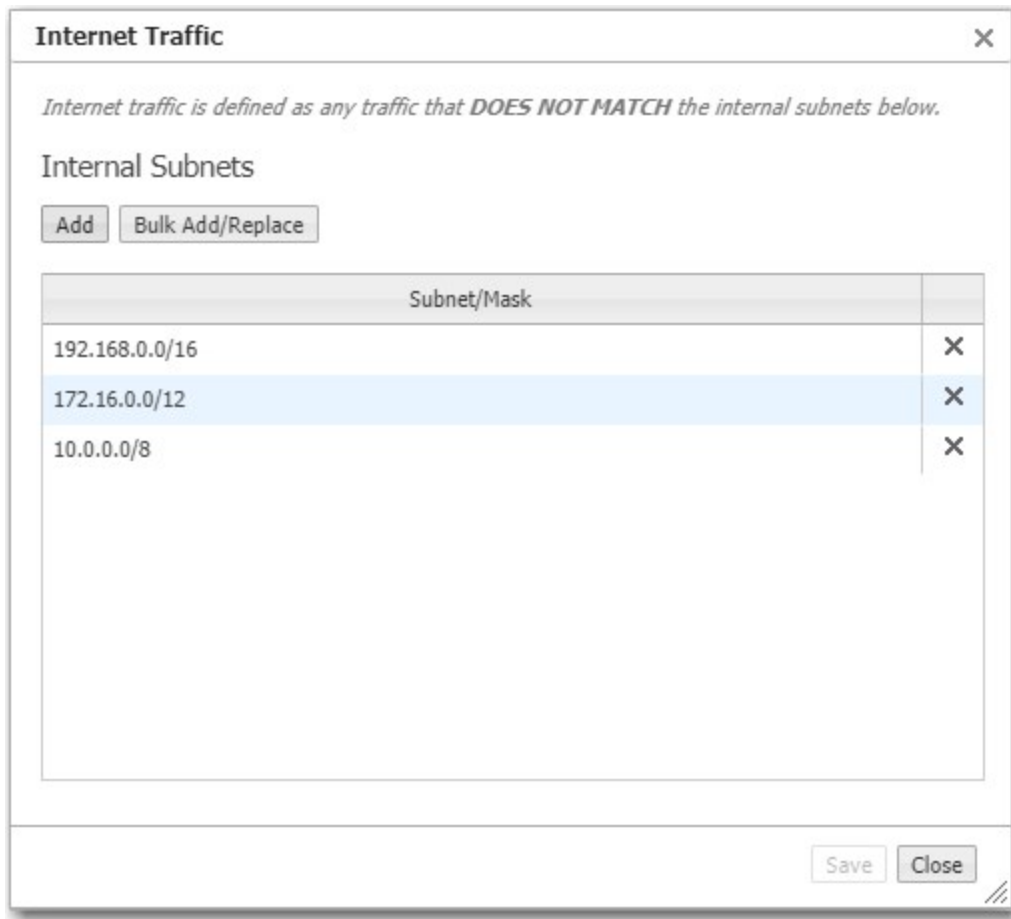
NOTE "Default" will always be the initial default zone. You cannot have another zone named "Default".

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

Internet Traffic

Configuration > Overlays & Security > Internet Traffic Definition

Internet traffic is any traffic that **does NOT match** the internal subnets listed on this dialog box.



IPSec Pre-Shared Key Rotation

Configuration > Overlays & Security > Security > IPSec Key Rotation

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

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 do not 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 will not activate the new key material until all active appliances are reachable and have received the new key material or if the

maximum activation wait time has exceeded. 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.

Schedule IPSec Key Rotation Dialog Box

The Schedule IPSec Key Rotation dialog box enables you to schedule your key rotation. The following tables provide details about the two sections on this dialog box.

SD-WAN IPSec UDP Key Material Rotation Section

Field	Description
Enable Key Rotation	Select this check box to enable key rotation.
Persist Key Material	If enabled, key material is stored on each appliance, ensuring data plane tunnels are built quickly after an appliance reboot (no dependency on Orchestrator). If disabled, new key material from Orchestrator is required after any reboot (Orchestrator reachability is critical).
Max Activation Wait	Maximum time Orchestrator waits before activating the new key material when there are appliances that are not reachable. After this set time, the new key material is activated on all reachable appliances.
Rotation Period	Click the edit icon to set the rotation and the time you want the key material rotation to begin. Click Force Rotate to immediately start a new key material rotation.
Key Material Lifetime	Amount of time a key material lasts. CAUTION The lifetime must be at least three times the amount of the set Rotation Period.

SD-WAN IPSec Pre-shared Key Rotation Section

Field	Description
Enable	Select this check box to enable.
Period	Click the edit icon to set the time when you want the key rotation to begin.

Intrusion Detection System (IDS)

Configuration > Overlays & Security > Security > Intrusion Detection System (IDS)

The Intrusion Detection System (IDS) can monitor traffic for potential threats and malicious activity and generates threat events based on preconfigured rules. Packets are copied and inspected against signatures downloaded to

Orchestrator from Cloud Portal. Orchestrator sends appliances the signature file and any rules that have been added to the allow list. Traffic is designated for inspection using matching rules enabled in the zone-based firewall.

Use the Intrusion Detection System tab to view status or modify the IDS configuration for appliances selected in the appliance tree. The following information is displayed for selected appliances:

Appliance	Name of the appliance.
Status	Indicates whether or not IDS is enabled on the selected appliance.
Events	Click Show Last 100 Events to see the 100 most recent IDS events on the selected appliance.
Stats	Click Show Stats to see the following IDS statistics for the selected appliance: Decoder Packets, Kernel Drops, Alerts Detected, and Decoder Bytes.

Intrusion Detection System ?

Enable IDS on Appliances IDS Allow List

4 Rows Search

Appliance	Status	Events	Stats
Appliance 1	Enabled	Show Last 100 Events	Show Stats
Appliance 2	Enabled	Show Last 100 Events	Show Stats
Appliance 3	Enabled	Show Last 100 Events	Show Stats
Appliance 4	Disabled	Show Last 100 Events	Show Stats

Prerequisites

Note the following requirements about using IDS:

- IDS can be enabled only on appliances with a minimum of four cores and 16 GB of RAM.
- IDS can be enabled only on appliances running ECOS 9.1.0.0 or higher, and appliances running an earlier version of ECOS will not be displayed on the Intrusion Detection System tab.
- IDS is a licensed feature and can be enabled only on appliances that have been assigned the Advanced Security license (see help text on the **Configuration > Overlays & Security > Licensing > Licenses** tab).

NOTE IDS alarms are logged in standard syslog format. You can configure a logging facility for IDS and remote log receiver to send logs to a 3rd party for additional review and analytics (see **Advanced Reporting and Analytics** below).

Enable or Disable IDS on Appliances

Click **Enable IDS on Appliances** to add (enable) or remove (disable) IDS on all appliances displayed in the table.

Enable / Disable IDS

4 Rows Search

IDS ☐ ☒

Appliance	IDS State
Network-1-Proxy	Yes <input type="button" value="Remove"/>
Network-2-Proxy	Yes <input type="button" value="Remove"/>
Network-3-Proxy	Yes <input type="button" value="Remove"/>
Network-4-Proxy	No

1. Select the **Add** check box to enable IDS on the appliances or select the **Remove** check box to disable IDS on the appliances.

The proposed change in state, if any, is displayed for each appliance in the IDS State column.

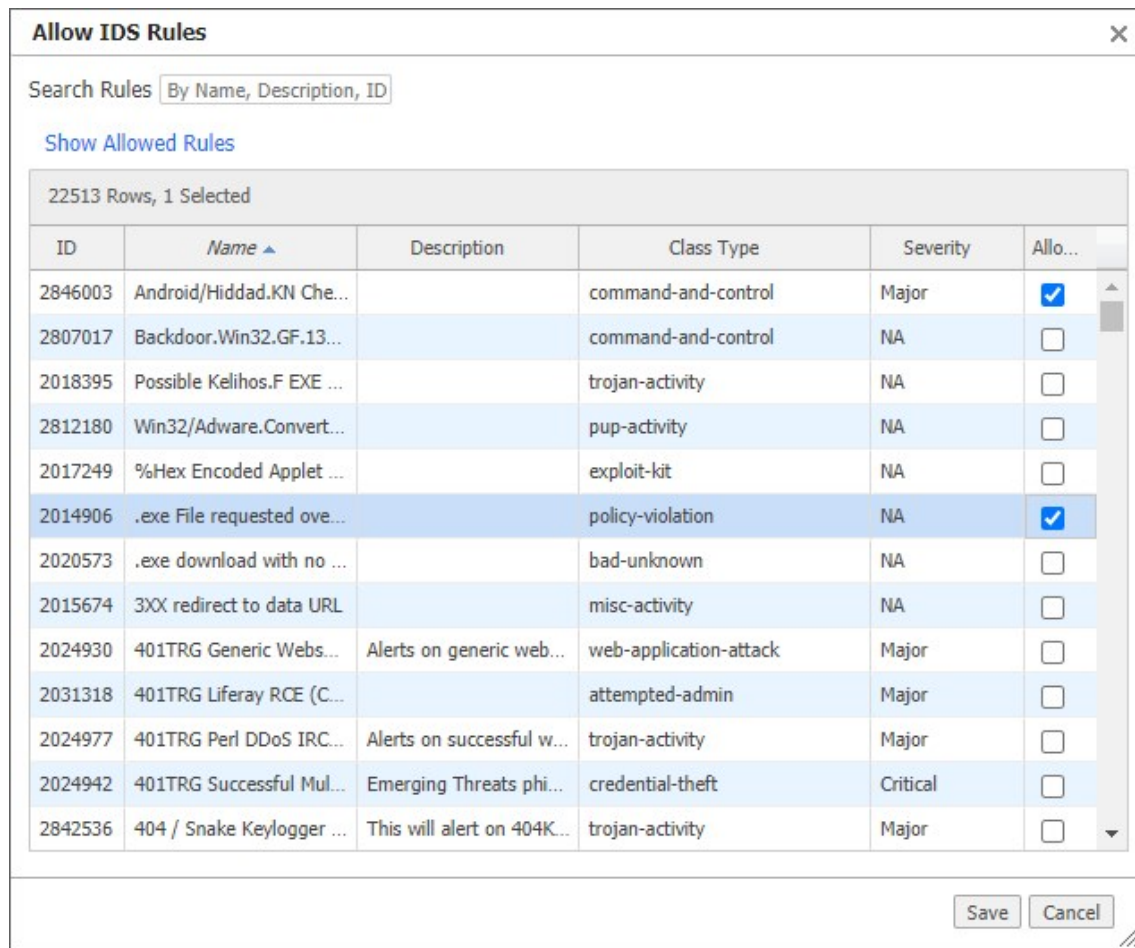
2. Click **Save** to apply your changes or click **Cancel** to close the dialog box without making any changes.

Enable or Disable Rules with the IDS Allow List

By default, all rules included in the IDS signature list are enabled on all appliances where IDS is enabled. For certain traffic or in some specific cases, however, you might want to disable logging and alarms for a rule by adding it to the IDS allow list.

1. To manage which IDS rules are enabled and disabled, click **IDS Allow List**.

The Allow IDS Rules dialog box opens.



- Use the search field at the top of the table to filter the list of rules. You can click **Show Allowed Rules** or **Show All Rules** to display only disabled rules or all rules, respectively.

NOTE If you disable or enable any rules, and then toggle the display between allowed and all rules without saving, your changes will be undone.

- Use the check box in the **Allow** column to disable or enable rules:
 - To disable a rule and add it to the allow list, select the check box.
 - To enable a rule and remove it from the allow list, clear the check box.
- Click **Save** to apply your changes or click **Cancel** to close the dialog box without making any changes.


Specify Traffic to Be Inspected

You can specify the traffic to be inspected according to source and destination zone, as well as specify detailed match criteria, using Firewall Zone Security Policies.

From Zone Default to Zone UNTRUSTED

Source Segment: Destination Segment:

2 Rows, 1 Selected

Priority ▲	Match Criteria	Action	Enabled
20000	Match Everything 	allow ▼	<input checked="" type="checkbox"/>
65535	Match Everything	allow deny inspect	<input checked="" type="checkbox"/>

With the addition of IDS, firewall actions have the following meanings:

- **allow:** Allow traffic and do not inspect
- **deny:** Deny traffic and do not inspect
- **inspect:** Allow traffic and inspect

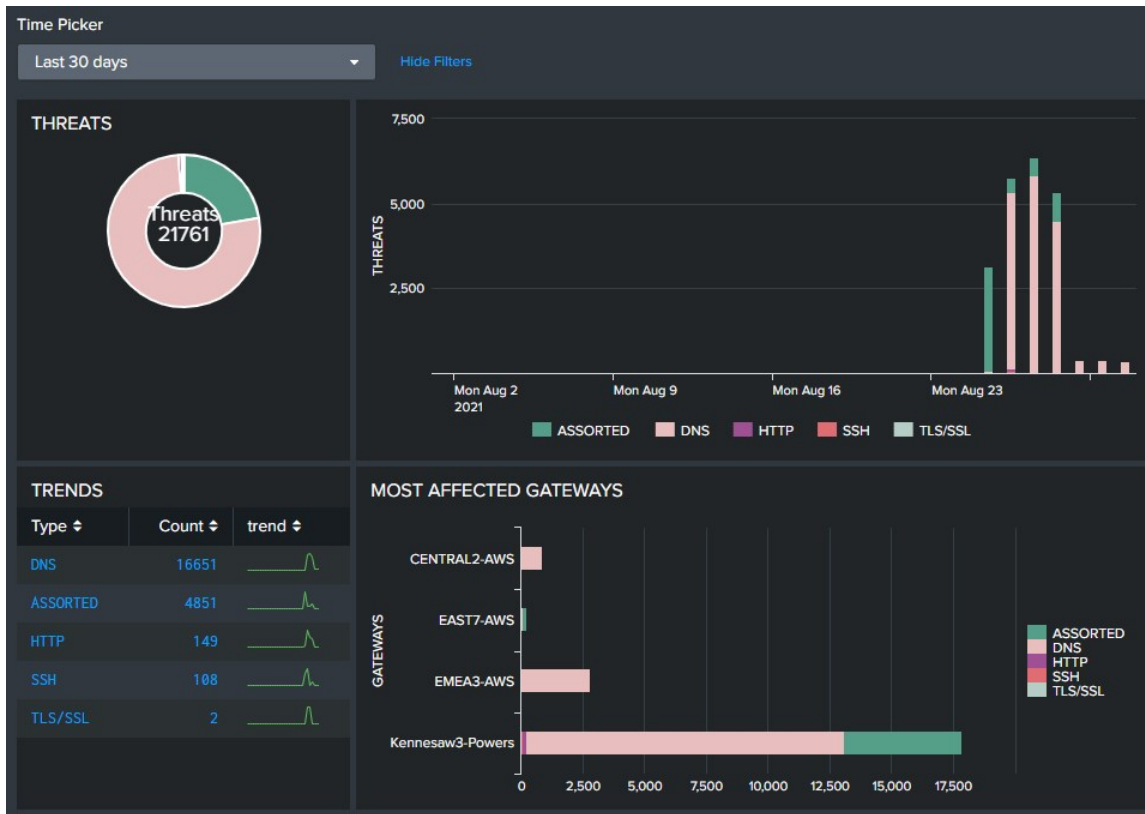
NOTE No traffic will be inspected until rules with the **inspect** action are specified in the security policy.

For more information, see the following tabs in Orchestrator:

- **Templates (Security Policies):** Configuration > Templates & Policies > Templates
- **Routing Segmentation:** Configuration > Networking > Routing > Routing Segmentation (VRF)

Advanced Reporting and Analytics

For users who are using or trying Splunk, you can install the Aruba EdgeConnect app to enable advanced reporting and analytics using the IDS alarms forwarded from EdgeConnect appliances. Search Splunkbase for "EdgeConnect" or click [this link](#) to search in your browser.



Follow the instructions provided to install and configure the app.

SSL Certificates Tab

Configuration > Overlays & Security > 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.

Edit	Appliance Name	Issuer	Issued To	Certificate	Expiration Date
	DM-VX-B	dm	dm	View	Jan 1 23:19:39 2015 G...

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 EdgeConnect 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 also must 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 & Security > 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' management interface. At the top, there are tabs for 'Topology' and 'SSL CA Certificates'. Below the tabs, there's a search bar and a 'Manage SSL CA Certificates with Templates' button. A table lists the certificates. The first entry is 'chateau', issued by 'Silver Peak SSL Proxy', with an expiration date of 'Dec 31 23:59:59 2034 GMT'. A 'View' link is next to the certificate. An arrow points from the 'View' link to a 'View Certificate Content' dialog box. The dialog box displays the certificate details, including version, serial number, issuer, validity dates, subject, and public key information.

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

Showing 1 to 1 of 1 entries

View Certificate Content

```

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:
      4F:BC:F0:A1:FE:AA:
  
```

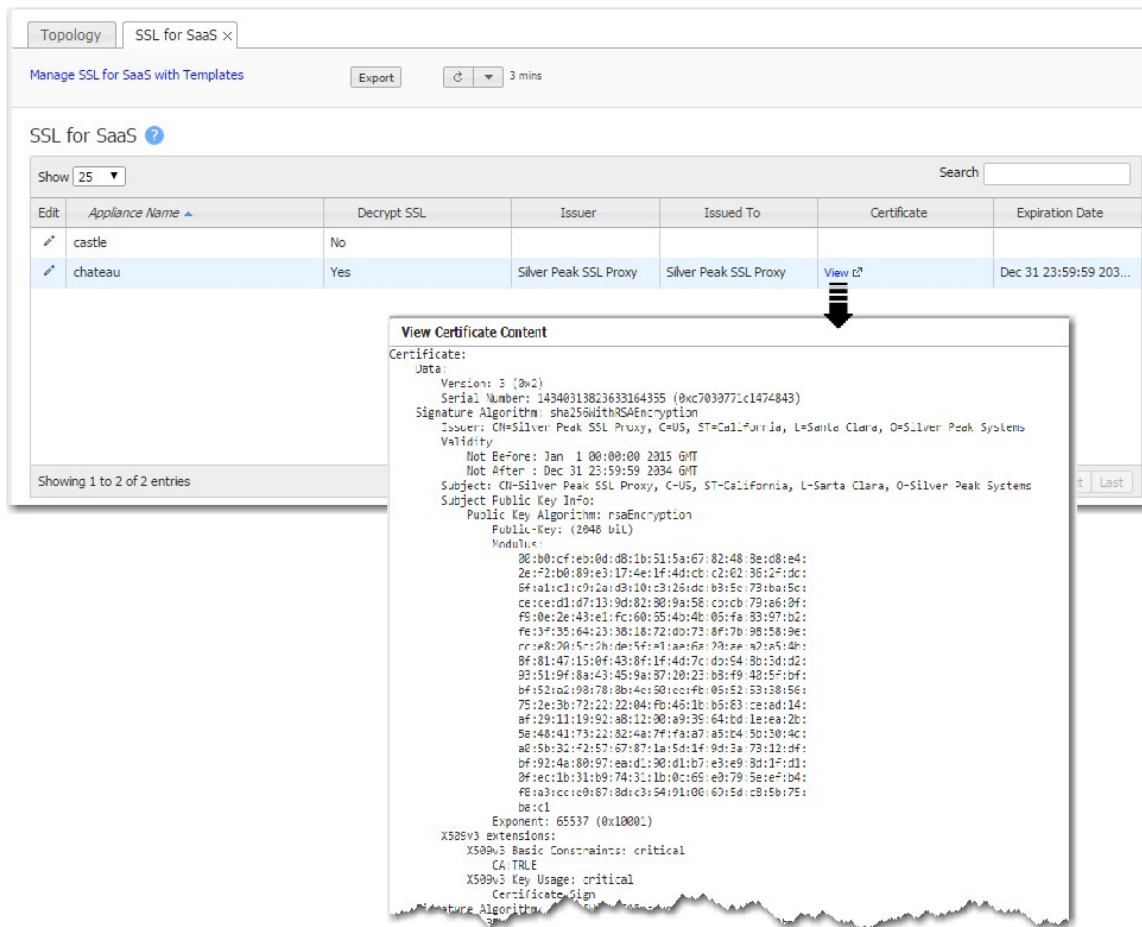
If the enterprise certificate that you used for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, you must add those CA certificates. If the browser cannot validate up the chain to the root CA, it will warn you that it cannot trust the certificate.

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

SSL for SaaS Tab

Configuration > Overlays & Security > SSL > SSL for SaaS

This report lists the signed substitute certificates for the appliances.



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, 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 **do not** already have a subordinate CA certificate, you can access any appliance's **Configuration > Templates & Policies > Applications & SaaS > SaaS Optimization** page and generate a Certificate Signing Request (CSR).

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

Discovered Appliances

Configuration > Overlays & Security > Discovery > Discovered Appliances

This tab lists each appliance that Orchestrator discovers.

Discovered Appliances											
Discovered Devices											
Discovery Email Recipients: <input type="text" value="rsinha@silver-peak.com"/> Save											
Show Denied Devices Show Approved Orchestrators											
Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Reachability	Approve	Deny	Software Version	Model
0100AAB81016	GMS-7.3	128.242.109.226	128.242.109.226	Milpitas, California, US	Unassigned	10-Feb-17 11:07		Approve	Deny	8.0.10.31334	GX-V
001BBC090805	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	11-Jan-17 11:50	Unreachable	Approve	Deny	8.1.4.0_62991	EC-V
001BBC090805	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	11-Jan-17 10:50	Unreachable	Approve	Deny	8.1.4.0_62991	EC-V
001BBC090805	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	10-Jan-17 16:58	Unreachable	Approve	Deny	8.1.4.0_62991	EC-V
0100AAB81016	GMS-7.3	128.242.109.226	128.242.109.226	Milpitas, California, US	Unassigned	06-Jan-17 16:51		Approve	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	Approve	Deny	8.1.4.0_62991	CPX
001BBC090805	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	06-Jan-17 10:29	Unreachable	Approve	Deny	7.3.9.0_62228	EC-V
001BBC090805	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	05-Jan-17 12:55	Unreachable	Approve	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	Approve	Deny	8.1.4.0_62697	EC-V
000C2913A38E	SaaSaaS-VXA	172.25.40.215	128.242.109.226	Milpitas, California, US	Unassigned	21-Oct-16 22:15	Unreachable	Approve	Deny	8.1.4.0_62671	VX-1000
000C29490805	traceroute-vnb	10.0.248.33	128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Reachable	Approve	Deny	8.1.1.0_60893	VX-1000
000C29FC7634	traceroute-vxa	10.0.248.28	128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Reachable	Approve	Deny	8.1.1.0_60893	VX-1000
001BBC091E21	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Unreachable	Approve	Deny		EC-V
000C29790872	rsinha-vnd	10.0.233.134	10.0.233.134	rsinha-vnd	Unassigned	20-Oct-16 11:44	Unreachable	Approve	Deny	8.0.2.0_59017	VX-2000
000C29B0008F	rsinha-vvc	10.0.233.132			Unassigned	19-Oct-16 11:41	Unreachable	Approve	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 does not 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											
Denied Devices											
Discovery Email Recipients: <input type="text" value="rsinha@silver-peak.com"/> Save											
Show Denied Devices											
Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Reachability	Approve	Software Version	Model	Comment
000C290AF241	rsinha-vxa	10.0.233.164			Unassigned	19-Oct-16 11:42	Reachable	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	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	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	Approve	8.0.6.0_61853	EC-XS	Appliance was delet...
001BBC090805	rsinha-ecv	10.0.233.94	128.242.109.226	Milpitas, California, US	Unassigned	11-Jan-17 10:46	Unreachable	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	Approve	8.1.4.0_62991	CPX	Appliance was delet...
000C290D1901	rsinha-vnb	10.0.233.169			Unassigned	19-Oct-16 11:46	Reachable	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

Serial Number	Hostname	IP Address	Public IP Address	Location	Tag	Discovered Time	Deny	Software Version	Model
0030AAB01016	GH5		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Deny		
0030AAB01016	GH5		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Deny	99.99.99.30654	GX-V
0100AAB01016	GH5-7.3		128.242.109.226	Milpitas, California, US	Unassigned	20-Oct-16 11:44	Deny	7.3.10.30880	GX-V

Preconfigure Appliances

Configuration > Overlays & Security > Discovery > Preconfiguration

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

Preconfigure Appliances x

Preconfigure Appliances ⓘ ⓘ 24 mins

New Clone from Existing

1 Rows

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
36  #           country - Country

```

Save Validate Cancel

Field	Description
Name	Assigns a name to the preconfiguration file.
Comment	Optional descriptive field.
Auto Approve when Discovered	When selected , Orchestrator finds the appliance that matches the Discovery Criteria and automatically loads it without needing user intervention. When deselected , the user will be prompted to manually approve the association of the preconfiguration file to the appliance.
Serial	Serial number associated with the appliance that is to receive this configuration.
Appliance Tag	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 & Security > Discovery > Configuration Wizard

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

NOTE Orchestrator assumes you will 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.

Recommended Configuration

Silver Peak recommends configuring the following items before running the Appliance Setup Wizard

Labels	<i>WAN</i> MPLS, Internet, LTE <i>LAN</i> Voice, Data
Deployment Profiles	No Profiles Created
Overlays	
Templates	Default Template Group, my templates
DHCP Pool	DHCP Pool not configured

Continue to Wizard

Appliance Wizard

Appliance Setup

1

2

3

4

Hostname*

Chennai

Group*

Asia

Admin Password

Confirm Password

Software Version

10.10.10.10.10.10

Site Name

Contact Name

Contact Email

Serial Number*

10.10.10.10.10.10

Location

Address 1

Address 2

City

Chennai

State

Zip Code

Country

India

Region

Asia

Map

Map data ©2018 Google Terms of Use Report a map error

< Previous

Next >

Apply

Appliance Setup

1234

Deployment Profile Current Configuration

Router

Bridge

Server

LAN Interfaces +Add

Interface	VLAN	FW Zone	Label	IP/Mask
lan0		Default	None	10.17.17.20/24 No DHCP

WAN Interfaces +Add

Interface	VLAN	FW Zone	FW Mode	Label	IP/Mask
wan0		Default	Allow All		10.17.18.20/24 MPLS
wan1		Default	Stateful+		DHCP10.0.184.154/24 Internet

Bandwidth (Kbps)

Next Hop

100,000

10.17.18.1

100,000

DHCP10.0.184.1

Σ Calc

Total Outbound 200,000 Kbps

Total Inbound 200,000 Kbps

EC Base

Boost 5,000 Kbps

< Previous

Next >

Apply

Appliance Wizard

×

Appliance Setup

1

2

3

4

Add Local Routes ?

The subnet containing the Silver Peak appliance will be automatically shared with other appliances in your network. Add additional subnets for this location below, and they will be shared as well.

☒ Use shared subnet information

☒ Automatically advertise local LAN subnets

☒ Automatically advertise local WAN subnets

Add

Subnet/Mask	Next Hop	Metric	Advertise to Peers	Exclude	

< Previous

Next >

Apply

Appliance Setup

Add Business Intent Overlays to this Site
Overlays build and manage connections between sites, as well as define how traffic is routed and prioritized throughout the network. The Deploy Overlays tab allows you to view and manage overlays on each appliance.

- ☒ RealTime
- ☒ Interactive
- ☒ Default
- ☐ BusinessCritical

Select Template Groups to be applied to this Site
Templates are used to configure appliance settings including: Authentication, SSL Certificates, Threshold Crossing Alerts, DNS, SaaS Optimization and Date/Time.

- ☐ Default Template Group
Access Lists, Shaper, System
- ☐ Demo
Access Lists, Security Policies
- ☐ North America
DNS, Shaper
- ☒ NTP
Date/Time, Logging
- ☒ Security1
Security Policies
- ☐ trial
Access Lists, Security Policies

< Previous Next > Apply

Licenses

Configuration > Overlays & Security > Licensing > Licenses

- This page lists the appliance model, serial number, hostname, EdgeConnect (EC) license term, Boost license term, and feature licenses, and license terms for the appliances selected in the appliance tree.
- You can add, edit, or revoke EC licenses (EC size, Boost, or feature licenses) from an appliance.
- For EC-Enterprise accounts: A license summary including the number of used licenses and total number of available licenses is displayed above the table. The expiration date of the Boost license and each feature license is also listed.

NOTE EdgeConnect stops passing traffic when a license expires.

- Click one of the following buttons to filter the list:

All	Display all appliances
EdgeConnect	Display appliances with EdgeConnect license (configured or granted)
Boost	Display appliances with Boost license (configured or granted)
Feature license	Display appliances with this feature license (configured or granted)

Assign a License to an Appliance

1. In the appliance tree, select one or more appliances to display in the table.
2. Do one of the following:
 - To assign licenses to one appliance, click the **Edit** icon next to that appliance.
 - To assign licenses in bulk (to all appliances in the table), click **Assign Licenses to Appliances**.

NOTE To assign licenses in bulk, all appliances must be on the same software version.

The **Assign Licenses to Appliances** dialog box opens.

3. Complete the following elements as needed:

Field	Description
EC	Select the Add/Replace check box, and then select the EC size from the list: Mini , Base , Base + Plus , 50 Mbps , 200 Mbps , 500 Mbps , 1 Gbps , 2 Gbps , or Unlimited .
Boost	Select the Add/Replace check box, and then enter the amount of Boost to apply to the EC.
Feature licenses	<ol style="list-style-type: none"> 1. To add a feature license, select the Add/Replace check box. 2. If required, select a license option from the list and specify a quantity, such as amount of bandwidth.

4. To revoke a license or Boost, select the **Revoke** check box next to the license or Boost you want to revoke.

NOTE If you revoke an EC license from an appliance, Silver Peak will revoke the Boost license and all feature licenses from that appliance.

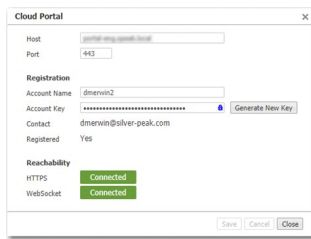
NOTE You must revoke the license from an appliance before you can RMA it. For more information on how to RMA an appliance, see [RMA Wizard](#).

5. Click **Apply**.

Cloud Portal

Configuration > Overlays & Security > 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:** portal.silverpeak.com
- Port:** 443
- Registration:**
 - Account Name:** dmervin2
 - Account Key:** [Redacted] with a 'Generate New Key' button.
 - Contact:** dmervin@silver-peak.com
 - Registered:** Yes
- Reachability:**
 - HTTPS:** Connected
 - WebSocket:** Connected

At the bottom 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 appliances.
- The cloud portal populates the **Contact** field from information included in your purchase order.
- Use of these services requires that your appliances can access the cloud portal via the Internet.

Network Configuration Tabs

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

Deployment Tab

Configuration > Networking > Deployment

This tab provides summary and detailed views of the selected appliance's deployment settings.

To change an appliance's deployment settings, click the **Edit** icon next to the name of the desired appliance.

The following table describes the fields on the Summary view of this tab.

Field	Description
Appliance Name	Name of the deployed appliance.
HA	Name of the appliance with which this appliance is paired for EdgeConnect High Availability (HA).
Mode	Indicates the deployment mode for the appliance: <ul style="list-style-type: none">■ Inline Router – Uses separate LAN and WAN interfaces to route data traffic.■ Bridge – Uses a virtual interface, bvi, created by binding the WAN and LAN interfaces.■ Server – Both management and data traffic use the mgmt0 interface.
Outbound Bandwidth	Deployment's total outbound bandwidth in Kbps.
Inbound Bandwidth	Deployment's total inbound bandwidth in Kbps.
WAN Labels Used	Identify the service, such as MPLS or Internet .
LAN Labels Used	Identify the data, such as data , VoIP , or replication .
Segment	Names of the segments used for this appliance deployment.
Details	Select the information icon to view further deployment details of an appliance.

The following table describes the fields on the Details view of this tab.

Field	Description
Appliance Name	Name of the deployed appliance.
Interface	Name of the LAN or WAN interface.
Label	Label mapped to the interface. 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.

Field	Description
Zone	Firewall zone applied to the interface.
Segment	Name of the segment used for this interface.
IP/Mask	Interface's IP address and subnet mask.
WAN/LAN Side	Indicates that the interface is WAN-side or LAN-side.
Next Hop	Deployment interface's next-hop router address.
Public IP	Public IP address.
Inbound	Interface's inbound bandwidth in Kbps.
Outbound	Interface's outbound bandwidth in Kbps.
NAT	Indicates whether the appliance is behind a NAT-ed interface.
Firewall Mode	Indicates the firewall mode for the appliance's WAN-side interface: <ul style="list-style-type: none"> ■ Allow All – Permits unrestricted communication. ■ Stateful – Only allows communication from the LAN-side to the WAN-side. Used if the interface is behind the WAN edge router. ■ Stateful+SNAT – Applies Source NAT to outgoing traffic. Used 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 an EdgeConnect appliance. For traffic outbound to the WAN, the appliance only allows IPSec tunnel packets and management traffic that terminate on an EdgeConnect appliance.
DHCP	Indicates whether the interface's IP address is obtained from the DHCP server.
HA Interface	Indicates whether the interface is part of an EdgeConnect High Availability (HA) link.
Comment	Additional information for this deployment interface.

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

The following table describes the fields on this tab.

Field	Description
Appliance Name	Name of the appliance of an interface.
Name	Name of the LAN/WAN interface selected.
Status	Status of the interface (up or down.)
IP Address/Mask	IP address.
Public IP	Public IP address.
Segment	Name of the configured segment being used.
DHCP	Whether this interface's IP address is obtained from the DHCP server. Displays as Yes , No , No data (not configured), or Invalid data (error condition).
Speed	Current interface speed.
Duplex	Current interface duplex.
MTU	Maximum number of packets being transmitted.
MAC Addresses	MAC addresses applied to an interface.

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

Terminology

Interface	Description
blan	Bonded LAN interfaces (as in lan0 + lan1).
bvi0	Bridge Virtual Interface. When the appliance is deployed in-line (Bridge mode), it is the routed interface that represents the bridging of wan0 and lan0 .
bwan	Bonded WAN interfaces (as in wan0 + wan1).
tlan	10-Gbps fiber LAN interface.
twan	10-Gbps fiber WAN interface.

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 also can 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	Description
Priority	Order in which the rules are executed; the lower the priority, the higher the chance your NAT rule will be applied.
LAN Interface	Name of the LAN interface the NAT rule is using. This is configurable for an outbound NAT rule only.
Segment	Name of the segment being used.
Direction	Select the direction the traffic is going: <ul style="list-style-type: none"> ■ Outbound (LAN to Fabric) ■ Inbound (Fabric to LAN)
Protocol	Type of protocol being used for each NAT.
Source	Original source IP address of the IP packet.
Destination	Address of the LAN/WAN interface where the traffic is going to.
Translated Source	Translated source IP address when the NAT rule is applied.
Translated Destination	Translated destination IP address when the NAT rule is applied.
Enabled	Select this check box to enable your customized NAT rule. Direction can be both inbound or outbound.
Comment	Any comment you want to add pertaining to your NAT rule.
Criteria	Match: LAN interface, direction, source, destination Set: 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 on 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 about your Pool.

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

VRRP Tab

Configuration > Networking > VRRP

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

In an out-of-path deployment, one method for redirecting traffic to the EdgeConnect appliance is to configure VRRP on a common virtual interface. 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 in which 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**.

WCCP Tab

Configuration > Networking > WCCP

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

Topology		WCCP								
Export										
WCCP										
5 Rows										
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.

INFO Refer to the [Silver Peak Network Deployment Guide](#) and the [SD-WAN Deployment Guide](#) for examples, best practices, and deployment tips.

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 ×

Export Deployment | Interfaces | PPPoE

PPPoE ?

7 Rows Search

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

When configuring a PPPoE connection, complete the following fields:

Field	Description
Ethernet Device	Specifies the 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	Name is ppp followed by a numerical suffix from 0 to 9 .
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 might be asked to configure some of the **Optional Fields**, below.

Field	Description
ACNAME	Access Concentrator Name. Provided by ISP.
Connect Poll	Specifies how many times to try to establish the link. The default value is 2 .
Connect Timeout	When trying to establish the link, this specifies how many seconds until the effort times out. The default value is 30 seconds.
Default Route	If the check box is selected, the connection uses the default gateway provided by the ISP.
DNS Type	This specifies the resolver to use: <ul style="list-style-type: none"> ■ NOCHANGE – Do not accept or configure the ISP's Domain Name Server (DNS). Use the DNS configured on the Administration > General Settings > Setup > DNS tab. ■ SERVER – Accept the ISP's DNS. This then overrides Silver Peak's DNS configuration. ■ SPECIFY – Use DNS1 and DNS2 to resolve domain names.
LCP Failure	Link Control Protocol Failure. Specifies the number of times the keep-alive can fail before the link goes down. The default value is 3 .
LCP Interval	The default value for this keep-alive interval is 20 seconds.
Service Name	Provided by ISP.

Loopback Interfaces

Configuration > Networking > Loopback Interfaces

The loopback feature enhances reliability and security by enabling you to access your network using a single static IP address. If one interface goes down, you can access all interfaces through the single static IP address.

To add a loopback interface to your network:

1. Navigate to **Configuration > Networking > Loopback Interfaces**.

The Loopback tab opens.

2. Click the edit icon next to the appliance to which you want to add a loopback interface.

The Loopback Interfaces dialog box opens.

3. Click **Add**.

The Add Interface dialog box opens.

4. Configure the following elements as needed:

Field	Description
Segment	Name of the segment, if enabled.
Interface	Name of the loopback interface.
IP/Mask	IP address for the loopback interface.
Admin	Select whether the admin status is up or down.
Label	Label of the loopback interface.
Zone	Zone you want to apply to your loopback interface.

5. Click **Add**.

Loopback Orchestration

Configuration > Networking > Loopback Orchestration

You can create a pool of loopback addresses for Orchestrator to automatically create one or more loopback interfaces. You also can assign IP addresses from the pool to each appliance in the network. Complete the following steps to create the range for your loopback interfaces.

1. Select **+Add Loopback Interface**. The **Loopback Interface** window opens.
2. Specify the **Label** from the drop-down menu. This is optional. If no label is selected, "None" is assigned. Additionally, **Label** only displays the LAN side interface labels configured on the **Interface Labels** tab.
3. Specify the firewall zone if you want the loopback interface to be part of a specific firewall zone.

4. Select the management check box if you want the interface to be used by management applications running on the appliance.

NOTE You can only select one loopback interface as management if you configure multiple loopbacks.

5. Click **Add**.

The following table represents the fields for loopback orchestration.

Field	Description
Segment	Associated segment that has loopback orchestration applied.
Label	Label of the LAN interface being used.
Zone	Firewall zone associated with the loopback interface.
Management IP	Loopback interface selected as the management interface.
Loopback Pool	Pool of loopback addresses representing each device.
Allocated / Total	Number of loopback IP addresses allocated from the pool out of the total number of IP addresses in the pool.
Deleted	Number of loopback interfaces deleted.

NOTE You can only delete an interface from an appliance in the Appliance Manager.

Virtual Tunnel Interface

A Virtual Tunnel Interface (VTI) 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.

Click 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. Click **Add**.

The **Add VTI Interface** window appears.

2. Complete the following fields with the appropriate information.

Field	Description
Segment	Name of the segment, if enabled.
Interface	ID of the VTI.
NOTE IDs 20000 through 30000 are reserved for Orchestrator.	
IP/Mask	IP address and subnet mask of the VTI.
Admin	Select whether the interface is up or down.
Passthrough Tunnel	Name of the passthrough tunnel associated with the VTI.
Interface Type	Interface type (lan or wan).
Label	If you want to apply a label to the VTI, select one from the list of those available.
Zone	Select the firewall zone to which the VTI should apply from the drop-down list.


3. When all fields are complete, click **Add**.

DHCP Server Defaults

Configuration > Networking > DHCP Server Defaults

You can reduce your workload by using this tab 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 EdgeConnect appliance acts as a **DHCP Relay** between a DHCP server at a data center and clients needing an IP address.
- On the **Configuration > Overlays & Security > 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	Description
DHCP Pool Subnet/Mask	Enter the DHCP pool subnet and mask IP addresses.
Subnet Mask	Mask that specifies the default number of IP addresses reserved for any subnet. For example, entering 24 reserves 256 IP addresses.
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.
Default lease, Maximum lease	Specify, in hours, how long an interface can keep a DHCP-assigned IP address.
Default gateway	Indicates whether the default gateway is being used.
DNS server(s)	Specifies the associated Domain Name System servers.
NTP server(s)	Specifies the associated Network Time Protocol servers.
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	<p>NetBIOS node type of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types:</p> <ul style="list-style-type: none"> ■ B-node – 0x01 Broadcast ■ P-node – 0x02 Peer (WINS only) ■ M-node – 0x04 Mixed (broadcast, then WINS) ■ H-node – 0x08 Hybrid (WINS, then broadcast)

Field	Description
DHCP failover	Enables DHCP failover. To set up DHCP failover, click the Failover Settings link.

DHCP/BOOTP Relay

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

DHCP Leases

Configuration > Networking > DHCP Leases

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

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

DHCP Failover

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

1. Select the **DHCP Failover** check box to enable the DHCP Failover feature.
2. Select whether you are configuring the failover settings for either the Primary or Secondary server.

3. Configure the remaining settings in the table below.

DHCP Failover Fields

Field	Description
My IP	IP address of the LAN interface.
My Port	Port number of the LAN interface.
Peer IP	IP address of the DHCP peer.
Peer Port	Port number of the DHCP peer.
MLCT	Optional. If selected, the default is 60 minutes. This field cannot be zero.
SPLIT	Optional. If selected, determines which peer (primary/secondary) should process the DHCP requests.
Max Response Delay	Optional. If selected, determines how many seconds the DHCP server can pass without receiving a message from its failover peer before it assumes the connection has failed.
Max Unacked Updates	Tells the remote DHCP server how many BNDUPD messages it can send before it receives a BNDACK from the local system.
Load Balance Max Seconds	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. It only works with clients that correctly implement the secs field.

DHCP Failover State

Configuration > Networking > 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 appliance 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 that 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 State Fields

Field	Description
Appliance Name	Name of the EdgeConnect appliance that is part of the DHCP failover configuration.





Field	Description
Failover Group Name	Failover group name that is the same for all the tagged and untagged interfaces corresponding to one physical interface.
My State	Failover endpoint state of the selected primary appliance. The states are: Normal , Communications-Interrupted , Partner-Down , Recover , Recover-wait , Recover-done .
My State Time	Date and time when the selected appliance's DHCP server entered the specified state in the table.
Partner State	Failover endpoint state of the partner appliance. The states are: Normal , Communications-Interrupted , Partner-Down , Recover , Recover-wait , Recover-done .
Partner State Time	Date and time when the partner appliance entered the specified state in the table.
MCLT	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.

Link Aggregation

Configuration > Networking > Link Aggregation

This tab displays the link aggregation details for all appliances selected in the appliance tree to the left.

Link Aggregation

4 Rows		Search <input type="text"/>		
Edit	Appliance Name ▲	Channel Groups	Interfaces	MTU
	Kennesaw3-Powers	Link aggregation was not configured on the Appliance.		
	Kennesaw4-Powers	bwan0	wan1,wan2,none,none	1500
	Kennesaw5-Powers	Link aggregation was not configured on the Appliance.		
	Kennesaw6-Powers	Link aggregation was not configured on the Appliance.		

Link aggregation combines data from multiple interfaces into a channel group that provides a single high-speed link. Configuring link aggregation also adds failover redundancy to the interfaces in the group.

View Aggregation Details

To view link aggregation details for one or more appliances, select an appliance or group of appliances. The following information is displayed for each of the selected appliances:

Column	Description
Appliance Name	Name of the EdgeConnect appliance.
Channel Groups	<p>If any channel groups have been configured on the selected appliance, the channel group names are listed here.</p> <p>NOTE You can create up to four channel groups per appliance; two each on the LAN side (blan0, blan1) and WAN side (bwana0, bwana1).</p>
Interfaces	<p>Physical or virtual interfaces that are included in the channel group. A channel group can contain two, three, or four interfaces.</p> <p>NOTE You cannot add an interface that has already been deployed on the appliance deployment page, and an interface can only be included in one channel group.</p>
MTU	The MTU size configured for the channel group. The configured MTU will override any existing MTU settings when the channel group is deployed. The default MTU is 1500.

Modify Link Aggregation

To add, change, or delete channel groups on an appliance, click the edit icon to the left of the appliance name.

Link Aggregation - EdgeConnect Powers

Add

1 Rows Search

Edit	Channel Groups ▲	Interfaces	MTU	
	bwana0	wana1, wana2	1500	

Apply Cancel

Add a Channel Group

To add a channel group, follow the steps below:

1. Click **Add** above the table of channel groups. The Add link aggregation dialog box opens.

A screenshot of the 'Add link aggregation' dialog box. The title bar says 'Add link aggregation - *Network > Policies*'. Inside, there are four fields: 'Channel Group (Interface) Name' with a dropdown menu showing 'blan0', 'Interfaces to Be Grouped' with four dropdown menus each showing 'none', and 'MTU' with a text input field showing '1500'. At the bottom right are 'Add' and 'Cancel' buttons.

2. Select a name for the channel group from the list of those available (blan0, blan1, bwan0, bwan1).
3. Select two, three, or four of the available interfaces to be grouped.
4. Specify the MTU to be applied to all interfaces in the group. The default MTU is 1500.
5. Click **Add**.

Modify a Channel Group

To modify an existing channel group, click the edit icon to the left of the group. You cannot modify an existing group name, but you can change the interfaces in the group and the MTU.

Delete a Channel Group

To delete an existing channel group, click the delete icon (X) to the left of the group.

Regions

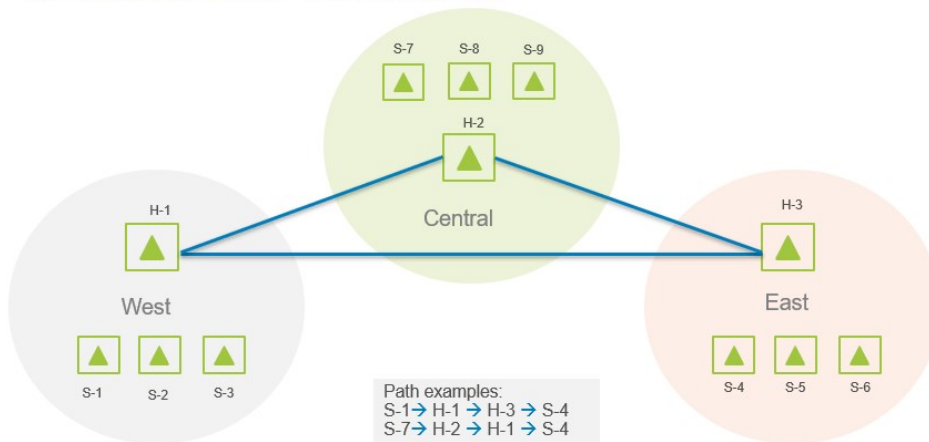
Configuration > Overlays & Security > Regions

Use this tab 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

When enabled, regional routing enables 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 BIOs per region. The following diagrams show examples of different regional network topologies you can build by enabling regional routing.


REGIONAL BIO TOPOLOGY



REGIONAL MULTI-HUB BIO TOPOLOGY



OPTIMIZED REGIONAL BIO TOPOLOGY



You can enable regional routing within your Orchestrator UI. Navigate to the Regional Routing window and click **Enable Regional Routing** in the header and move the toggle.

View Status

Click **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. Click **Edit Regions**.
2. Click **New Region**.
3. Enter the name of your new region in the **Region Configuration** dialog box.
4. Click **Save**.

You also can edit an existing region.

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

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

Routing Segmentation

Configuration > Networking > Routing > Routing Segmentation (VRF)

Use this tab to enable and disable routing segmentation across your network and apply unique configuration to your segments. Routing segmentation allows for the configuration of VRF (Virtual Routing and Forwarding) style

layer-3 segmentation in your SD-WAN deployments. Note the following before configuring routing segmentation in Orchestrator:

- You must upgrade all EdgeConnect appliances and Orchestrator to version 9.0.
- All EdgeConnects must be configured to Inline Router Mode.
- If a new appliance has been added to your network, or if an existing appliance has been replaced, you need to upgrade the appliance software to the appropriate version running in the network.
- After upgrading, segmentation is **disabled** by default. You will have to enable it on this tab.
- Regardless of whether segmentation is enabled or disabled, a **Default** segment is automatically created when you upgrade to 9.0.
- The system-generated Default segment cannot be deleted.
- After you enable routing segmentation, all existing configuration across your network is associated with the Default segment.

Add a New Segment

Before adding a segment, you must enable segmentation by moving the toggle at the top of the page. If Routing Segmentation is not enabled, you cannot make any modifications to the Default segment or add any new segments.

To add a new segment, click **+Add Segment** and enter a **Segment Name**. You can make further specifications by clicking the edit icon or by selecting the **+Add** icon in any of the columns in the table.

Segment Configuration

You can uniquely configure your segments by specifying the following on this page:

- Overlays & Breakout Policies
- Firewall Zone Policies
- Inter-Segment Routing & D-NAT
- Inter-Segment SNAT
- Loopback

NOTE **Inter-Segment Routing & DNAT** and **Inter-Segment Routing & SNAT** are applicable only if you are using different segments.

The following sections provide more details.

Overlays & Breakout Policies for Segments

Use this window to configure overlays and breakout policies for your segments. This configuration determines the overlays used by each segment when traffic is originating from that segment and sent over the SD-WAN fabric to

other sites. This configuration is also used when traffic breaks out locally to the Internet and Cloud Services using the Preferred Policy Order on the **Business Intent Overlay** (BIO) tab. For traffic to match what is on the specified BIO tab, ensure the following two conditions are true:

- BIO must include the defined segment policy
- The BIO match criteria must match the new flow

The overlays are arranged by priority defined in the **Match** field in the **Overlay Configuration** window on the **BIO** page. You can specify if you want to include or skip the segment for each overlay by clicking **Include** or **Skip** icon in the table cell. By default, all overlays are included for all configured segments.

Include and Skip

If you want to skip an overlay, click the enabled **Include** icon and **Skip** appears grayed out. The segment will not be applied to the specified overlay. Click **Skip** again to include the segment; it will turn back to green. If an overlay is set to Skip, traffic will not match that overlay and moves to the next prioritized BIO. Additionally, if no BIOs match, traffic is dropped.

TIP If overlay is set to **Skip**, Flow Details on the **Flows** tab displays the list of skipped overlays.

Firewall Zone Policies

Use this tab to enable and associate firewall zones to your segments. With segmentation enabled, firewall zone security policies are orchestrated and there is no need for Firewall Security Templates. After migration, deactivate the Security Policies Template in all Template Groups. If left active, the template will override any default-default segment security policies configured on this tab.

Before you begin Firewall Zone configuration, note the following:

- Review your existing security policies.
- Create a new security templates group with the new firewall zoning policies that only includes zones associated with LAN and WAN interfaces.
- Delete all rules in your previous Security Policy Template on the **Apply Template Group** tab.
- Ensure you have selected the **Replace** option in the previous Security Policy Template.
- Save the previously used Security Policy Template. This deletes the security policy rules on your appliances.

Complete the following steps to set a rule or policy to your firewall zones within your segment.

1. Select the cell of the segment you want to update in the Matrix View. The **From Zone To Zone** window opens.

NOTE If you are already in Table View, click **Add Rule**.

2. Enter the Source Segment in the **Source Segment** field. This is the segment that the firewall is starting from.

3. Enter the Destination Segment in the **Destination Segment** field. This is the segment where the firewall is going to.
4. Select **Add Rule**.
5. Complete the content in the table.

Field	Description
Priority	Enter the priority amount.
Match Criteria	Click the edit icon in this column to modify and create the match criteria for each zone.
Action	Select Allow or Deny to determine whether this zone will apply the selected segment.
Enable	Select the check box to enable or clear it to disable.
Logging	Determines the filter for the zone-based firewall drop logging levels. You can select one of the following levels to apply: None , Emergency , Alert , Critical , Error , Warning , Notice , Info , or Debug .
Tag	Use tags to categorize or identify the purpose of a rule.
Comment	Any additional details about the firewall zone.

NOTE Firewall zones are unique to each segment. For example, the default zone in Segment X will not be the same default zone in Segment Y.

Inter-Segment Routing & DNAT

Use this tab to configure inter-segment routing and DNAT rules when traffic is crossing between segments. Click **+Add** and the **Inter-Segment Routing & DNAT** window opens. Click **+Add** again and select any rule in the table to modify the following:

Field	Description
Source Segment	Name of the segment traffic is initiating from.
Matches Destination IP	IP address got the source segment. This is used to match the packet destination IP address before the packet goes through DNAT.
Send to Segment	Name of the segment the packets are translated to from the matched destination IP address.
Translated Destination	IP address of the DNAT IP address when the segment is translated.
Enabled	Whether or not this is enabled or disabled within your segment.
Comment	Any additional information.

Inter-Segment Routing & SNAT

This window enables you to enable source network address translation to your segments.

NOTE The default setting for SNAT is enabled for inter-segment traffic.

Field	Description
Source	Name of the segment that the SNAT is starting from.
Destination	Name of the segment that SNAT is translated to.
SNAT	Whether SNAT is enabled or disabled.

Loopback

Click **+Add** and you are redirected to the **Loopback Orchestration** tab. Select the segment you want to apply a loopback interface from the table, and then click **+Add Loopback Interface**.

Appliances

This column represents the amount of appliances the selected segment is enabled on.

Comment

Click the cell in the **Comment** column to add a comment including any additional information for that particular segment.

Deleting a Segment

WARNING Segmentation involves drastic changes to your physical network. Deleting segments can be service affecting. Carefully read this section before deleting any of your segments.

Deleting a segment removes all the segmentation configuration from all the appliances within your network. When you delete a segment, Orchestrator automatically deletes the following:

- The segment's association with the overlay and break-out policies
- The intra-segment and inter-segment firewall zone policies
- The inter-segment routing & D-NAT rules
- The inter-segment S-NAT rule
- The loopback interfaces associated with the segment
- The VTI interfaces associated with the segment
- All the interface and VLAN interfaces

Manual Tasks to Complete Before Deleting a Segment

The following configuration is disassociated from the segment and you need to manually delete the following:

- Any manual created tunnels
- BGP peers in the segment
- Internal subnet table rules
- Overlay ACL rules associated to the deleted segment

To delete a segment, click the **X** in the last column in the table. A Delete Routing Segment warning appears. Click **Delete** or **Cancel**.

Disable a Segment

To disable routing segmentation across your network, you need to delete all configured segments in the network, except the default segment (which cannot be deleted). After all the segments are deleted, navigate to this tab and move the toggle at the top of the page to disable.

Management Services

Configuration > Networking > Routing > Management Services

Use this tab to configure your management services. Management Services can be configured irrespective of whether routing segmentation is enabled or disabled. When routing segmentation is disabled, all the interfaces are available for configuration. When routing segmentation is enabled, Management Services are functional in the associated segment based on the selected interface.

NOTE Management Services will still function if Routing Segmentation is not enabled on your Orchestrator. In this case, you will **only** be able to use the default configuration: **Any** interface with the **Default** segment.

The Management Routes tab works similar to Management Services. The Management Routes tab enables users to configure static routes for management services traffic from the EdgeConnect appliance (egress traffic). This tab, Management Services, enables users to select the source IP address used for each service. It is recommended that you use the management services configuration, but if deterministic source IP address for “Management Services” is not a requirement, you can continue to use Management Routes configuration.

Field	Description
Appliance Name	Name of the appliance selected.
Management Service	Management service being used by your appliance.
Source IP Address	IP address of the interface being used by the selected management service.
Segment	Name of the associated segment that is applied to the management service when your source IP address is selected.

Click the edit icon to get started.

Management Services Dialog Box

Click the **Any** field in the **Source IP Address** column. A drop-down list displays all the interfaces configured on your appliance. When a source IP address is selected, the **Segment** column automatically updates and provides the associated segment.

The following table describes the behaviors of management services.

Service	Behavior
<ul style="list-style-type: none"> ■ HTTP(S) ■ Cloud Portal ■ Orchestrator ■ SaaS Opt 	<p>These services use the selected interface's Interface for the Source IP Address as the source address to establish reachability and WebSocket connections to the Cloud Portal and Orchestrator. HTTP/HTTPS uses the Interface for the Source IP Address for connection as well.</p> <p>If routing segmentation is enabled, SaaS Opt packets are sent from the Interface for the Source IP Address segment interface.</p> <hr/> <p>CAUTION If routing segmentation is enabled, make sure to provide Internet connectivity from the segment to the Interface for source IP Address associated with the segment.</p>
<ul style="list-style-type: none"> ■ NTP ■ NetFlow ■ SNMP ■ SSH ■ Syslog 	<p>Each of these management services uses Interface for the Source IP Address as the source IP address. The source interface configured from the management route table is ignored if the Interface for Source IP Address is not "Any".</p>

Inter-segment DNAT Exceptions

Use this tab to configure inter-segment routing and DNAT rules when traffic is crossing between segments. Click the edit icon and the **Inter-Segment Routing & DNAT** window opens. Click **+Add** and select any rule in the table to modify or define the following:

Field	Description
Appliance Name	Name of the appliance that the DNAT exception is being applied to.
Source Segment	Name of the segment traffic is initiating from.
Matches Destination IP	IP address that matches the destination segment IP address, before DNAT. The IP address is included in the defined policy match criteria.
Send to Segment	Name of the segment the packets are translated to from the matched destination IP address. This is included in the set criteria.

Field	Description
Translated Destination	IP address of the DNAT IP address when the segment is translated. NOTE If DNAT is not needed, this field is empty.
Enabled	Whether inter-segment DNAT is enabled or disabled within your segment.
Comment	Any additional information.

INFO This tab only pushes the inter-segment DNAT exceptions to one appliance, selected in the left toolbar.

Inter-segment SNAT Exceptions

Use this window to enable source network address translation to your segments. Select an appliance or group of appliances from the left menu to apply your SNAT exceptions.

NOTE The default setting for SNAT is enabled for Inter-Segment traffic.

Field	Description
Source	Name of the segment that the SNAT is starting from.
Destination	Name of the segment that SNAT is translated and going to.
SNAT	Whether SNAT is enabled or disabled for the specified segment.

BGP Tab

Configuration > Networking > Routing > BGP

On this tab, you can configure **BGP (Border Gateway Protocol)** for appliances and add their BGP peers (also known as BGP "neighbors"). You also can add and modify peer-based advertisement and redistribution rules. Silver Peak has the following behaviors relative to **communities**:

- Although Silver Peak does not configure BGP communities, it propagates existing communities.
- Appliances can display up to ten 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.
- BGP Equal-cost multi-path (ECMP) is supported for eBGP and iBGP. Multiple next-hops will be installed for the same prefix if all BGP path attributes are the same, enabling BGP to load balance egress traffic across multiple peers. A maximum of 20 BGP peers is supported per appliance, with 16 next-hops supported per interface.

Click the **Summary** button on the BGP tab to display configuration details associated with the local appliance, such as its local AS number and router ID. Click the icon in the **BGP State Details** column to display a summary, including the number of routes learned and advertised via BGP by this appliance.

Click the **Peers** button on the BGP tab to display information about all configured peers for the appliances selected in the appliance tree. Click the icon in the **Peer Details** column to display the connection status of each peer that is configured for the appliance.

The table below describes the fields displayed for the BGP configuration.

Field	Description
Appliance Name	Name of the appliance.
Segment	Name of the segment being used, if enabled.
Peer IP	IP address of the Silver Peak peer.
Local Interface	A list of the interfaces that can be chosen: Any , lan0 , wan0 , or wan1 .
Peer ASN	Peer's Autonomous System Number.
Peer State	State of the peer. A peer state of Established indicates that full adjacency has been established and routes can be advertised to and learned from that peer.
Soft Reset	Allows new changes to be incorporated without taking the entire BGP session down.
Established Time	Final peer state that indicates neighbor connection as complete.
Type	Governs what kinds of routes the appliance is allowed to advertise to this BGP peer. These routes are itemized as Route Export Policies.
Inbound Route Map	Route map being used for the inbound traffic.
Outbound Route Map	Route map being used for the outbound traffic.

Field	Description
Local Preference	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.
MED	<p><i>Multi Exit Discriminator.</i> When BGP chooses 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 is chosen. With MED, the lowest value is preferred.</p> <p>NOTE If you configured the Metric Delta parameter in an earlier version of our software, this value has been translated into a MED value.</p>
Input Metric	Metric that is advertised with the route when shared.
Enable Imports	Allows the learning of routes from this specific BGP peer.
AS Prepend Count	Learned path from an external prepend between a remote BGP site to local BGP peers.
Next-Hop-Self	Advertised route connected to a CE router that an EdgeConnect appliance learns from the eBGP with a PE router.
Keep Alive Timer	Interval, in seconds, between keep alive signals to a peer.
Hold Timer	When availability to a peer is lost, this specifies how long to wait before dropping the session.
Peer Details	Any additional details about a peer or its state.

To edit the BGP configuration for one of the listed appliances, click the edit icon in the left column of the table.

BGP Information

Use this window to enable BGP for your appliances and to configure BGP peers. Complete the following steps to start BGP configuration.

1. Move the toggle to **Enable BGP**.
2. Complete the following fields.

Field	Description
Autonomous System Number (ASN)	Configure this number as needed for your network.

Field	Description
Router ID	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of BGP.
Graceful Restart	<p>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. They will be deleted and replaced with newly received routes.</p> <ul style="list-style-type: none"> ■ Max Restart Time – 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. ■ Stale Path Time – Specifies maximum time (in seconds) following a peer restart that Silver Peak waits before removing stale routes associated with that peer.
AS Path Propagate	Select this check box to enable this appliance to send the full AS path, associated with a prefix to other routers and appliances, avoiding routing loops. This will provide the learned path from an external prepend between a remote BGP site to local BGP peers.

To add a BGP peer, select **Add**. The **Add Peer** dialog box opens.

Add Peer

Complete the following fields to add a BGP peer.

Field	Description
Peer IP	IP address of the Silver Peak peer.
Local Interface	You can specify the source address or interface for a specific BGP peer. Select the interface from the drop-down list: any , lan0 , wan0 , or wan1 .
Peer ASN	Peer's Autonomous System Number.
Peer Type	Select the type of peer from the drop-down list: Branch or PE-router .
Admin Status	Select whether you want the Admin Status UP or DOWN .
Next-Hop-Self	Select this check box to enable the next-hop-self.
Inbound route map	Route map for inbound traffic. Select the edit icon to load or configure inbound route maps.
Outbound route map	Route map for outbound traffic. Select the edit icon to load or configure outbound route maps.
Keep Alive Timer	Interval, in seconds, between keep alive signals to a peer.

Field	Description
Hold Timer	Specified time to wait before dropping the session when the reachability to a peer is lost.
Enable MD5 Password	Select this check box to add a password to authenticate the TCP session with the peer.

BGP Inbound and Outbound Route Redistribution Maps

Route Maps are policies that can be applied to static, OSPF, BGP, and SD-WAN fabric learned routes. These policies have match and set criteria. A route map is applied to routes during route redistribution between routing protocols and allows for filtering routes or modifying route attributes.

- Specify up to 20 BGP route maps (inbound and outbound).
- Apply up to 128 rules per route map.

You can add, delete, rename, or clone route maps using this window. You can add rules to your route map to further specify routing protocols by clicking **Add Rule**. Use rules to allow or deny routes based on numerous matching criteria.

NOTE Prefix match criteria is 'exact match + less than'. Both the prefix specified and any subnets of that prefix will be matched. This behavior will be updated in a future release to allow for selection of 'exact,' 'greater than,' or 'less than' criteria.

To permit a default-route, deny 0.0.0.0/1, deny 128.0.0.0/1, and then permit any.

You can specify the following fields in each rule for the selected route map:

	Inbound
Priority	<ul style="list-style-type: none"> If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from 1000 – 9999 before applying its policies. You can create rules with higher priority than Orchestrator rules (1 – 999) and rules with lower priority (10000 – 19999 and 25000 – 65534). <p>NOTE The priority range from 20000 to 24999 is reserved for Orchestrator.</p> <ul style="list-style-type: none"> When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

Source Protocol

BGP	<p>Enter the prefix (list of subnets separated by commas) and your BGP communities.</p> <ul style="list-style-type: none"> Prefix BGP Communities
-----	---

Outbound**Match Criteria**

Source Protocol	Fields (based on protocol chosen)
Local/Static	<p>Enter the prefix (list of subnets separated by commas) and your BGP communities.</p> <ul style="list-style-type: none"> Prefix
SD-WAN (Local/Static)	<ul style="list-style-type: none"> Prefix BGP Communities
BGP	<ul style="list-style-type: none"> Prefix BGP Communities
OSPF	<ul style="list-style-type: none"> Prefix OSPF Tag
SD-WAN (BGP)	<ul style="list-style-type: none"> Prefix BGP Communities
SD-WAN (OSPF)	<ul style="list-style-type: none"> Prefix OSPF Tag

Inbound and Outbound**Set Actions**

Permit	Enable or disable.
BGP Local Preference	Best BGP destination. The default value is 100.
Metric	Metric for the route.

BGP Communities	Label of extra information that is added to one or more prefixes advertised to BGP neighbors.
Nexthop	Advertised route connected to a CE router that an EdgeConnect appliance learns from the eBGP with a PE router.
ASN Prepend Count	Original route path that was used.

The following table describes the redistribution commands supported in the BGP routing protocol.

Command	Redistribution Support
Match prefix	Yes
Set metric	Yes
Set tag	Yes

Routes Tab

Configuration > Networking > Routing > Routes

Each appliance builds a route table with entries that are added automatically by the system, added manually by a user, or learned from a routing protocol (SD-WAN Fabric Subnet Sharing, BGP, or OSPF).

Route Maps

Orchestrator supports the ability to apply route maps to various routing protocols. This provides more control to import and export routes to and from the SD-WAN fabric. You can configure your route maps to modify information of a route through ACLs and applying tags by using commands. Each route map has a **match** command and **set** command. The match command verifies the attributes of the original route the protocol supports. The set command modifies information that is redistributed into the target protocol.

NOTE Prefix match criteria is 'exact match + less than'. Both the prefix specified and any subnets of that prefix will be matched. This behavior will be updated in a future release to allow for selection of 'exact,' 'greater than,' or 'less than' criteria.

To permit a default-route, deny 0.0.0.0/1, deny 128.0.0.0/1, and then permit any.

Silver Peak supports route mapping for the following protocols and the direction of those protocols:

- Local, static to SD-WAN fabric
- BGP, OSPF to SD-WAN fabric
- SD-WAN fabric to BGP Outbound peers

- Local, BGP, OSPF to BGP outbound peers
- Local BGP Peers to EdgeConnect BGP sessions

The following table lists the routing protocols and the associated commands supported.

Command	Redistribution Support	BGP	OSPF	SD-WAN	Local/Static
Match prefix	Yes	Yes	Yes	Yes	Yes
Set metric	Yes	Yes	Yes	Yes	Yes
Set tag	Yes	Yes	Yes	Yes	Yes

You can filter the type of routes displayed by clicking **All**, **Local / Static**, **SD-WAN Fabric**, **BGP**, or **OSPF**.

Import

Click **Import** to import route details from a CSV file into the selected appliance. The CSV file should contain values for the following fields in the exact order specified: Subnet, Mask Length, Metric, Is Local, Advertise to Silver Peak Peers, Advertise to BGP Peers, Next Hop, Advertise to OSPF Neighbors, Interface Name, Segment.

NOTE The CSV file should not contain a header row, and it should have no spaces after commas. You can specify only the Subnet, Mask Length, and Metric, and Orchestrator uses default values for the remaining fields. If you include values in any of the remaining fields, however, all fields must have a value (that is, none can be blank).

The following lines illustrate what two rows in a CSV import file might look like:

```
10.1.0.0,16,50,TRUE,FALSE,TRUE,10.1.0.1,FALSE,lan0,Default
10.2.0.0,16,50,,,,,,,,
```

Export

Click **Export** to save the contents of the Routes table 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 on 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

Segment

The segments you have configured on the Routing Segmentation tab are listed in the Segment field. After you specify the segment, the Routes table displays only the routes belonging to that segment.

The following information is displayed for each route listed in the table:

Field	Description
Appliance Name	Name of the appliance.
Segment	Routes displayed belonging to this segment.
Subnet/Mask	Actual subnet to be shared or learned.
Next Hop	Next-hop IP address for the route. A maximum of 16 next-hops are supported per logical interface.
Interface	Interface for outgoing traffic. Display only.
State	Shows whether the route is up or down.
Metric	Metric of the subnet. Value must be between 0 and 100. When a peer has more than one tunnel with a matching subnet (for example, in a high availability deployment), it chooses the tunnel with the lower numerical value.
Advertise to Peers	<p>Select to share subnet information with categories of peers. Select from the following options:</p> <ul style="list-style-type: none"> ■ Advertise to Silver Peak Peers ■ Advertise to BGP Peers ■ Advertise to OSPF Peers <p>Peers then learn the subnets. To add a subnet to the table without divulging it to peers, clear this option.</p>
Type (indicates one of the following route types)	
Auto (System)	Automatically added subnets of interfaces on this appliance.
Auto (SaaS)	Automatically added subnets from SaaS services.

Field	Description
Added by user	Subnets manually added or configured on this appliance.
SP: Hostname	Subnets added by exchanging information with peer appliances. If the peer has learned the subnet from a remote BGP or OSPF peer, that information is appended.
<BGP peer Type>: <BGP peer ip>	Subnets added by exchanging information with local BGP peers.
OPSF: OPSF neighbor IP	Subnets added by exchanging information with local OSPF peers.
Additional Info (indicates any tags for restricting route lookups)	
Tag FROM_LAN	Used to restrict route lookups to traffic arriving on a LAN-side interface.
Tag FROM_WAN	Used to restrict route lookups to traffic arriving on a WAN-side interface.
Comment	Any additional information you would like to include.

To edit a route, select the edit icon in the Routes table.

Route Table Lookup Criteria

Each Route table has lookup criteria that is used in the following order:

- Longest Prefix Match
- Route Table admin distance of the source protocol (lower the better)
- Metric (lower the better)
- Use peer priority (if configured) as a tie-breaker

If there are two or more routes that match all the above criteria, use multiple routes.

Admin Distance Configuration

You can configure the admin distance by using the Admin Distance template on the **Templates** tab. The default settings in this template determine the most reliable route with the use of admin distance. See the table below for the various default admin distances per route type.

Route Type	Default Admin Distance
Local	1
SD-WAN Fabric - Static	10
SD-WAN Fabric - BGP	15
SD-WAN Fabric - OSPF	15

eBGP	20
OSPF	110
iBGP	200

Navigate to the **BGP** and **OSPF** tabs for more information about applying or configuring your route maps.

Edit or Add Routes

The following table describes the elements in the Routes dialog box. They represent various features you can apply to your route.

Field	Description
Automatically advertise local LAN subnets	Indicates whether the system-created LAN subnets of your appliance should be advertised to your peers.
Automatically advertise local WAN subnets	Indicates whether the system-created local WAN subnets of your appliance should be advertised to your peers.
Metric for automatically added routes	Metric assigned to subnets of interfaces on this appliance. Specify a value from 0 to 100. The default value is 50. 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 metric value.
Redistribute routes to SD-WAN Fabric	Route redistribution map for the SD-WAN Fabric. Click the edit icon next to this field and specify the appropriate route redistribution map.
Filter Routes From SD-WAN Fabric With Matching Local ASN	Indicates whether to filter routes from the SD-WAN fabric with matching local Autonomous System Number (ASN).
Include BGP Local ASN to routes sent to SD-WAN Fabric	Indicates whether all routes must carry local ASN over subnet sharing to remote Silver Peak peers.
Tag BGP communities to routes	Indicates whether all routes carry BGP Local Communities over subnet sharing to remote Silver Peak peers. Additionally, all routes advertised to local BGP peers include local communities in the BGP route advertisements.
Communities	BGP communities to share. A community must be a combination of two numbers (0 to 65535) separated by a colon. For multiple communities, use a comma to separate them. You can have up to nine communities per route shared with subnet-sharing.
Use SD-WAN Fabric Learned Routes	Indicates whether to use SD-WAN fabric learned routes.
Enable Equal Cost Multi Path (ECMP)	Indicates whether you want to enable Equal Cost Multi-Path routing support.

Add Routes

Use the Add Routes dialog box to add a user-defined route to an appliance's route table.

1. On the Routes dialog box, click **Add Routes**.

The Add Route dialog box opens.

2. Configure the following elements as needed.

Field	Description
Subnet/Mask	Subnet IP address and mask (for example, 4.4.4.4/32).
Next Hop	Next-hop IP address for the route. If you specify a next hop, you cannot select a zone for the route. (Optional)
Interface	Interface for outgoing traffic. Click in the field and select the appropriate interface. If you specify an interface, you cannot select a zone for the route. (Optional)
Zone	Firewall zone to apply to the route. Select the appropriate firewall zone from the drop-down list. Initially, this field is set to Default. If you specify a next hop or an interface, you cannot select a zone for the route; the field automatically sets to None and cannot be changed. (Optional)
Metric	Metric for the subnet. Specify a value from 0 to 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 metric value. The default value is 50.
Tag	Tag for restricting route lookups. It is primarily used to filter routes from being redistributed in a routing loop. Select one of the following options from the drop-down list: <ul style="list-style-type: none"> ■ ANY – Allows route lookups for traffic arriving on a LAN-side or WAN-side interface. ■ FROM_LAN – Restricts route lookups to traffic arriving on a LAN-side interface. ■ FROM_WAN – Restricts route lookups to traffic arriving on a WAN-side interface.
Comments	Additional information you want to provide about this route. (Optional)

3. Click **Add**.

Import Subnets

Do the following to import route details from a CSV file into the selected appliance.

1. Click **Choose File**.
2. Locate and select the CSV file on your local machine, and then click **Open**.
3. Click **Import**.

Orchestrator imports the information from the selected file and the Routes table displays new or updated route details.

OSPF Tab

Configuration > Networking > Routing > OSPF

This tab manages **OSPF (Open Shortest Path First)** on LAN and WAN interfaces.

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

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.

Field	Description
Appliance Name	Name of the appliance.
Enable	[Route Metric] Cost associated with a route. The higher the value, the less preferred.
Router ID	This router identifier is the IPv4 address by which the remote peer can identify this appliance for purposes of OSPF.
Redistribute Routes to OSPF	Redistribution map being used to redistribute routes to OSPF.
Details	Any additional details about your route.

Select the edit icon in the OSPF table to edit and enable OSPF.

OSPF Edit Row

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

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

Field	Description
Enable OSPF	When enabled, the appliance has access to use the OSPF protocol.
Router ID	IPv4 address of the router that the remote peer uses to identify the appliance for purposes of OSPF.
Redistribute routes to OSPF	Redistributing routes into OSPF from other routing protocols or from static will cause these routes to become OSPF external routes. Select the edit icon to the left of this field and select the OSPF route redistribution maps you would like to select.

To add an additional interface to an OSPF route, click **Add** in the **Interfaces** section.

To configure or modify an OSPF route map, select the edit icon next to the Redistribute routes to OSPF field.

Add Interface

Complete the following fields to add an interface to OSPF.

Field	Description
Interface	Indicates whether a Backup Designated Router (BDR) is specified for the Designated Router (DR). Options are Yes or No .
Area ID	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.
Cost	The cost of an interface in OSPF is an indication of the overhead required to send packets across a certain interface. It is used in the OSPF path calculation to determine link preference.
Priority	Router priority. (If two or more best routes are subnet shared, peer priority is used as the tie-breaker.)
Admin Status	Indicates whether the interface is set to admin UP or DOWN .
Hello Interval	Specifies the length of time, in seconds, between the hello packets that a router sends on an OSPF interface.
Dead Interval	Number of seconds that a router's Hello packets have not been seen before its neighbors declare the OSPF router down.
Transmit Delay	Number of seconds required to transmit a link state update packet. Valid values are 1 to 65535.
Retransmit Interval	Amount of time (in seconds) the router will wait to send retransmissions if the router receives no acknowledgment.
Authentication	<ul style="list-style-type: none"> ■ None – No authentication. ■ Text – Simple password authentication allows a password (key) to be configured per area. ■ MD5 – 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.
Comment	Any information you want to include for your own use.

OSPF Route Redistribution Maps

Route Maps are policies that can be applied to static, OSPF, BGP, and SD-WAN fabric learned routes. These policies have match and set criteria. A route map is applied to routes during route redistribution between routing protocols and allows for filtering routes or modifying route attributes.

- You can specify up to 20 OSPF route maps.
- You can apply up to 128 rules per route map.

You can add, delete, rename, or clone route maps using this window. You can add rules to your route map to further specify routing protocols by clicking **Add Rule**. Use rules to allow or deny routes based on numerous matching criteria.

NOTE Prefix match criteria is 'exact match + less than'. Both the prefix specified and any subnets of that prefix will be matched. This behavior will be updated in a future release to allow for selection of 'exact,' 'greater than,' or 'less than' criteria.

To permit a default-route, deny 0.0.0.0/1, deny 128.0.0.0/1, and then permit any.

You can specify the following fields in each rule for the selected route map:

Source Protocol	Fields (based on protocol chosen)
Priority	<ul style="list-style-type: none"> ■ If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from 1000 – 9999 before applying its policies. ■ You can create rules with higher priority than Orchestrator rules (1 – 999) and rules with lower priority (10000 – 19999 and 25000 – 65534). <p>NOTE The priority range from 20000 to 24999 is reserved for Orchestrator.</p> <ul style="list-style-type: none"> ■ When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.
Local/Static	<ul style="list-style-type: none"> ■ Prefix
SD-WAN Routes	<ul style="list-style-type: none"> ■ Prefix ■ BGP Communities ■ OSPF Tag

BGP

- Prefix
- BGP Communities

Set Actions

- **Permit** – Specify if you want to allow or deny the route map.
- **OSPF Tag** – Value of OSPF tag to set in routing information sent to destination.
- **OSPF Metric Type** – Filters redistributed routes to OSPF.
- **Metric** – The metric for the route.
- **Comment** – Any additional comments you would like to include.

Multicast

Configuration > Networking > Routing > 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. Orchestrator provides four views of multicast status, each accessible by one of the corresponding buttons at the top of the Multicast tab: **Summary**, **Interfaces**, **Neighbors**, and **Routes**.

Descriptions of fields on the Summary view follow:

Field	Description
Appliance Name	Name of the appliance (also selected in the left menu) associated with the multicast configuration.
Enable	Indicates whether multicast is enabled.
Rendezvous Point IP	IP address of the centralized, source router distributing the packet of traffic to each router involved in multicast.

Click the edit icon to enable or disable multicast, add an interface for multicast, or edit an existing interface.

Multicast Dialog Box

From the Summary, Interfaces, Neighbors, or Routes view on the Multicast tab:

1. Click the edit icon next to the appliance for which you want to set up multicast.
The Multicast dialog box opens.
2. Move the **Enable Multicast** toggle to the right to enable multicast.
3. In the **Rendezvous Point IP Address** field, enter the appropriate IP address.

Interfaces

Field	Description
Interface	Name of the interface you want to connect.
PIM Enabled	Indicates whether Protocol Independent Multicast is enabled. This allows routers to communicate through the unidirectional shared trees within multicast through the shortest path.
IGMP Enabled	Indicates whether Internet Group Management Protocol is enabled. This establishes the other routers in the multicast group.
DR Priority	Designated router priority of the given interface.
DR Router IP	IP address of the designated router within your network.

To add an interface:

1. Click **Add**.
The Add Interface dialog box opens.
2. Select the desired interface from the **Interface** drop-down list.
3. Select the **Enable PIM** check box if you want to enable it.
4. Select the **Enable IGMP** check box if you want to enable it.
5. Click **Add**.

Neighbors

Field	Description
Interface	Name of the interfaces you want to connect.
Neighbor DR Priority	Designated router priority of the neighbor.
Neighbor IP	IP address of the neighbor.

Routes

Field	Description
Source	Transmitter of the multicast data.
Group	IP address of the multicast group.
Incoming Interface	Interface that receives inbound traffic.
Outgoing Interfaces	Interface that receives outbound traffic.

On the Multicast tab, you can click **Export** to export an excel file of the multicast report. You also can click the refresh button to update information displayed on the tab.

Peer Priority Tab

Configuration > Networking > Routing > 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, the appliance randomly distributes flows among multiple peers.
- The lower the number, the higher the peer's priority.

Click the edit icon to configure a peer and its peer 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.

Peer Priority Edit Row

This dialog box displays a list of configured peers. The peer priority and advertise metric are displayed for each peer.

- Peer priority controls the peer to which traffic is sent when route ties occur. It acts similar to BGP's local preference.
- Advertise metric controls the return path of a flow back toward the local appliance. It adjusts the metric of all routes sent to Peer Name. Advertise metric announces different metrics to different fabric peers. It acts similar to BGP's Multi Exit Discriminator (MED). The default setting is *preserve existing* (do nothing).

Both peer priority and advertise metric impact all routes sent and received from Peer Name.

To add a peer:

1. Click **Add Peer**.
2. In the new row that is added to the table, enter the peer name, peer priority, and advertise metric.
3. To delete a peer, click the **X** in the far-right column of the peer's row.
4. When finished, click **Apply**.

Admin Distance Tab

Configuration > Networking > Routing > Admin Distance

This tab shows 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.

The following table displays the values associated with various types of **Admin Distance**.

Field	Description
Appliance Name	The name of the appliance.
Local	A manually configured route, or one learned from locally connected subnets.
EBGP	External BGP: exchanging routing information with a router outside the company-wide network.
IBGP	Internal BGP: exchanging routing information with a router inside the company-wide network.
Subnet Shared - Static Routes	A route learned from a Silver Peak peer.
Subnet Shared - BGP Remote	A route shared from a Silver Peak peer from an external network.
OSPF	A route learned from an OSPF (Open Shortest Path First) neighbor.
Subnet Shared - OSPF Remote	A route learned from a Silver Peak peer.

To edit these fields, click the edit icon.

Management Routes Tab

Configuration > Networking > Routing > Management Routes

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

Management Routes x

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	100
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 might 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**, 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**, 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, that IP address is used instead.

Tunnels Tab

Configuration > Networking > Tunnels > Tunnels

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

- If you have deployed an SD-WAN network, **Business Intent Overlays** (BIOs) govern tunnel creation and properties.
 - Overlay tunnels consist of bonded underlay tunnels.

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

Add a Tunnel

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

Field	Description
Appliance	Name of the selected appliance.
Segment	Name of the segment, if enabled.
Overlay Tunnel	Designated overlay tunnel.
Overlay	Tunnels are applied to this designated overlay.
Admin Status	Indicates whether the tunnel has been set to admin Up or Down .

Field	Description
Status	<p>Indications are as follows:</p> <ul style="list-style-type: none"> ■ Down – The tunnel is down. This can be because the tunnel administrative setting is down or the tunnel cannot communicate with the appliance at the other end. Possible causes are: <ul style="list-style-type: none"> • Lack of end-to-end connectivity / routability (test with <i>iperf</i>). • Intermediate firewall is dropping the packets (open the firewall). • Intermediate QoS policy (be packets are being starved. Change control packet DSCP marking). • Mismatched tunnel mode (udp / gre / ipsec / ipsec_udp). • IPSec is misconfigured: (1) enabled on one side (see <i>show int tunnel configured</i>), or mismatched pre-shared key. ■ Down - In progress – The tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel. ■ Down - Misconfigured – The two appliances are configured with the same System ID (see <i>show system</i>). ■ Up - Active – The tunnel is up and active. Traffic destined for this tunnel will be forwarded to the remote appliance. ■ Up - Active - Idle – The tunnel is up and active, but it has not had recent activity in the past five minutes, and it has slowed the rate of issuing keep-alive packets. ■ Up - Reduced Functionality – The tunnel is up and active, but the two endpoint appliances are running mismatched software releases that give no performance benefit. ■ UNKNOWN – The tunnel status is unknown. This can be because the appliance is unable to retrieve the current tunnel status. Try again later.
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. Auto allows the tunnel MTU to be discovered automatically. It overrides the MTU setting.
Uptime	How long since the tunnel has been up.
Underlay Tunnels	Designated underlay tunnel.
Live View	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 are expecting tunnels to be built?*

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

2. *Are the appliances on which you are expecting the Overlays to be built using Release 8.0 or later?*

View the active software releases on **Administration > Software > Upgrade > Software Versions**.

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

Verify this on 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 **do not** want those appliances to connect directly. To view the list of Site Names, navigate to the **Configuration > Networking > Tunnels > Tunnels** tab, and then click **Sites** at the top.

Use 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 and compatible third-party IPSec tunnels from EdgeConnect devices to non-EdgeConnect 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 Exception

Configuration > Networking > Tunnels > Tunnel Exception

Orchestrator includes a tunnel exception feature that enables 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.

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



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 Use the description field to add a comment if you want to indicate why you are adding an exception.

Schedule Auto MTU Discovery

Configuration > Networking > Tunnels > Auto MTU Discovery

Use this dialog box to schedule when to discover Auto MTU.



Policy Configuration Tabs

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

DNS Proxy Policies

Configuration > Networking > DNS Proxy

The DNS (Domain Name Server) Proxy stores public IP addresses with their associated domain name. Server A is used primarily 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 also is used for public (cloud services) to breakout traffic. See the table below for the field descriptions on this tab.

Field	Description
Appliance Name	Name of the appliance associated with DNS proxy.
DNS Proxy Enabled	Whether the DNS Proxy is enabled. Select True or False .
Interface	Name of the interface associated with the DNS proxy.
Server A Addresses	IP addresses of Server A.
Server A Domains	Domain addresses of Server A.
Server A Caching	Whether you configured the server to be cached.
Server B Addresses	IP addresses of Server B.
Server B Domains	Domain addresses of Server B.
Server B Caching	Whether you configured the server to be cached.

Configure 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 whether 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 whether 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 > Templates & Policies > 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 are 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 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 also might 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 are 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

- If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies.
- You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

NOTE The priority range from **20000** to **24999** is reserved for Orchestrator.

- When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

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 & Policies > ACLs > 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** check box.

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 > Templates & Policies > Policies > QoS Policies

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.

QoS Policies ×								
Manage QoS Policies with Templates								
Export <input checked="" type="checkbox"/> Display active policies <input type="button" value="↻"/> 1 min								
QoS Policies ?								
120 Rows Search								
Edit	Appliance Na...	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

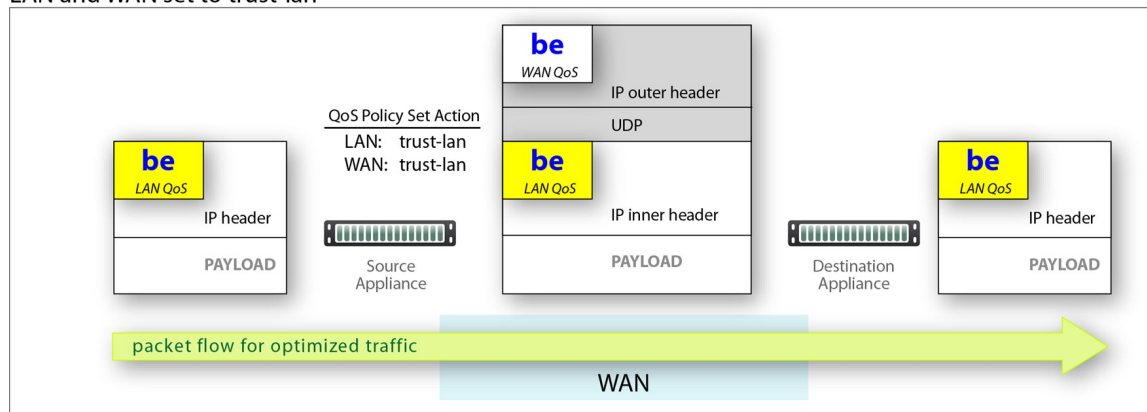
Handle and Mark DSCP Packets

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

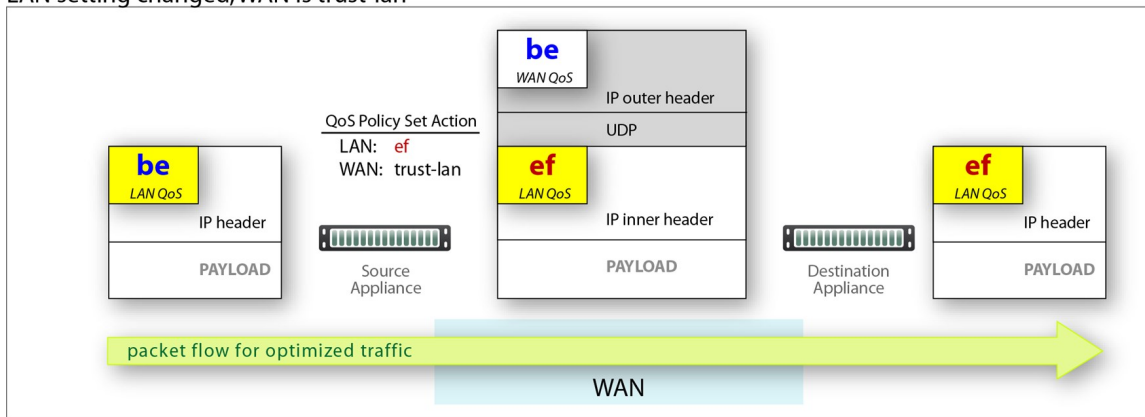
Apply DSCP Markings to Optimized (Tunnelized) Traffic

- 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.

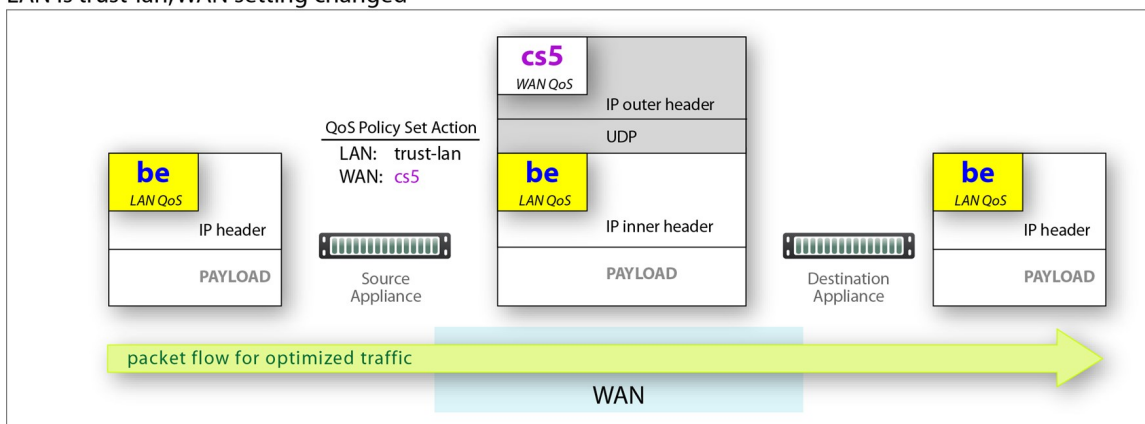
LAN and WAN set to trust-lan



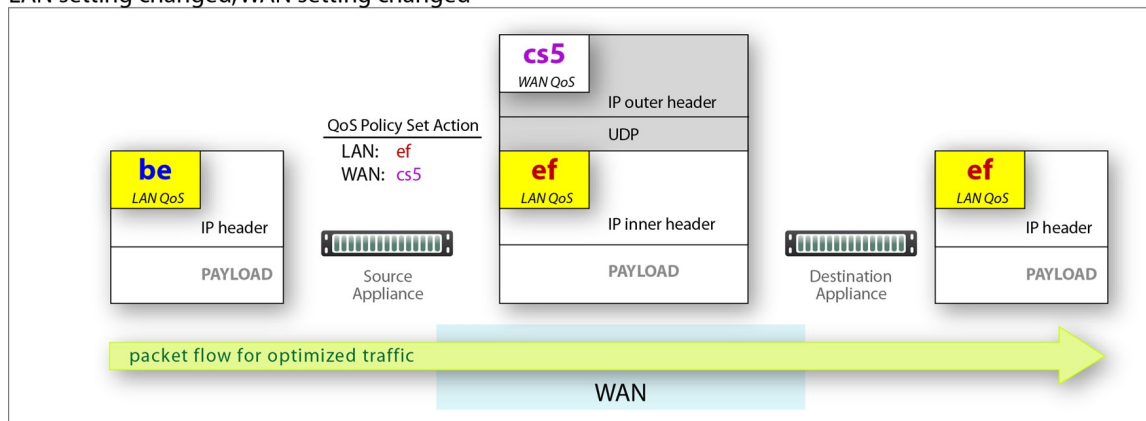
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



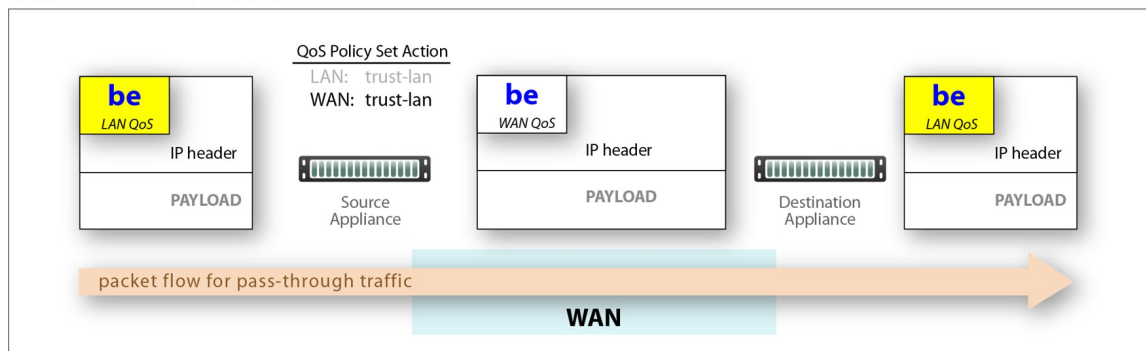
LAN setting changed, WAN setting changed



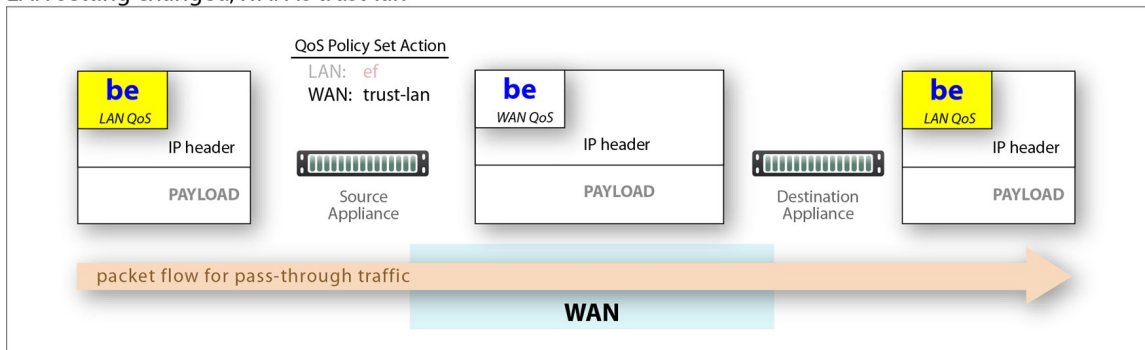
Apply DSCP Markings to Pass-through Traffic

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

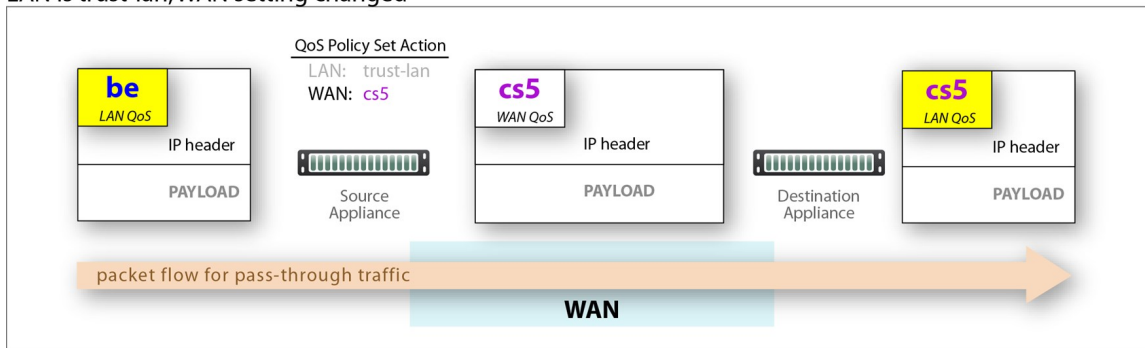
LAN and WAN set to trust-lan



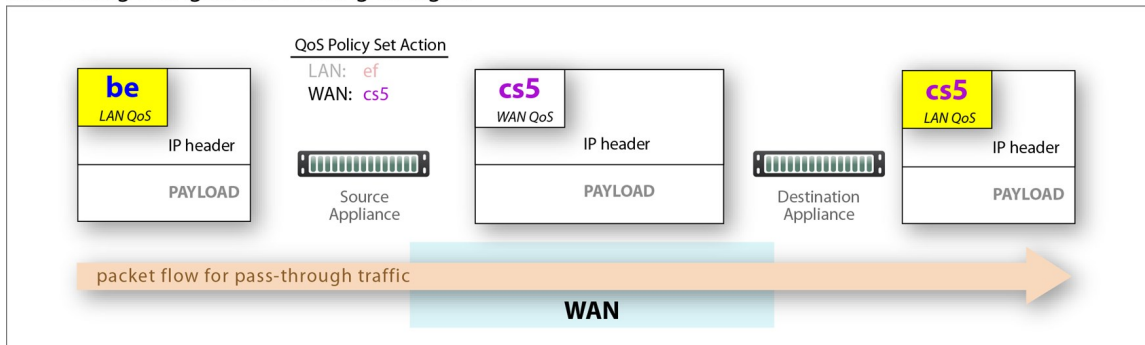
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



LAN setting changed, WAN setting changed



Priority

- If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies.
- You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

NOTE The priority range from **20000** to **24999** is reserved for Orchestrator.

- When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

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 & Policies > ACLs > 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** check box.

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 > Templates & Policies > Policies > Schedule QoSMap 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

Add

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

[Schedule QoSMap](#) [Cancel](#)

Before using this option, verify the following:

- The desired Template Group has the QoS maps you need.
- You have applied the Template Group to the appliances you want to schedule.

TIP To specify the timezone for scheduled jobs and reports, use the Schedule Timezone window (**Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**).

Optimization Policies Tab

Configuration > Templates & Policies > 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.

Edit	Appliance Name	Map	Priority	Match Criteria	Network H.	IP Header	Payload Co.	TCP Accel	TCP Accel D.	Protocol Ac.	Comment
	Chennai	map1 (active)	10000	Protocol tcp, Destination Port 139	balanced	Yes	Yes	Yes		cfs	
	Chennai	map1 (active)	10010	Protocol tcp, Destination Port 445	balanced	Yes	Yes	Yes		cfs	
	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		cbrix	
	Chennai	map1 (active)	10040	Protocol tcp, Destination Port 1494	balanced	Yes	Yes	Yes		cbrix	
	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		cfs	
	London	map1 (active)	10010	Protocol tcp, Destination Port 445	balanced	Yes	Yes	Yes		cfs	
	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		cbrix	
	London	map1 (active)	10040	Protocol tcp, Destination Port 1494	balanced	Yes	Yes	Yes		cbrix	

Priority

- If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from **1000 – 9999** before applying its policies.
- You can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

NOTE The priority range from **20000** to **24999** is reserved for Orchestrator.

- When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

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 & Policies > ACLs > 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** check box.

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	Description
Network Memory	<p>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> <ul style="list-style-type: none"> ▪ Maximize Reduction – Optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP, where bandwidth savings are the primary concern. ▪ Minimize Latency – Ensures that Network Memory processing adds no latency. This might come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It is also appropriate when the primary objective is to fully utilize the WAN pipe to increase the LAN-side throughput, as opposed to conserving WAN bandwidth. ▪ Balanced – Is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types. ▪ Disabled – Turns off Network Memory.
IP Header Compression	<p>Process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end. It is 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.</p>
Payload Compression	<p>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.</p>
TCP Acceleration	<p>Uses techniques such as selective acknowledgments, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links.</p> <hr/> <p>INFO The slow LAN alert goes off when the loss has fallen below 80% of the specified value configured in the TCP Accel Options window.</p> <hr/> <p>For more information, see TCP Acceleration Options.</p>
Protocol Acceleration	<p>Provides explicit configuration for optimizing CIFS, SSL, SRDF, Citrix, and iSCSI protocols. In a network environment, it is possible that not every appliance has the same optimization configurations enabled. Therefore, the site that initiates the flow (the client) determines the state of the protocol-specific optimization.</p>

TCP Acceleration Options

TCP acceleration uses techniques such as selective acknowledgment, 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 Block 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"/> ▼
Per-Flow Buffer	
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)
Keep Alive Timer	
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 $TCP\ MSS = Tunnel\ MTU - Tunnel\ Packet\ Overhead$.</p> <p>This feature is enabled by default so that the maximum value of the end host MSS is always coupled to the Tunnel MSS. If the end host MSS is smaller than the tunnel MSS, the end host MSS is used instead.</p> <p>A use case for disabling this feature is when the end host uses Jumbo frames.</p>
Auto Reset Flows	<p>NOTE Whether this feature is enabled or not, the default behavior when a tunnel goes Down is to automatically reset the flows.</p> <p>If enabled, it resets all TCP flows that are not accelerated, but should be (based on policy and on internal criteria like a Tunnel Up event).</p> <p>The internal criteria can also include:</p> <ul style="list-style-type: none"> ■ Resetting all TCP accelerated flows on a Tunnel Down event. ■ Resetting <ul style="list-style-type: none"> ■ TCP acceleration is enabled. ■ SYN packet was not seen (so this flow was either part of WCCP redirection or it already existed when the appliance was inserted in the data path).
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 EdgeConnect appliances in the flow's data path, and optimizes accordingly.</p> <p>Disable this feature if there is 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 ON (Default), TCP on the local appliance synchronizes this graceful shutdown of the local LAN side with the remote Silver Peak's LAN side. When this feature is OFF (Default TCP), no such synchronization happens and the two LAN segments at the ends gracefully shutdown, independently.</p>
IP Block Listing	<p>If selected, and if the appliance does not receive a TCP SYN-ACK from the remote end within five seconds, the flow proceeds without acceleration and the destination IP address is blocked for one minute.</p>
Keep Alive Timer	<p>Allows changing the Keep Alive timer for the TCP connections.</p> <ul style="list-style-type: none"> ■ Probe Interval – Time interval in seconds between two consecutive Keep Alive probes. ■ Probe Count – Maximum number of Keep Alive probes to send. ■ First Timeout (Idle) – Time interval until the first Keep Alive timeout.

Option	Description
LAN Side Window Scale Factor Clamp	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 in which there are many out-of-order packets being received from the LAN side. These out-of-order packets cause much buffer utilization and maintenance.
Per-Flow Buffer	(Max LAN to WAN Buffer and Max WAN to LAN Buffer) This setting clamps the maximum buffer space that can be allocated to a flow, in each direction.
Persist timer Timeout	Allows the TCP to terminate connections that are in Persist timeout stage after the configured number of seconds.
Preserve Packet Boundaries	Preserves the packet boundaries end-to-end. If this feature is disabled, the appliances in the path can coalesce consecutive packets of a flow to use bandwidth more efficiently. It is enabled by default so that applications requiring packet boundaries to match do not fail.
Route Policy Override	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 this case, you might need to configure flow redirection to ensure optimization of TCP flows.
Slow LAN Defense	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 so 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.
Slow LAN Window Penalty	This setting (OFF by default) penalizes flows that are slow to send data on the LAN side by artificially reducing their TCP receive window. This causes less data to be received and helps to reach a balance with the data sending rate on the LAN side.
WAN Congestion Control	Selects the internal Congestion Control parameter: <ul style="list-style-type: none"> ■ Optimized – This is the default setting. This mode offers optimized performance in almost all scenarios. ■ Standard – In some unique cases, it might be necessary to downgrade to Standard performance to better interoperate with other flows on the WAN link. ■ Aggressive – Provides aggressive performance and should be used with caution. Recommended mostly for Data Replication scenarios.
WAN Window Scale	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 > Templates & Policies > Policies > SaaS NAT Policies

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

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

NAT Policies ×

Manage NAT Policies with Templates

BasicAdvancedExport↺▼

NAT Policies ?

9 Rows

Search

Edit	Appliance Name	NAT All Inbound			NAT All Outbound		
		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 Name	Map	Priority	Match Criteria	NAT Type	NAT Direct...	NAT IP	Fallback	Comment
✓	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

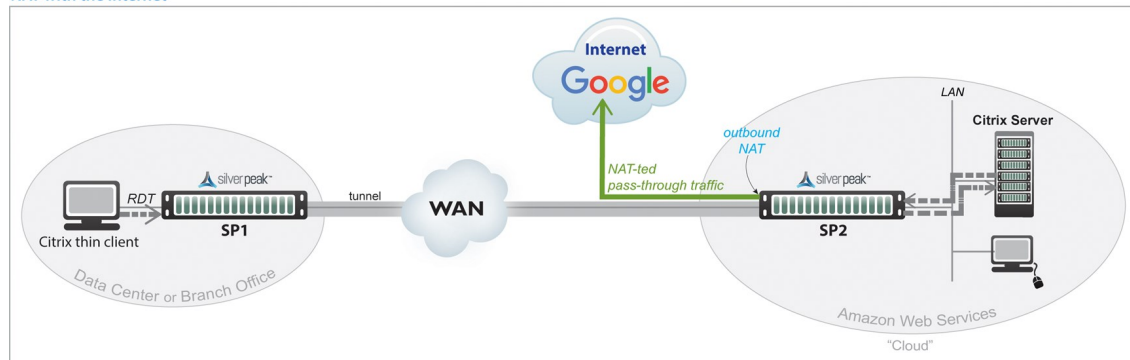
servers has a return path to the appliance from which that traffic originated.

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 does not 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 does not 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 & Policies > ACLs > 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** check box.

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	Description
NAT Type	no-nat	Is the <i>default</i> . No IP addresses are changed.
	source-nat	Is the <i>default</i> . No IP addresses are changed.
NAT Direction	inbound	NAT is on the LAN interface.
	outbound	NAT is on the WAN interface.
	none	Only option if the NAT Type is no-nat .
NAT IP	auto	Select if you want to NAT all traffic. The appliance then picks the first available NAT IP/Port.
	tunnel	Select if you only want to NAT tunnel traffic. Applicable only for inbound NAT, as outbound does not support NAT on tunnel traffic.
	[IP address]	Select if you want to make NAT use this IP address during address translation.
Fallback		If the IP address is full, the appliance uses the next available IP address.

When you select a specific IP, 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 > Overlays & Security > Security > 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 on whether 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 is 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. Additionally, inbound port forwarding does not support TFTP servers.

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.

Field	Description
Source IP/Subnet	Source of the WAN device managing the LAN device(s) specified in the destination.
Destination IP/Subnet	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. Select the **Translate** check box to enable Translate mode.
5. Complete each field with the appropriate information.

Field	Description
Source IP/Subnet	Source of the WAN device managing the LAN device(s) specified in the destination.
Destination IP/Subnet	Address of the WAN interface IP.
Destination Port/Range	Port/range of the LAN device(s) that are managed remotely.
Protocol	Select the protocol you want to apply: UDP, TCP, ICMP, Any . If you select Any , the Destination and Translated Ports have a default value that need to be between 0-100. If the value exceeds, 100 a warning appears.
Translated IP	IP address of the LAN device accessed inside your network.
Translated Port/Range	Port/range of the LAN device accessed inside your network.
Source Interface	Source interface name.
Segment	Name of the segment being used.
Comment	Any additional details.

Additional Information

- **Interface Modes**

Port forwarding is used only 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 a protocol option only on versions or 8.1.9.4 and later.

Security Policies Tab

Configuration > Overlays & Security > Security > Firewall Zone Security Policies

This tab displays the Security Policies, which manage traffic between firewall zones.

- 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.
- When Routing Segmentation (VRF) is enabled, by default, traffic is allowed between interfaces labeled with the same zone and the same segment. Any traffic between different zones or between different segments is dropped.
- When segmentation is enabled, define your security policies from the Routing Segmentation (VRF) tab.
- When segmentation is enabled, do not use templates. If a security policy template is applied while segmentation is enabled, it will only apply within the default segment. It will override the default-default security policy defined on the Routing Segmentation (VRF) tab. This behavior is designed to prevent a disruption in traffic when segmentation is enabled for the first time, and during a migration to segments. After the migration process is complete, the security policy template should be removed.
- If segments are disabled, define your security policies by creating templates. You can then apply template groups to appliances.
- Clicking the edit icon opens the Security Policy that has been applied. Any changes made here are local to that appliance. Silver Peak does not recommend making changes from this tab.
- 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.
- Define your Security Policies by creating **templates**. You can then apply templates to Interfaces or Overlays.
- Clicking the edit icon opens the Security Policy that has been applied. Any changes made here are **local** to that appliance.
- Click **Firewall Drops** to see statistics on various flows, packets, and bytes dropped or allowed by a zone-based firewall for a given time range.
- Click **Manage Security Policies with Templates** to define policies on all appliances within your network. You can use the matrix and table view to further specify your policies. If segmentation is enabled, do not use templates. Manage from the Routing Segmentation (VRF) tab instead.

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 > Templates & Policies > Policies > ACLs > Access Lists

This tab lists the configured **Access Control List** (ACL) rules. An **ACL** is a reusable MATCH criteria for filtering flows. It is associated with an action: **permit** or **deny**. An ACL can be a MATCH condition in more than one policy—Route, QoS, or Optimization.

Field	Description
Appliance Name	Name the appliance selected.
ACLs	Access Control Lists. A list of one or more ordered access control rules.
	NOTE An ACL only becomes active when it is used in a policy.
Priority	<ul style="list-style-type: none"> ■ If you are using Orchestrator templates to add rules, Orchestrator will delete all entries from 1000 – 9999 before applying its policies. ■ You can create rules with higher priority than Orchestrator rules (1 – 999) and rules with lower priority (10000 – 19999 and 25000 – 65534). <p>NOTE The priority range from 20000 to 24999 is reserved for Orchestrator.</p> <ul style="list-style-type: none"> ■ When adding a rule, the priority is incremented by 10 from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.
Match Criteria	Configured ACL match criteria associated to the appliance. See below for more information about Match Criteria.
Permit	<p>Whether the ACL is set to Permit or Deny.</p> <ul style="list-style-type: none"> ■ Permit allows the matching traffic flow to proceed to the policy entry's associated SET actions. ■ Deny prevents further processing of the flow by that ACL, specifically. The appliance continues to the next entry in the policy.
Comment	Any additional information about the ACL.

Click the edit icon to make add, delete, or modify rules to your ACLs.

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 & Policies > ACLs > 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** check box.

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.

Address Groups

Configuration > Templates & Policies > ACLs > Address Groups

Use the Address Groups tab to view and manage address groups in your SD-WAN network. An address group is a logical collection of IP hosts or subnets that can be referenced in source or destination matching criteria in the zone based firewall and security policies (route, QOS, optimization, and so forth).

NOTE Orchestrator supports up to 500 address groups.

Address Groups ? 4 mins

-- Add Group Delete Group Bulk Import Export CSV

Add Rule

Memory Consumed: 30 bytes

2 Rows, 1 Selected Search

Edit	Name	Includes	Excludes	Comment
	AuthorizedDNS	8.8.8.8, 8.8.4.4, 1.1.1.1, 1.1.0.0		Firewall denies all but these endpoints
	Loopbacks-Default	128.41.0.0/16		Pool for loopback orchestration

Add an Address Group

Follow the steps below to create a new address group:

1. Click **Add Group** to open the Add Address Group dialog box.

Add Address Group ×

Group name

Rules

IPs to include

IPs to exclude

Groups to include

Comment

Add Cancel

2. Provide the following details in the fields provided:

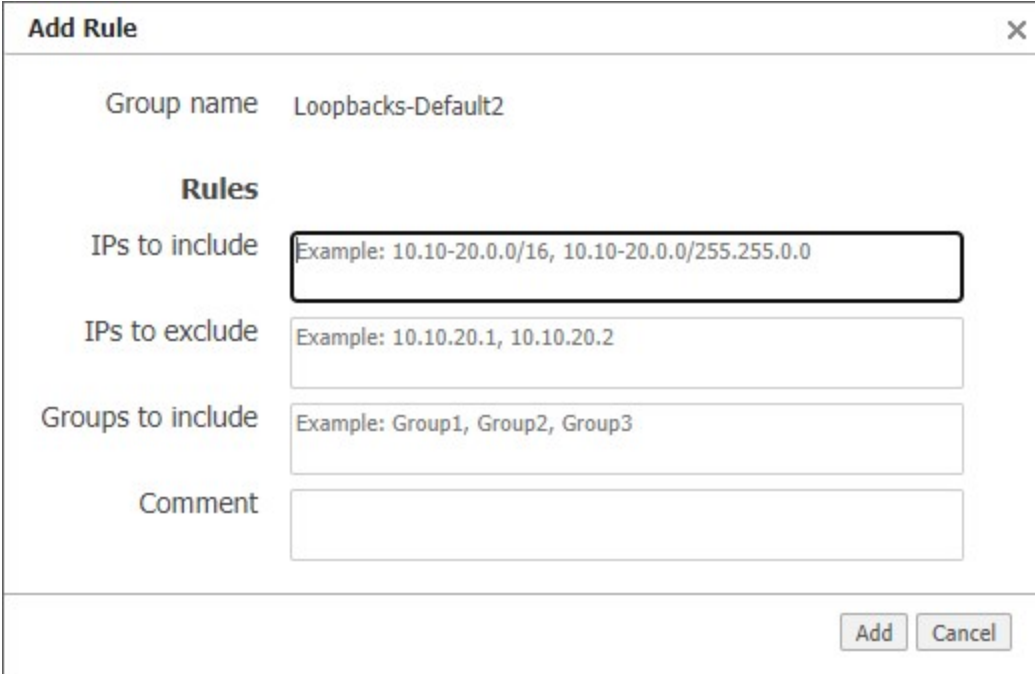
Field	Description
Group name	Enter a unique name for the group, up to 64 characters long. NOTE Group names can only contain uppercase and lowercase letters, numbers, dots, underscores, and hyphens.
IPs to include	Enter one or more IP addresses or subnets to include in the group (see Address Group Formats below).
IPs to exclude	Enter one or more IP addresses to exclude, in the case where you are including an IP range.
Groups to include	Enter the name of one or more address groups to include. NOTE Group inclusion only supports two levels of nesting. For example, if Group1 includes Group2 and Group2 includes Group3, you could not include Group1 anywhere because it already contains two levels of nested groups.
Comment	Enter an optional comment that describes the address group and how it might be used.

3. Click **Add** to create the address group or click **Cancel** to close the dialog box without making any changes.

Add a Rule to an Address Group

Follow the steps below to add a rule to an existing address group:

1. Select the address group to which you want to add a rule from the drop-down list above the table.
2. Click **Add Rule** to open the Add Rule dialog box.

The image shows a dialog box titled "Add Rule" with a close button (X) in the top right corner. Inside the dialog, there is a label "Group name" followed by the text "Loopbacks-Default2". Below this is a section header "Rules". Under "Rules", there are four input fields: "IPs to include" with the example text "Example: 10.10-20.0.0/16, 10.10-20.0.0/255.255.0.0", "IPs to exclude" with the example text "Example: 10.10.20.1, 10.10.20.2", "Groups to include" with the example text "Example: Group1, Group2, Group3", and a "Comment" field. At the bottom right of the dialog are two buttons: "Add" and "Cancel".

Add Rule

Group name Loopbacks-Default2

Rules

IPs to include Example: 10.10-20.0.0/16, 10.10-20.0.0/255.255.0.0

IPs to exclude Example: 10.10.20.1, 10.10.20.2

Groups to include Example: Group1, Group2, Group3

Comment

Add Cancel

3. Provide the details for the new rule in the fields provided (see field descriptions in **Add an Address Group**).
4. Click **Add** to create the rule or click **Cancel** to close the dialog box without making any changes.

Delete an Address Group

Follow the steps below to delete an address group:

1. Select the address group you want to delete from the drop-down list above the table.
2. Click **Delete Group**.

A confirmation dialog opens.

3. Click **Delete** to confirm your choice and permanently remove the selected group and all of its rules. Otherwise, click **Cancel** to return to the list without deleting the group.

Export Address Groups

You can export the current address groups to a CSV file as a backup to make bulk modifications outside of the Orchestrator UI. Follow the steps below to export address groups.

1. Click **Export CSV**.
2. In the save dialog box, browse to the location where you want to save the file, provide a name for the file, and then click **Save**.

3. Open the saved file in Excel or another program to view or modify its contents.

	A	B	C	D	E
1	Name	IncludedIPs	ExcludedIPs	IncludedGroups	Comment
2	AuthorizedDNS	8.8.8.8,8.8.4.4,1.1.1.1,1.1.0.0			Firewall denies all but these endpoints
3	Loopbacks-Default	128.41.0.0/16			Pool for loopback orchestration
4					

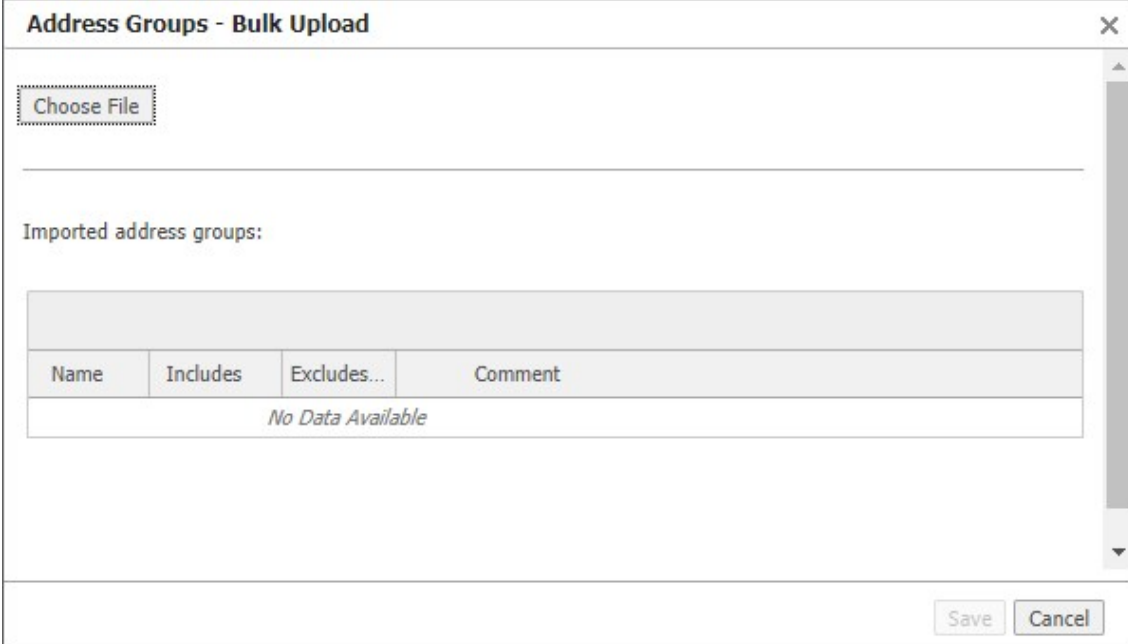
NOTE When editing exported rules and address groups, you can modify the included or excluded IPs, included groups, or comments to overwrite the same rule when imported. If you modify the group name on a rule, however, it will create a new rule when imported.

Import Address Groups

Follow the steps below to import address groups from a CSV file:

NOTE You can import a file that was exported and modified, or a new file that contains data in the same rows and columns as the exported file. Columns are ordered as Name, Included IPs, Excluded IPs, Included Groups, and Comment. The first row of the import file will be ignored.

1. Click **Bulk Import** to open the Address Groups - Bulk Upload dialog box.



The dialog box titled "Address Groups - Bulk Upload" features a "Choose File" button at the top. Below it, the section "Imported address groups:" contains a table with the following structure:

Name	Includes	Excludes...	Comment
No Data Available			

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

2. Click **Choose File**, locate and select the CSV file to be imported, and then click **Open**.

3. Review the groups and rules to be imported.
4. Click **Save** to import the file and merge with or replace the existing address groups, or click **Cancel** to close the dialog box without making any changes.

View a Single Address Group

By default, all address groups are displayed in the table on the Address Groups tab. To filter the table to a single address group, select the group from the drop-down list above the table.

NOTE You can only add rules to an existing group when viewing a single address group. You cannot add a group with the same name as an existing group.

Edit or Delete a Rule

To edit or delete an existing rule, click the edit icon to the right of the rule. The Edit Rule dialog box opens.

Edit Rule

Group name Loopbacks-Default2

Rules

IPs to include 128.41.0.0/24

IPs to exclude Example: 10.10.20.1, 10.10.20.2

Groups to include Example: Group1, Group2, Group3

Comment Pool for loopback orchestration

Save Delete Cancel

- To edit the rule, modify the available fields, and then click **Save**.
- To delete the rule, click **Delete**.

Using Address Groups in Match Criteria

When specifying match criteria for IP/Subnet, you can use an address group by enabling the **Src:Dest** and **Groups** options.

Address Group Formats

An address group can include IP addresses, subnets, address groups, or any combination thereof. For IPs and subnets, the following formats are allowed:

- One or more IP addresses: 10.10.10.1 or 10.10.10.2, 10.10.10.2, 10.10.10.3
- IP subnet: 10.10.0.0/16 or 10.10.0.0/255.255.0.0
- IP range: 10.10.10.10-20
- IP range and subnet: 10.10-20.0.0/16, 10.10-20.0.0/255.255.0.0
- IP wildcard: 10.10.10.* (you can use the wildcard in any octet)
- Wildcard and mask: 10.*.0.0/16, 10.*.0.0/255.255.0.0

Service Groups

Configuration > Templates & Policies > ACLs > Service Groups

Use the Service Groups tab to view and manage service groups in your SD-WAN network. A service group is a logical collection of protocols and ports that can be referenced in source or destination matching criteria in the zone based firewall and security policies (route, QOS, optimization, and so forth).

NOTE Orchestrator supports up to 500 service groups.

Service Groups 2 mins

--
Add Group
Delete Group
Bulk Import
Export CSV

Add Rule

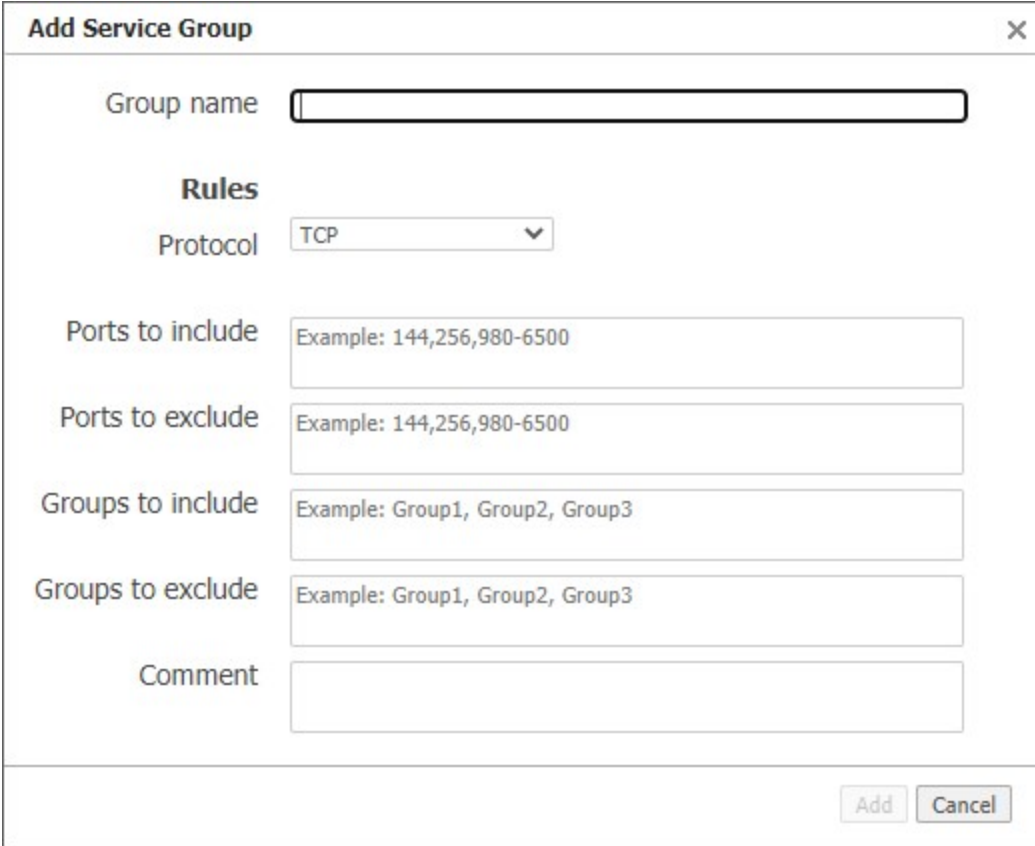
Memory Consumed: 320 bytes

15 Rows					Search <input type="text"/>
Edit	Name	Protocol	Includes	Excludes	Comment
	ICMP-Echo	ICMP	8		ping request
	ICMP-EchoReply	ICMP	0		ping reply
	ICMP-DestinationUnreacha...	ICMP	3		destination unreachable
	ICMP-TTLExceeded	ICMP	11		TTL Expired in Transit (traceroute, routing loops,...
	ICMP-Redirect	ICMP	5		redirect messages for better gateway
	SystemPorts-UDP	UDP	1-1023		system, well-known, privileged ports
	SystemPorts-TCP	TCP	1-1023		system, well-known, privileged ports
	EphemeralPorts-TCP	TCP	49152-65535		ephemeral ports
	EphemeralPorts-UDP	UDP	49152-65535		ephemeral ports
	RegisteredPorts-TCP	TCP	1024-49151		registered and unprivileged ports
	RegisteredPorts-UDP	UDP	1024-49151		registered and unprivileged ports
	Unidirectional-UDP	UDP	514, 2055, 67, 162		apps that continuously send in only one direction
	multi-services	TCP	1, 2, 3-5, ICMP-Echo	100-200, SystemPorts-TCP	
	multi-services	GRE			
	multi-services	UDP	512	123	

Add a Service Group

Follow the steps below to create a new service group:

1. Click **Add Group**. The Add Service Group dialog box opens.

The image shows a dialog box titled "Add Service Group" with a close button (X) in the top right corner. The dialog contains several input fields: "Group name" (a text box), "Rules" (a section header), "Protocol" (a dropdown menu with "TCP" selected), "Ports to include" (a text box with the example "144,256,980-6500"), "Ports to exclude" (a text box with the example "144,256,980-6500"), "Groups to include" (a text box with the example "Group1, Group2, Group3"), "Groups to exclude" (a text box with the example "Group1, Group2, Group3"), and "Comment" (a text box). At the bottom right, there are "Add" and "Cancel" buttons.

Add Service Group

Group name

Rules

Protocol

Ports to include

Ports to exclude

Groups to include

Groups to exclude

Comment

2. Provide the following details in the fields provided:

Field	Used in	Description
Group name	All	Enter a unique name for the group, up to 64 characters long. NOTE Group names can only contain uppercase and lowercase letters, numbers, dots, underscores, and hyphens.
Protocol	All	Select a protocol from the list of those available.
Ports to include	TCP, UDP	Enter one or more ports to include in the group. A single port, multiple comma-separated ports, and a range of ports are supported (e.g., 20, 22, 24-30).

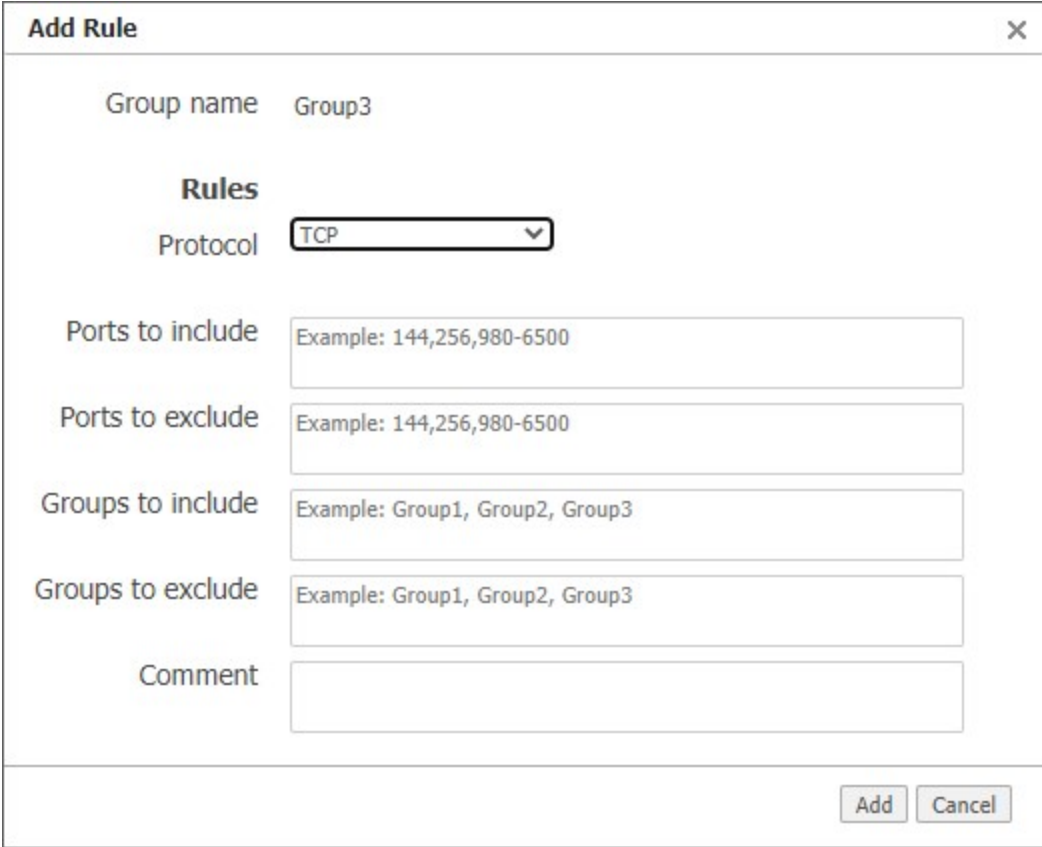
Field	Used in	Description
Ports to exclude	TCP, UDP	Enter one or more ports to exclude from the group, in the case where you are including a range of ports. A single port, multiple comma-separated ports, and a range of ports are supported (e.g., 20, 22, 24-30).
Groups to include	TCP, UDP	Enter the name of one or more service groups to include. NOTE Group inclusion only supports two levels of nesting. For example, if Group1 includes Group2 and Group2 includes Group3, you could not include Group1 anywhere because it already contains two levels of nested groups.
Groups to exclude	TCP, UDP	Enter the name of one or more service groups to exclude, in the case where you are already including a group that includes multiple groups.
ICMP types	ICMP	For ICMP, add one or more message types to include. Multiple types and ranges are supported (e.g., 1, 2, 4-8).
Comment	All	Enter an optional comment that describes the service group and how it might be used.

- Click **Add** to create the service group or click **Cancel** to close the dialog box without making any changes.

Add a Rule to a Service Group

Follow the steps below to add a rule to an existing service group:

- Select the service group to which you want to add a rule from the drop-down list above the table.
- Click **Add Rule**. The Add Rule dialog box opens.

A screenshot of the 'Add Rule' dialog box. The dialog has a title bar with 'Add Rule' and a close button. Inside, there are several fields: 'Group name' with the value 'Group3'; 'Rules' section with 'Protocol' set to 'TCP' in a dropdown; 'Ports to include' with placeholder text 'Example: 144,256,980-6500'; 'Ports to exclude' with the same placeholder; 'Groups to include' with placeholder text 'Example: Group1, Group2, Group3'; 'Groups to exclude' with the same placeholder; and a 'Comment' field. At the bottom right are 'Add' and 'Cancel' buttons.

Add Rule

Group name Group3

Rules

Protocol TCP

Ports to include Example: 144,256,980-6500

Ports to exclude Example: 144,256,980-6500

Groups to include Example: Group1, Group2, Group3

Groups to exclude Example: Group1, Group2, Group3

Comment

Add Cancel

3. Provide the details for the new rule in the fields provided (see field descriptions in **Add a Service Group**).
4. Click **Add** to create the rule or click **Cancel** to close the dialog box without making any changes.

Delete a Service Group

Follow the steps below to delete a service group:

1. Select the service group you want to delete from the drop-down list above the table.
2. Click **Delete Group**.

A confirmation dialog opens.

3. Click **Delete** to confirm your choice and permanently remove the selected group and all of its rules. Otherwise, click **Cancel** to return to the list without deleting the group.

Export Service Groups

You can export the current service groups to a CSV file as a backup to make bulk modifications outside of the Orchestrator UI. Follow the steps below to export service groups.

1. Click **Export CSV**.
2. In the save dialog box, browse to the location where you want to save the file, provide a name for the file, and then click **Save**.
3. Open the saved file in Excel or another program to view or modify its contents.

	A	B	C	D	E	F	G	H
1	Name	Protocol	IncludedPorts	ExcludedPorts	IncludedGroups	ExcludedGroups	IcmpTypes	Comment
2	ICMP-Echo	ICMP					8	ping request
3	ICMP-EchoReply	ICMP					0	ping reply
4	SystemPorts-UDP	UDP	1-1023					system, well-known, privileged ports
5	SystemPorts-TCP	TCP	1-1023					system, well-known, privileged ports
6	EphemeralPorts-TCP	TCP	49152-65535					ephemeral ports
7	EphemeralPorts-UDP	UDP	49152-65535					ephemeral ports
8	RegisteredPorts-TCP	TCP	1024-49151					registered and unprivileged ports
9	RegisteredPorts-UDP	UDP	1024-49151					registered and unprivileged ports
10	Unidirectional-UDP	UDP	514,2055,67,162					apps that continuously send in only one direction

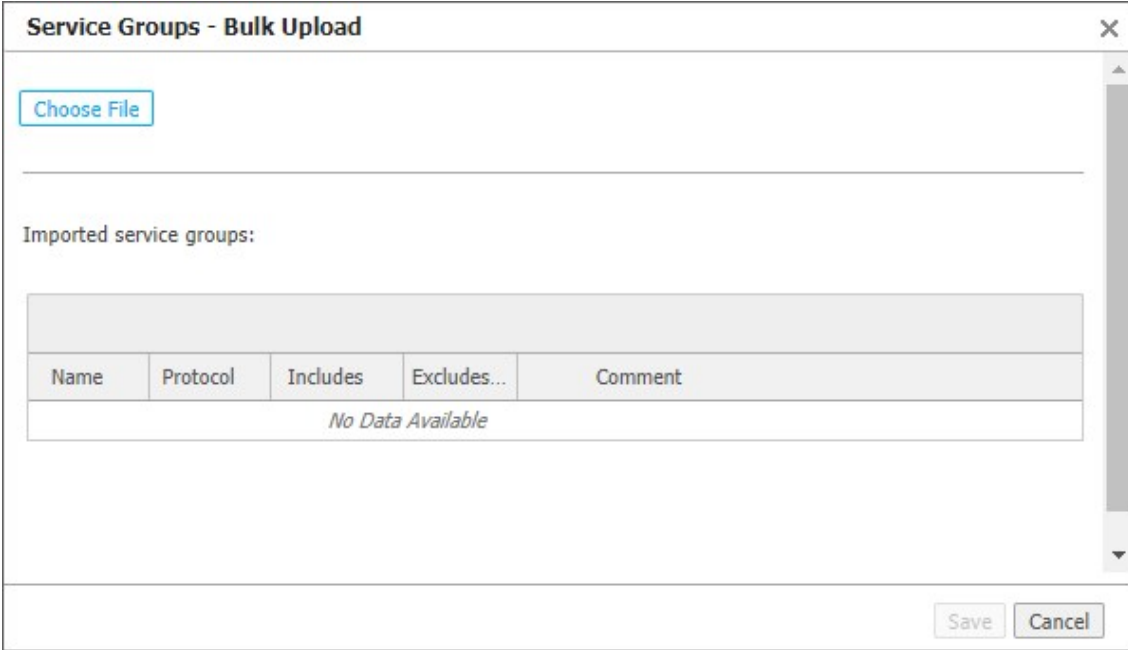
NOTE When editing exported rules and service groups, you can modify the protocol, inclusions, exclusions, ICMP types, or comments to overwrite the same rule when imported. If you modify the group name on a rule, however, it will create a new rule when imported.

Import Service Groups

Follow the steps below to import service groups from a CSV file:

NOTE You can import a file that was exported and modified, or a new file that contains data in the same rows and columns as the exported file. Columns are ordered as Name, Protocol, Included Ports, Excluded Ports, Included Groups, Excluded Groups, ICMP types, and Comment. The first row of the import file will be ignored.

1. Click **Bulk Import**. The Service Groups - Bulk Upload dialog box opens.



The dialog box is titled "Service Groups - Bulk Upload" and has a close button (X) in the top right corner. It contains a "Choose File" button in the top left. Below this is a section labeled "Imported service groups:" which contains a table. The table has five columns: "Name", "Protocol", "Includes", "Excludes...", and "Comment". The table is currently empty, with the text "No Data Available" centered in the body. At the bottom right of the dialog box are "Save" and "Cancel" buttons.

Name	Protocol	Includes	Excludes...	Comment
No Data Available				

2. Click **Choose File**, locate and select the CSV file to be imported, and then click **Open**.
3. Review the groups and rules to be imported.
4. Click **Save** to import the file and merge with or replace the existing service groups, or click **Cancel** to close the dialog box without making any changes.

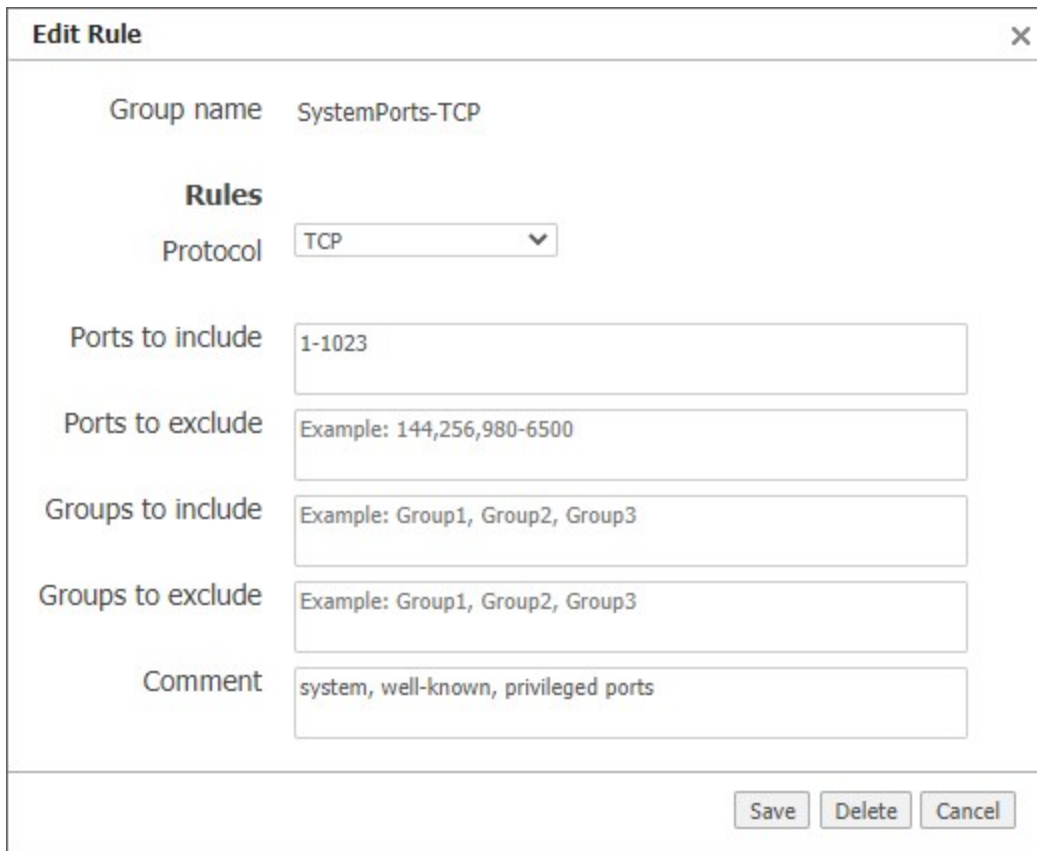
View a Single Service Group

By default, all service groups are displayed in the table on the Service Groups tab. To filter the table to a single service group, select the group from the drop-down list above the table.

NOTE You can only add rules to an existing group when viewing a single service group. You cannot add a group with the same name as an existing group.

Edit or Delete a Rule

To edit or delete an existing rule, click the edit icon to the right of the rule and the Edit Rule dialog box opens.



Edit Rule [X]

Group name SystemPorts-TCP

Rules

Protocol TCP [v]

Ports to include 1-1023

Ports to exclude Example: 144,256,980-6500

Groups to include Example: Group1, Group2, Group3

Groups to exclude Example: Group1, Group2, Group3

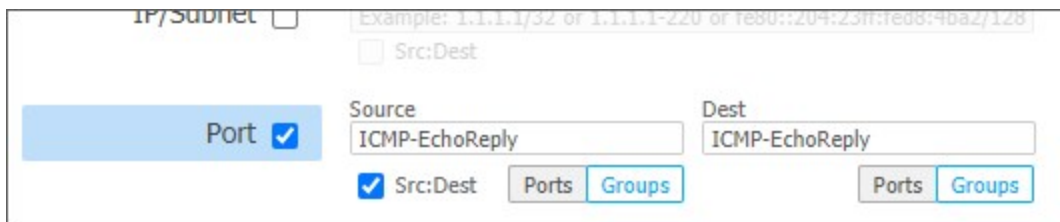
Comment system, well-known, privileged ports

[Save] [Delete] [Cancel]

- To edit the rule, modify the available fields, and then click **Save**.
- To delete the rule, click **Delete**.

Using Service Groups in Match Criteria

When specifying match criteria for Port, you can use a service group by enabling the **Src:Dest** and **Groups** options.



IP/Subnet [] Example: 1.1.1.1/32 or 1.1.1.1-220 or fe80::204:23ff:fe08:4ba2/128

[] Src:Dest

Port [x]

Source ICMP-EchoReply

Dest ICMP-EchoReply

[x] Src:Dest [Ports] [Groups] [Ports] [Groups]

Shaper Tab

Configuration > Templates & 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 ×

Manage Shaper settings with Templates Inbound Outbound Export 8 mins

Shaper 300 Rows

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_svr...	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 ×

Manage Shaper settings with Templates Inbound Outbound Export 8 mins

Shaper 30 Rows

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
✓	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 tunnelized and pass-through-shaped traffic, shaping it as it exits to the WAN.
- To manage Shaper settings for an appliance's system-level **WAN** Shaper, access the Shaper template.
- 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	Description
Excess Weighting	If there is bandwidth left over after satisfying the minimum bandwidth percentages, the excess is distributed among the traffic classes in proportion to the weightings specified in the Excess Weighting column. Values range from 1 to 10,000.
Interface Shaper	Enables a separate shaper for a specific WAN interface. <ul style="list-style-type: none"> For WAN optimization, the interface shaper can be used, but it is not recommended. For SD-WAN, it should never be used because overlay traffic is not directed to an interface shaper; traffic is always shaped by the default WAN shaper.
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 Max Wait Time 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%, lower-priority traffic classes might not receive their guaranteed bandwidth if it is all consumed by higher-priority traffic. If you set Min Bandwidth to a value greater than Max Bandwidth , then Max overrides Min .
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 - 1 is first, 10 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 0 (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 wan0 goes down, wan0 bandwidth is removed from the total bandwidth when recalculating.
Traffic ID	The number assigned to the traffic class.

Field	Description
Traffic Name	The name assigned to a traffic class, either prescriptively or by the user.

SaaS Optimization Tab

Configuration > Templates & 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 the edit icon. The SaaS Optimization window opens.

SaaS Optimization Configuration

161 Rows

Search

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	CCone	No	No	10 ms	myccportal.com, *.myccportal.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, eloquatraincenter.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.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

Application Definitions

Configuration > Templates & Policies > Applications & SaaS > Application Definitions

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

The screenshot displays the 'Application Definitions' tab in the Silver Peak Unity Orchestrator interface. At the top, there is a search bar with the text 'google' and a 'Search' button. Below the search bar, there are tabs for 'Application Definitions' and 'Application Groups'. The 'Application Definitions' tab is active, showing a list of 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. Below the application definitions table, there is a section for 'Hide Advanced App Definitions' with a dropdown menu. Below this, there is a table of 140 rows of protocols. The table has columns for 'Protocol', 'Name', 'Notes', 'Confidence', and 'Edit'. The rows list various network protocols and their corresponding names.

- 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 > Templates & Policies > Applications & SaaS > Application Groups

Application groups associate applications into a common group 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 Unity Orchestrator. On the left, there are two filter sections: 'Traffic Type' and 'Content Type'. The 'Interactive' traffic type is selected, showing 139 rows. The main table lists applications and their associated groups. The 'Interactive' group is highlighted in the sidebar.

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 blank.
- Group names are case-insensitive.
- An application group cannot contain another application group.
- A group name followed by * indicates a group defined by a user.
- You cannot change the name of a group provided by Silver Peak, but you can modify the applications those groups contain.

NOTE To avoid performance issues, it is strongly recommended that you assign an application to no more than three groups.

Threshold Crossing Alerts Tab

Configuration > Templates & Policies > TCAs > Threshold Crossing Alerts

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

Threshold Crossing Alerts ×

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
	Tallinn	File-system utilization	90%	85%	1	Yes	75%	75%	1	No
	Tallinn	LAN-side receive throughput	1000000 kbps	1000000 kbps	1	No	0 kbps	0 kbps	1	No
	Tallinn	Total number of flows	90%	85%	1	Yes	0%	0%	1	No
	Tallinn	Total number of optimized flows	90%	85%	1	No	0%	0%	1	No
	Tallinn	Tunnel OOP post-POC	100%	100%	1	No	0%	0%	1	No
	Tallinn	Tunnel OOP pre-POC	100%	100%	1	No	0%	0%	1	No
	Tallinn	Tunnel latency	1000 ms	850 ms	1	Yes	0 ms	0 ms	1	No
	Tallinn	Tunnel loss post-FEC	100%	100%	1	No	0%	0%	1	No
	Tallinn	Tunnel loss pre-FEC	100%	100%	1	No	0%	0%	1	No
	Tallinn	Tunnel reduction	100%	100%	1	No	0%	0%	1	No
	Tallinn	Tunnel utilization	100%	100%	1	No	0%	0%	1	No
	Tallinn	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 ×

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

The alerts are triggered with rising and falling threshold crossing events (that is, 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, navigate to **Configuration > Networking > Tunnels > 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 then click **OK**.

- To view globally applied system and tunnel alerts, click **System**.
- To view alerts that are specific to an individual tunnel, click **Tunnel**.

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



Rules:

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

ON by Default

- **Appliance Capacity** – Triggers when an appliance reaches 95% of its total flow capacity. It is not configurable and can be cleared only 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)

IP SLA Tab

Configuration > Templates & Policies > TCAs > IP SLA


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.

This tab displays all of the IP SLA rules configured on the selected appliances. To add or modify rules, click the edit icon to the left of any row in the table.

IP SLA Monitor Use Cases

The following examples describe five basic use cases for IP SLA monitoring.

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_ ▼


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

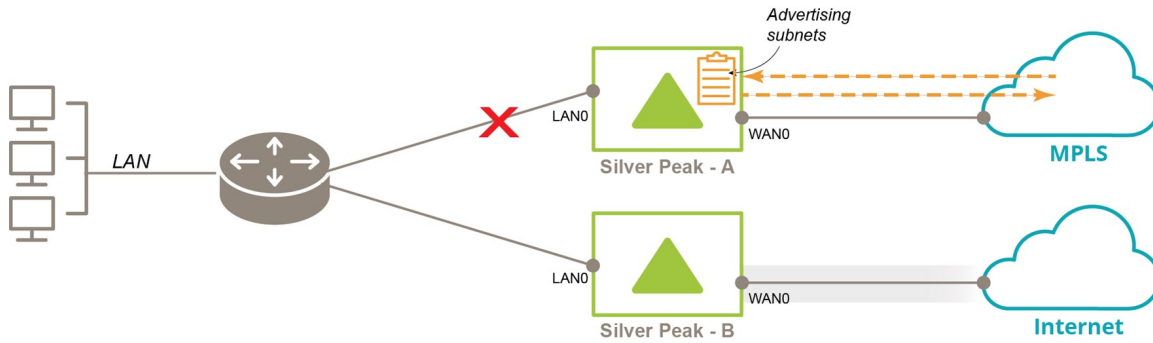
Example #2 – HTTP/HTTPS via Interface



- 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**.

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

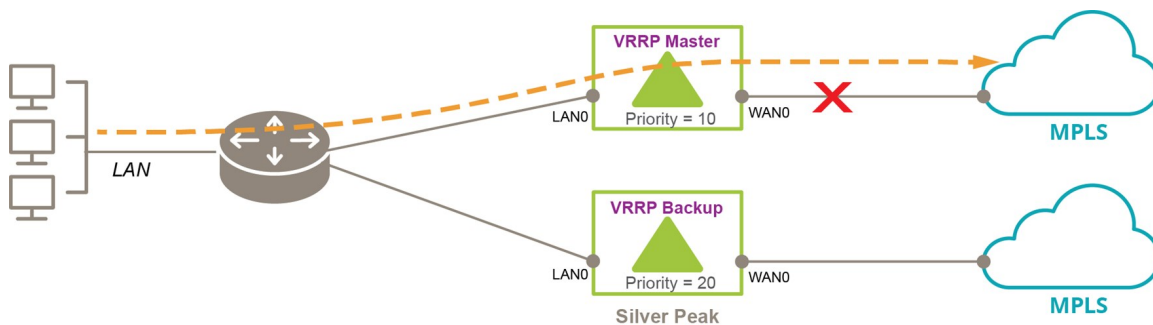
Example #3 – Monitor Interface (LAN0)



- 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.

The screenshot shows the 'IP SLA Rule' configuration window. The 'Monitor' section has 'Interface' set to 'lan0' and 'Interval' set to '30 (Sec)'. The 'Actions' section has 'Down Action' set to 'Disable Subnet Sharing' and 'Up Action' set to 'Default Subnet Action'. There is a 'Comment' field at the bottom.

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.

The screenshot shows the 'IP SLA Rule' configuration window. At the top right is a close button (X). Below the title bar, there are two toggle buttons: 'ON' (highlighted in blue) and 'OFF'. The window is divided into two main sections: 'Monitor' and 'Actions'.

Monitor Section:

- Monitor:** A dropdown menu set to 'Interface'.
- Interface:** A dropdown menu set to 'MPLS'.
- Interval:** A text input field containing '30', followed by '(Sec)'.

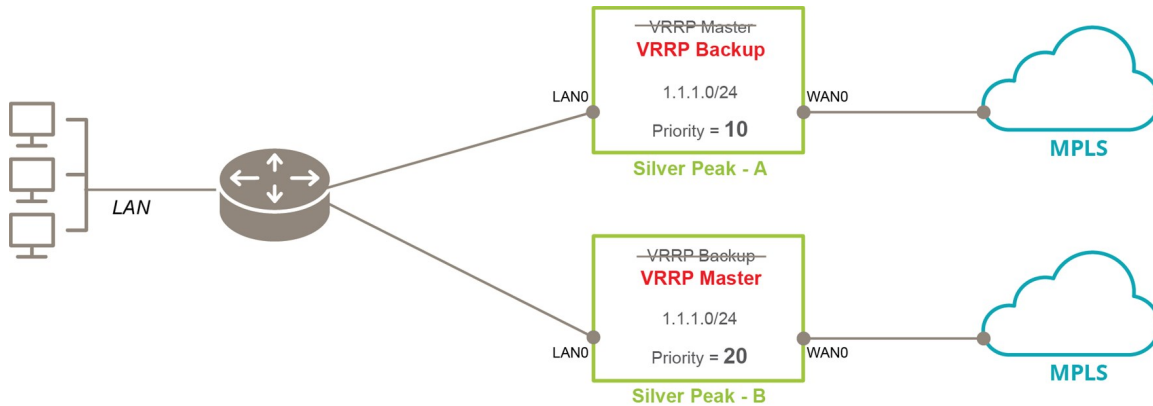
Actions Section:

- Down Action:** A dropdown menu set to 'Decrease VRRP Priority'.
- Interface:** A dropdown menu set to 'lan0'.
- Priority:** A text input field containing '30'.
- Up Action:** A dropdown menu set to 'VRRP Default'.
- Comment:** A text area for additional notes.

At the bottom right of the window are two buttons: 'Add' and 'Close'.

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 is **LAN0**.

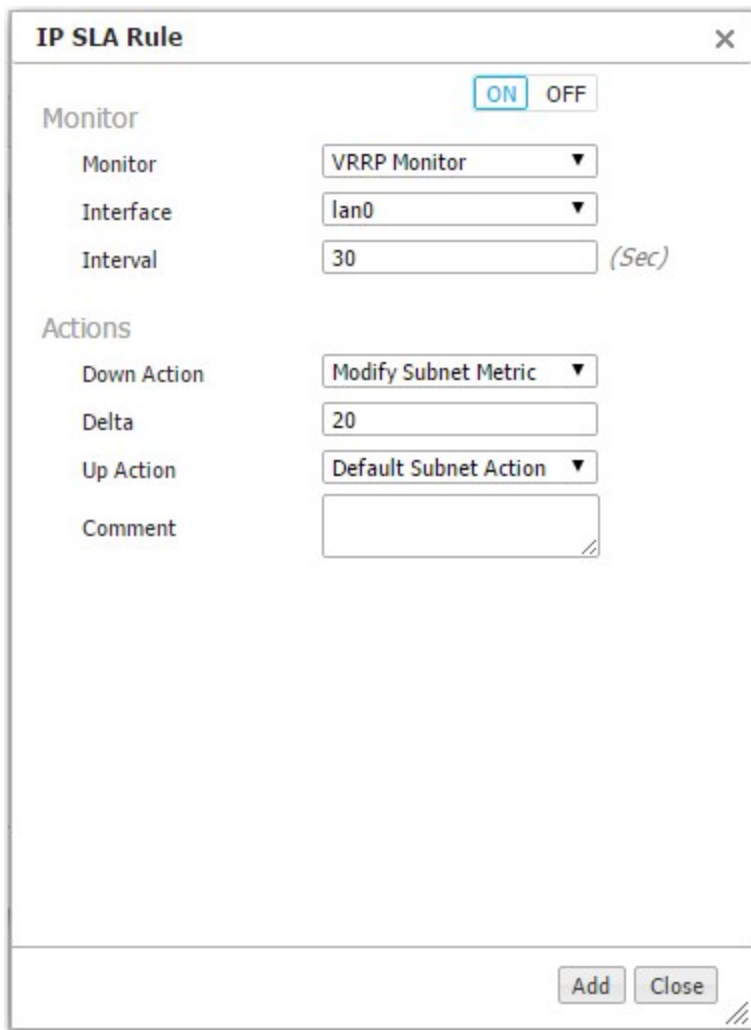
- Here we are looking at an instance where the VRRP role changes, but priority does not, for whatever reason.
- Its **IP SLA Rule** would look like this, with the **Default Subnet Action** being to revert to the original Priority.

The screenshot shows the 'IP SLA Rule' configuration window. The 'Monitor' section has a toggle switch set to 'ON'. Below it, the 'Monitor' dropdown is set to 'VRRP Monitor', the 'Interface' dropdown is set to 'lan0', and the 'Interval' is set to '30 (Sec)'. The 'Actions' section has three dropdowns: 'Down Action' set to 'Disable Subnet Sharing', 'Up Action' set to 'Default Subnet Action', and a 'Comment' field containing the text 'Monitors VRRP router state'. At the bottom right, there are 'Add' and 'Close' buttons.

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 in which 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 window titled "IP SLA Rule" with a close button (X) in the top right corner. The window 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 showing "VRRP Monitor".
- Interface:** A dropdown menu showing "lan0".
- Interval:** A text input field containing "30", followed by the unit "(Sec)".

Actions Section:

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

At the bottom right of the window are two buttons: "Add" and "Close".

Configuration Templates

This section describes templates and how to use them to manage and assign common configuration parameters to appliances.

Templates Overview

CAUTION After saving, templates are applied automatically and replace all settings on an appliance with those configured in the template. Some templates support a MERGE option. Refer to the Help for more information.

- You can edit only a template that appears under Active Templates.
- Click **Show All >** to view available templates that are not part of the selected template group.
- To add a template to Active Templates, double-click it or drag it from Available Templates.
- Click a template under Active Templates to modify it.
- To save the current Active Templates as a new template group, click **Save As**.

Template Groups

A Template Group contains one or more templates you can assign to some or all of the appliances in your network.

- To create a template group, click **+Add** below the template group drop-down list.
- To delete the selected template group, click **-Delete** below the template group drop-down list.
- When you apply a template group to an appliance, Orchestrator automatically keeps the templates in the group in sync with the appliance.
- To apply template groups, 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.
- When returning to the Templates page, Orchestrator displays the last template group viewed.

System Template

Use this page to configure system-level features.

Category	Field	Description
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, Dynamic Path Control determines where to send the flows. There are three options:</p> <ul style="list-style-type: none"> ■ fail-stick – When the failed tunnel comes back up, the flows do not return to the original tunnel. They stay where they are. ■ fail-back – When the failed tunnel comes back up, the flows return to the original tunnel. ■ disable – When the original tunnel fails, the flows are not routed to another tunnel.
Network Memory	Encrypt data on disk	Enables encryption of all the cached data on the disks. Disabling this option is not recommended.
Excess Flow Handling	Excess flow policy	Specifies what happens to flows when the appliance reaches its maximum capacity for optimizing flows. The default is to bypass flows. Or, you can choose to drop the packets.
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	Specifies whether 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 by using the Orchestrator. This activity can apply to the entire distributed network of EdgeConnect appliances or just to a specified group of appliances.
----------------------	---	--

Bridge Loop Test	Only valid for virtual appliances. When enabled, the appliance can detect bridge loops. If it detects a loop, the appliance stops forwarding traffic and raises an alarm. Appliance alarms include recommended actions.
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.
Enable default DNS lookup	Allows the appliance to snoop the DNS requests to map domains to IP addresses. This mapping then can be used in ACLs for traffic matching.
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.
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.
UDP flow timeout	Specifies how long to keep the UDP session open after traffic stops flowing. The default is 120 seconds (2 minutes).
Non-accelerated TCP Flow Timeout	Specifies how long to keep the TCP session open after traffic stops flowing. The default is 1800 seconds (30 minutes).
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.
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 to keep the mappings in the NAT device intact.
Tunnel Alarm Aggregation Threshold	Specifies the number of alarms to allow before alerting the tunnel alarm.
Maintain end-to-end overlay mapping	Enforces the same overlay to be used end-to-end when traffic is forwarded on multiple nodes.
IP Directed Broadcast	Allows an entire network to receive data that only the target subnet initially receives.
Allow WAN to WAN routing	Redirects inbound LAN traffic back to the WAN.

SNMP Template

EdgeConnect appliances support Management Information Base (MIB-II) as described in RFC 1213 for cold start traps, warm start traps, and Silver Peak proprietary MIBs. Appliances issue an SNMP trap during reset when loading a new image, recovering from a crash, or rebooting.

An appliance sends a trap every time an alarm is raised or cleared. Traps contain additional information about alarms, including severity, sequence number, a text-based description of the alarm, and the time the alarm was created. For more information, you can download a .zip archive containing supported MIBs at <https://www.silver-peak.com/download/latest/mibs.html>.

Use this page to configure the appliance's **SNMP** agent and trap receivers.

1. Select the **Enable SNMP** check box to activate configuration options for SNMP v1/v2, SNMP v3, and **Trap Receivers** details.
2. If you select the **Enable SNMP Traps** check box, the SNMP agent on the appliance sends traps to configured receivers.
3. Use the **Default Trap Community** field to specify the string the trap receiver uses to accept traps being sent to it. The default value is **public**. You can modify this value.

SNMP v1/v2

Configure the following fields for SNMP v1 and v2c.

Field	Description
Enable SNMP	Allows the SNMP agent on the appliance to send traps to configured receivers.
Read-Only Community	The SNMP application needs to present this text string (secret) to poll the appliance's SNMP agent. The default value is public . You can modify this value.

SNMP v3

For additional security, configure SNMP v3 if you want to authenticate without using clear text. To add an SNMP v3 user, click **Add** above the SNMP v3 table and configure the following properties:

Field	Description
Enabled	Select this check box to enable the selected user. Clear this check box to disable the user and maintain the configuration.
Username	Enter the username to identify the SNMP v3 user.
Authentication Type	Select the authentication type to use for SNMP requests from the user. NOTE Authentication type is required and SHA-1 is the only supported algorithm.
Authentication Password	Enter a password that the SNMP agent can use to authenticate requests sent by the user. NOTE The password must be at least 20 characters long.
Privacy Type	Select the encryption type to use for encrypting requests from the SNMP user. NOTE Encryption is required, and AES-128 is the only supported algorithm.
Privacy Password	Enter a password (key) to use for encrypting requests sent by the user. NOTE The password must be at least 20 characters long.

NOTE To delete an SNMP v3 user, click the **X** to the right of the entry in the table.

Trap Receivers

To configure a trap receiver, click **Add** above the Trap Receivers table and configure the following properties:

NOTE You can configure up to three trap receivers per appliance.

Field	Description
Host	IP address of the host where traps should be sent.
Version	Select the SNMP version of the trap receiver.
Community/Username	For v1 and v2c, enter the community string the receiver should use to accept traps. If left blank, the default community string (public) is used. If a different community string is configured on the trap receiver, enter it here. For v3, specify the SNMP v3 user that is sending traps to the receiver.
Enabled	Select this check box to enable the receiver. Clear this check box to disable the receiver and maintain the configuration.

NOTE To delete a receiver, click the **X** to the right of the entry in the table.

Flow Export Template

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

- The appliance exports flows against two virtual interfaces—**sp_lan** and **sp_wan**—that accumulate the total of LAN-side and WAN-side traffic, regardless of physical interface.
- These interfaces appear in SNMP and are, therefore, "discoverable" by NetFlow 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 are 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 want. The default is **WAN TX**.

DNS Template

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

Logging Template

Use this template to configure local and remote logging parameters.

Each requires that you specify the minimum severity level of event to log.

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

Minimum Severity Levels

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

EMERGENCY	System is unusable.
ALERT	Includes all alarms the appliance generates: CRITICAL , MAJOR , MINOR , and WARNING .
CRITICAL	Critical event.

ERROR	An error. This is a non-urgent failure.
WARNING	A warning condition. Indicates an error will occur if action is not taken.
NOTICE	A normal, but significant, condition. No immediate action required.
INFORMATIONAL	Informational. Used by Silver Peak for debugging.
DEBUG	Used by Silver Peak for debugging.
NONE	If you select NONE , no events are logged.

- The bolded part of the name is what displays in Silver Peak logs.
- If you select **NOTICE** (the default), 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, after they clear, list as the ALERT level in the **Event Log**.

Configure 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 might not accept as low a severity level as you are forwarding to it.
- In the **Log Facilities Configuration** section, assign each message/event type (System / Audit / Flow) to a syslog facility level (**local0** to **local7**).
- For each remote syslog server that you add to receive the events, specify the receiver's IP address, along with the messages' minimum severity level and facility level.

Banner Messages Template

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

Date/Time Setting

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

1. From the Time Zone list, select the appliance's geographical location.
 - If you select Manual, the appliance is matched to your web client system when the template is applied. This eliminates the delay between configuring time manually and applying the template.

To use an NTP server, select **NTP Time Synchronization** and complete the following steps.

1. Click **Add**.
2. Enter the IP address or host name of the server.
3. Select the version of NTP protocol to use.

NOTE The server is selected in the order listed when you list more than one NTP server.

Data Collection

- 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.

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 shows the 'Templates' configuration window for the 'HTTPS Certificate' template. The left sidebar contains a 'Template Group' dropdown set to 'Default Template Group', with '+Add -Delete' links. Below is a list of 'Active Templates' with a 'Show All >' link. The 'General Settings' section is expanded, showing 'System', 'HTTPS Certificate' (highlighted), 'Policies', 'Shaper', and 'Access Lists'. At the bottom of the sidebar are 'Save', 'Save As', and 'Cancel' buttons, with a note 'Applies to all templates in group' and a link 'Apply Template Groups'. The main content area is titled 'HTTPS Certificate' and has two radio button options: 'Self Signed Certificate (Issuer: Silver Peak)' (selected) and 'Custom Certificate' (with an 'Upload and Replace' button). Below these are labels for 'Issuer', 'Issued to', and 'Expiration'.

To use a custom certificate 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, and so forth.

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

Templates x

Template Group ?
Default Template Group ▼
+Add -Delete

Active Templates Show All >

General Settings
System
User Management
Policies
Shaper
Access Lists

Save Save As Cancel
Applies to all templates in group

Apply Template Groups

User Management ?

User Accounts
Add

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 user accounts, **admin** and **monitor**, that cannot be deleted.
- You can, however, assign a new password to either one and apply it to any appliances you want.

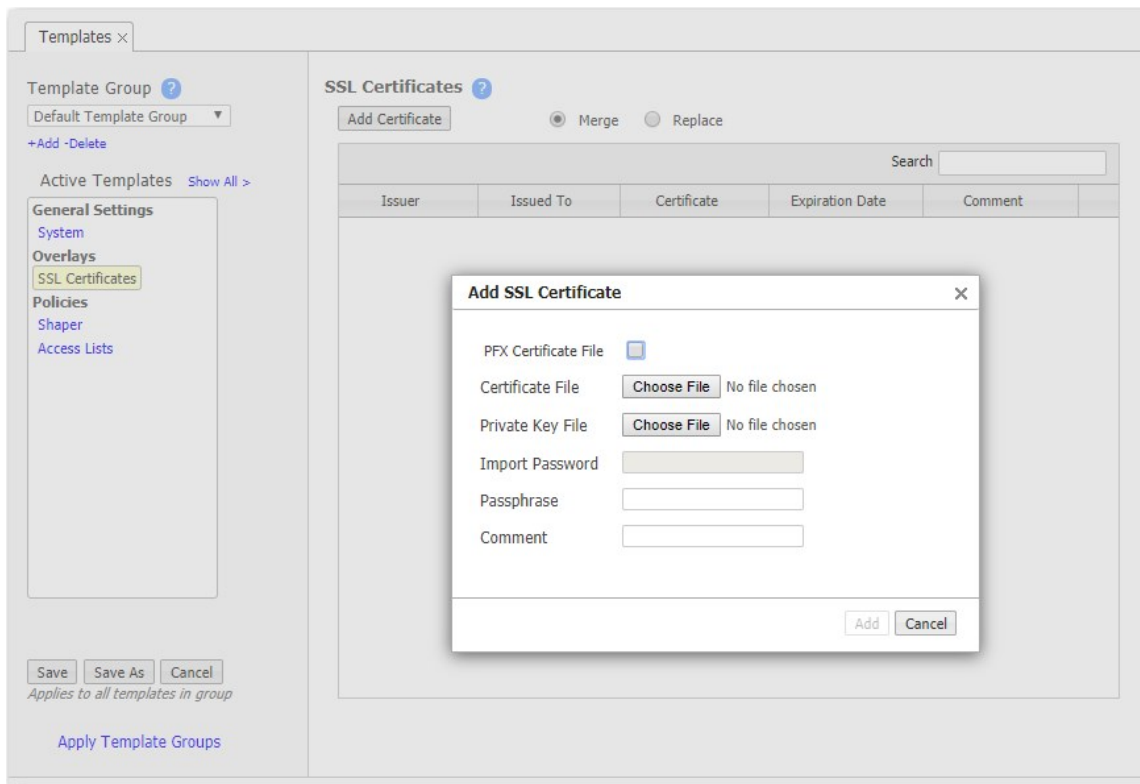
Command Line Interface Privileges

- 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 an EdgeConnect 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.

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.



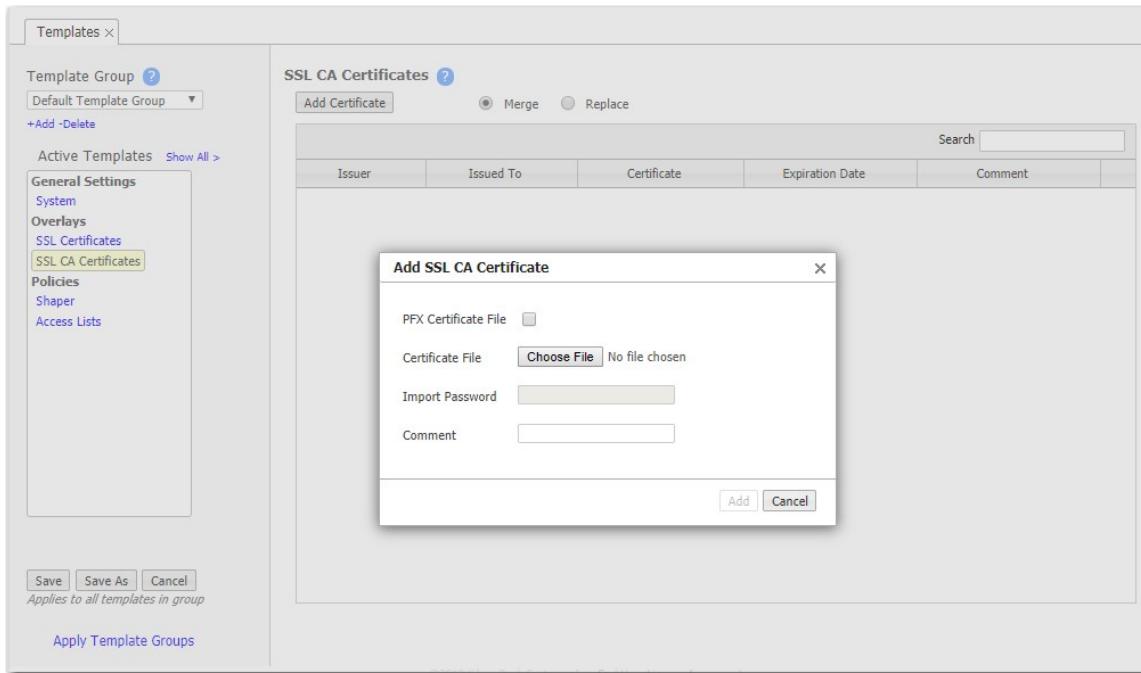
By supporting the use of SSL certificates and keys, Silver Peak provides deduplication for Secure Socket Layer (SSL) encrypted WAN traffic:

- Silver Peak decrypts SSL data using the configured certificates and keys, optimizes the data, and transmits data over an IPSec tunnel. The peer EdgeConnect 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 > Networking > Tunnels > 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 > Templates & Policies > Optimization Policies** page, and then review the **Set Actions**.
- If you choose to be able to decrypt the flow, optimize it, and send it in the clear between appliances, 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 you used for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, you must add those CA certificates here. If the browser cannot validate up the chain to the root CA, it will warn you that it cannot trust the certificate.

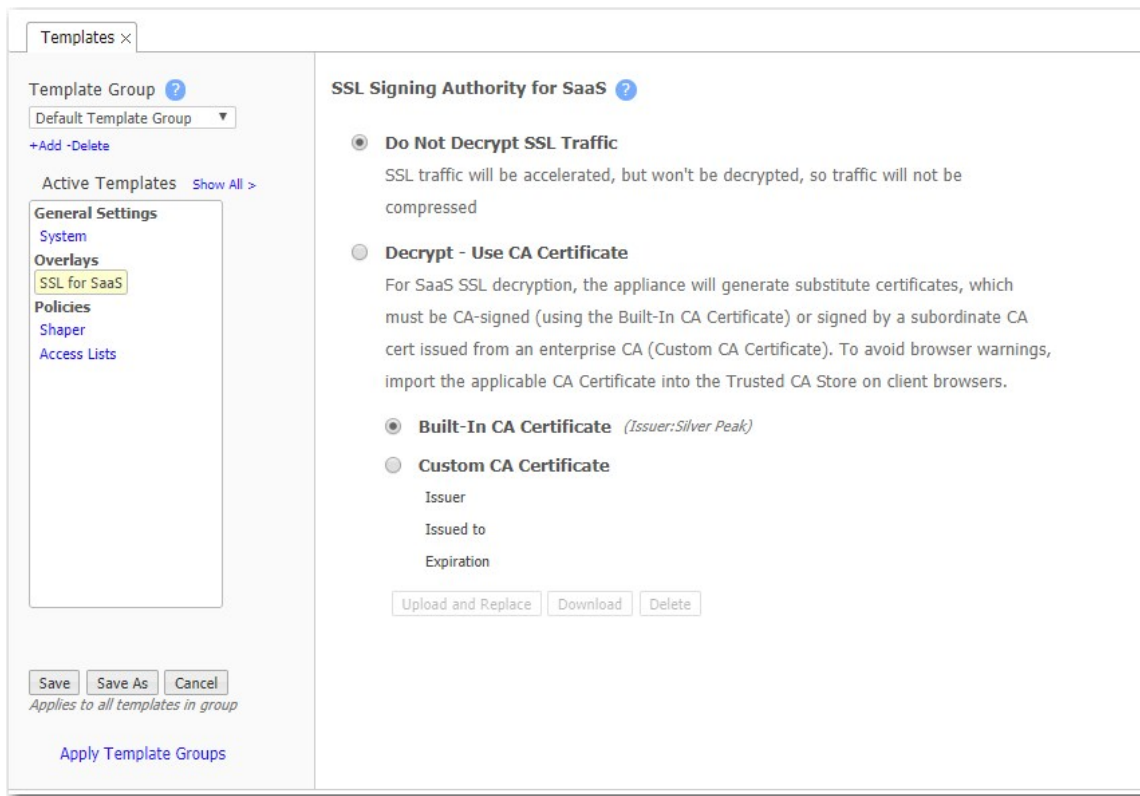


TIP For a historical matrix of Silver Peak security algorithms, click [here](#).

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 then must 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, 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 **do not** already have a subordinate CA certificate, you can access any appliance's **Configuration > Templates & Policies > Applications & SaaS > SaaS Optimization** page and generate a Certificate Signing Request (CSR).

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

Auth/Radius/TACACS+ Template

EdgeConnect 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, EdgeConnect appliances:

- Support a built-in, **local database**.
- Can be linked to a **RADIUS** (Remote Authentication 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. Refer to 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. Refer to 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

DNS Proxy Policies

Configuration > Templates & Policies > Templates

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

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

1. Choose whether 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 whether 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.

Tunnels Template

NOTE If you are deploying an SD-WAN network, the Business Intent Overlays (BIOs) govern tunnel properties. In this case, you do not need this template.

If you are not creating overlays, 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**.
- To **view**, **edit**, and **delete** tunnels, use the **Tunnels** tab. The **Mode** selected determines the tabs that display.

Tunnel Settings ? ×

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

Save
Close

Tunnels Template Settings

Field	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.

Tunnels Template Settings

Field	Description
DSCP	Determines the DSCP marking that the keep-alive messages should use.
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 brownout is:</p> $T_{wait} = Base + N * RTT_{avg}$ <p>where Base is a value in milliseconds, and N 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 brownout if it does not see a reply packet from the remote end within 300ms of receiving the most recent packet.</p> <p>In the Tunnel Advanced Options, Base is expressed as Fastfail Wait-time Base Offset (ms), and N is expressed as Fastfail RTT Multiplication Factor.</p> <ul style="list-style-type: none"> ■ Fastfail Enabled – This option is triggered when a tunnel's keepalive signal doesn't receive a reply. The options are disable, enable, and continuous. If the disqualified tunnel subsequently receives a keepalive reply, its recovery is instantaneous. <ul style="list-style-type: none"> • If set to disable, keepalives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost. • If set to enable, keepalives are sent every second, and a missed reply increases the rate at which keepalives are sent from one per second to ten per second. Failover occurs after one second. • When set to continuous, keepalives are continuously sent at ten per second. Therefore, failover occurs after one tenth of a second. ■ Thresholds for Latency, Loss, or Jitter are checked once every second. <ul style="list-style-type: none"> • Receiving three 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 three successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.
FEC	(Forward Error Correction) can be set to enable , disable , and auto .

Tunnels Template Settings

Field	Description
FEC Ratio	Is an option when FEC is set to auto that specifies the maximum ratio. The options are 1:2, 1:5, 1:10, or 1:20.
IPSec Anti-replay window	Provides protection against an attacker duplicating encrypted packets by assigning a unique sequence number to each encrypted packet. The decryptor keeps track of which packets it has seen on the basis of these numbers. The default window size is 64 packets.
IPSec Preshared Key	A shared, secret string of Unicode characters that is used for authentication of an IPSec connection between two parties.
Mode	Indicates whether the tunnel protocol is udp , gre , or ipsec .
MTU (bytes)	Maximum Transmission Unit (MTU) 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.
Reorder Wait (ms)	Maximum time the appliance holds an out-of-order packet when attempting to reorder. The 100ms default value should be adequate for most situations. FEC can introduce out-of-order packets if the reorder wait time is not set high enough.
Retry Count	Number of failed keep-alive messages that are allowed before the appliance brings the tunnel down.
UDP destination port	Used in UDP mode. Accept the default value unless the port is blocked by a firewall.
UDP flows	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 'Templates' configuration window in the Silver Peak Unity Orchestrator. The window is divided into two main sections. On the left, under the 'Template Group' tab, there is a dropdown menu for 'Default Template Group' and a '+Add -Delete' button. Below this is a list of 'Active Templates' with a 'Show All >' link. A sidebar on the left lists navigation options: 'General Settings', 'System', 'Networking' (which is expanded to show 'VRRP', 'Policies', 'Shaper', and 'Access Lists'), 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'. The right section is titled 'VRRP' and contains configuration fields: 'Admin' with a dropdown set to 'Up', 'Advertisement Timer' with a text input field and a range '(1..255)', 'Priority' with a text input field and a range '(1..254)', 'Preemption' with a checkbox, and 'Authentication String' with a text input field.

In an out-of-path deployment, one method for redirecting traffic to the EdgeConnect appliance is to configure VRRP on a common virtual interface. 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 in which 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	Description
Admin	Options are up (enable) and down (disable).
Advertisement Timer	Default is 1 second .
Authentication String	Clear text password for authenticating group members.
Preemption	Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
Priority	The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.

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, the appliance randomly distributes flows among multiple peers.
- The lower the number, the higher the peer's priority.

NOTE This feature requires appliance software 8.3.3.0 or higher for version 8 releases, and requires 9.0.2.0 or higher for version 9 releases.

Templates ×

Template Group ?

Default Template Group ▼

+Add -Delete

Active Templates Show All >

General Settings

System

Networking

Peer Priority

Policies

Shaper

Access Lists

Save

Save As

Cancel

Applies to all templates in group

Apply Template Groups

Peer Priority ?

Add Peer

1 Rows, 1 Selected Search

Peer Name	Priority ▼	
Type to select	0	×

Admin Distance Template

This table shows 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.

Field	Description
Local	A manually configured route, or one learned from locally-connected subnets.
Subnet Shared - Static Routes	A route learned from a Silver Peak peer.
Subnet Shared - BGP Remote	A route shared from a Silver Peak peer from an external network.
Subnet Shared - OSPF Remote	A route shared from a Silver Peak peer within the same network.

Field	Description
BGP Branch (pre-8.1.9.4)	A type of dynamic route learned from a local BGP branch peer before 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 before 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 PE (pre-8.1.9.4)	A type of dynamic route learned from a local BGP PE (Provider Edge) router before version 8.1.9.4.
OSPF	A route learned from an OSPF (Open Shortest Path First) neighbor.
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.

Route Redistribution Template

To use this template, you must have your route maps configured for either SD-WAN, BGP, and OSPF. See the **Routes** tab for more details about the configuration and defining rules for your route maps.

Merge and Replace

If you select **Merge**, new maps are added to the existing maps. If the map already exists, the new map will match appliance rules in the orchestrator range. If the configured rules do not match, the new map's rules are appended to the existing rules. **Replace** will take the new maps and replace all existing maps and not include the rules that match outside of the configured range.

Complete the following steps to redistribute a route map.

1. Select the direction of traffic you want to redistribute your routes to: **SD-WAN Fabric**, **BGP Inbound** and **Outbound**, and **OSPF**.
2. When selected, click **Add Map**.
3. Enter a **Map Name**, and then click **Add**.
4. Select **Add Rule**. The **Add Rule** window opens.

In this window, you define the rules applied to your route map, which includes the **Match Criteria** and the **Set Actions**. Each route map has a **match** command and **set** command. The match command verifies the attributes of the original route the protocol supports. The set command modifies information that is redistributed into the target protocol.

NOTE You can apply 128 rules per map.

5. Click **Add**.

Shaper Template

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

- The Shaper shapes traffic by allocating bandwidth as a percentage of the **system bandwidth**.
- The Shaper's parameters are organized into ten traffic classes. Four traffic classes are preconfigured and named **real-time**, **interactive**, **default**, and **best effort**.
- The system applies these QoS settings globally after compressing (deduplicating) all the outbound tunnelized and pass-through-shaped traffic, shaping it as it exits to the WAN.
- Applying the template to an appliance updates its system-level **wan** Shaper. If the appliance has any added, interface-specific Shapers, they are preserved.
- 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 Templates & Policies > Shaping > Shaper** page.

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

Save Save As Cancel


Applies to all templates in group

Apply Template Groups

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**. 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	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"> ■ For WAN optimization, the interface shaper can be used, but it is not recommended. ■ For SD-WAN, it should never be used because overlay traffic is not directed to an interface shaper; traffic is always shaped by the default WAN shaper.
Excess Weighting	If there is bandwidth left over after satisfying the minimum bandwidth percentages, the excess is distributed among the traffic classes in proportion to the weightings specified in the Excess Weighting column. Values range from 1 to 10,000.
Interface Shaper	Interface that 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 Max Wait Time are dropped.

Field	Description
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%, lower-priority traffic classes might not receive their guaranteed bandwidth if it is all consumed by higher-priority traffic. If you set Min Bandwidth to a value greater than Max Bandwidth , then Max overrides Min .
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 - 1 is first, 10 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 0 (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 wan0 goes down, wan0 bandwidth is removed from the total bandwidth when recalculating.
Traffic Name	Name assigned to a traffic class, either prescriptively or by the user.

QoS Policies Template

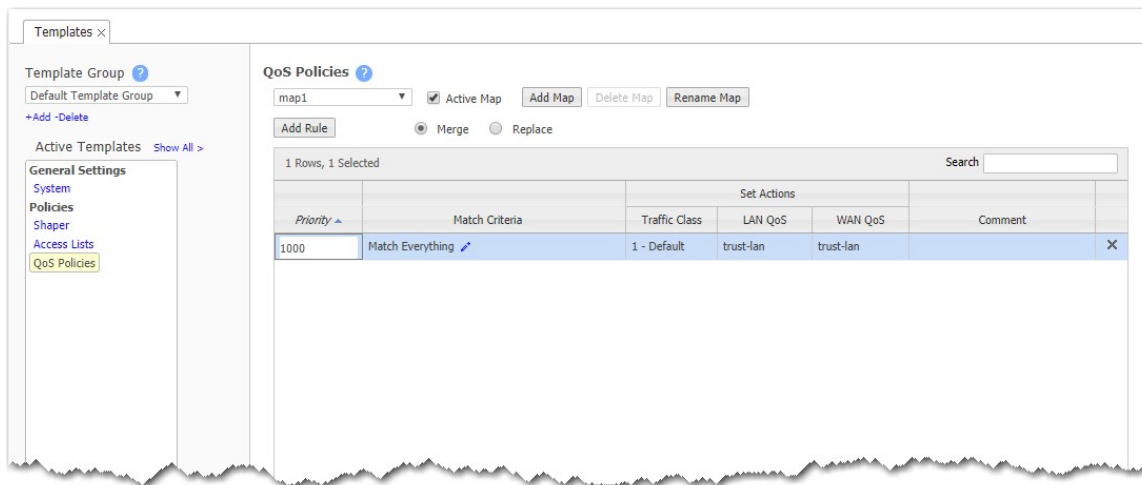
QoS Policy determines how flows are queued and marked.

The QoS Policy's SET actions determine two things:

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

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

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



Priority

- With this template, you can create rules with a priority from **1000 – 9999**. When the template is applied to an appliance, Orchestrator will delete all rules having a priority in that range before applying its policies.
- If you access an appliance directly, you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

NOTE The priority range from **20000** to **24999** is reserved for Orchestrator.

- When adding a rule, the priority is incremented by ten from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

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 & Policies > ACLs > 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** check box.

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.

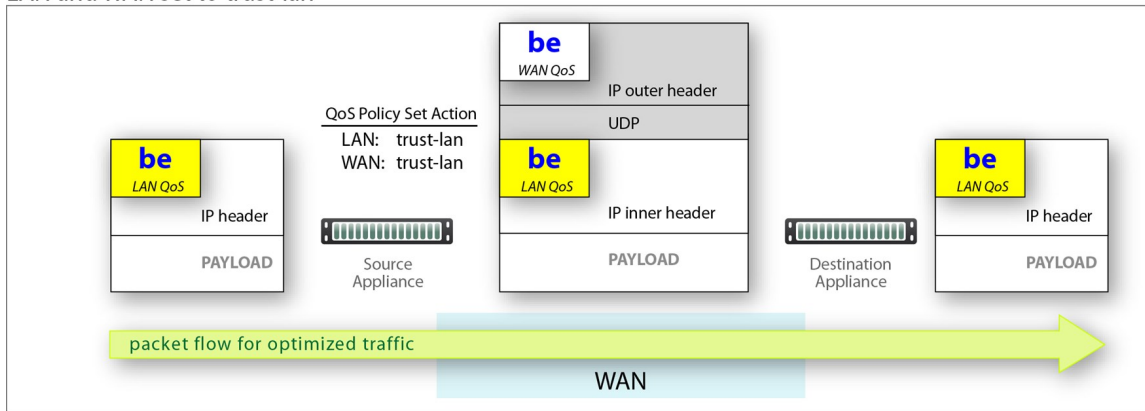
Handle and Mark DSCP Packets

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

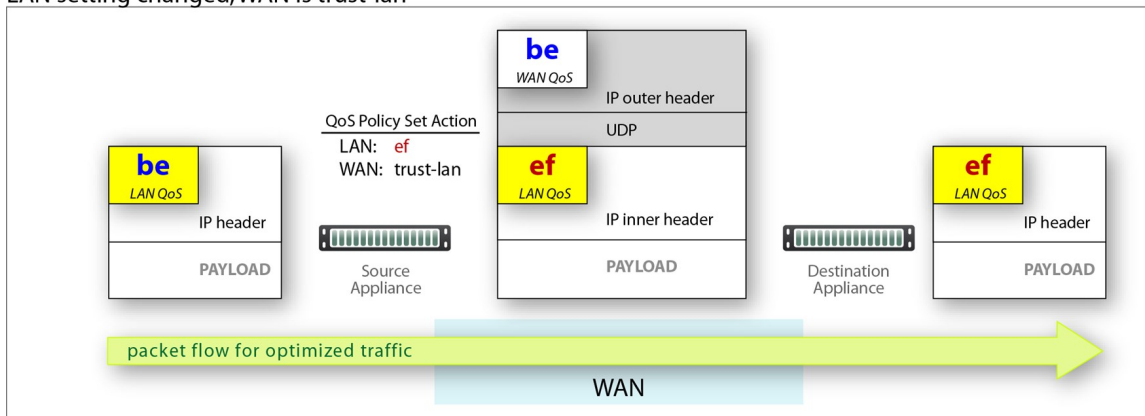
Apply DSCP Markings to Optimized (Tunnelized) Traffic

- 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.

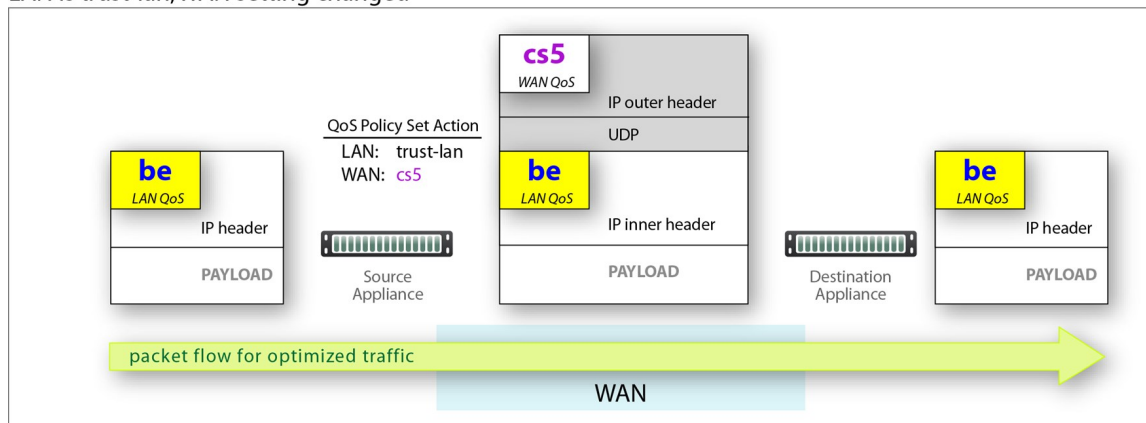
LAN and WAN set to trust-lan



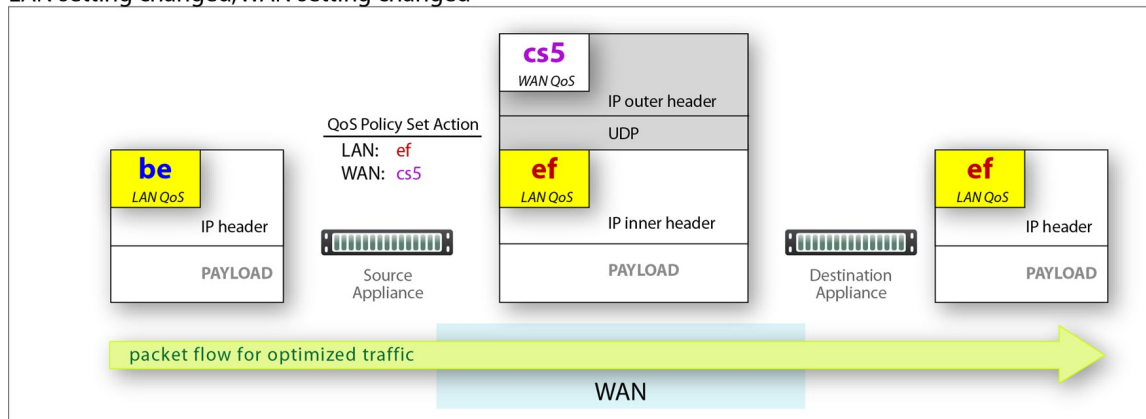
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



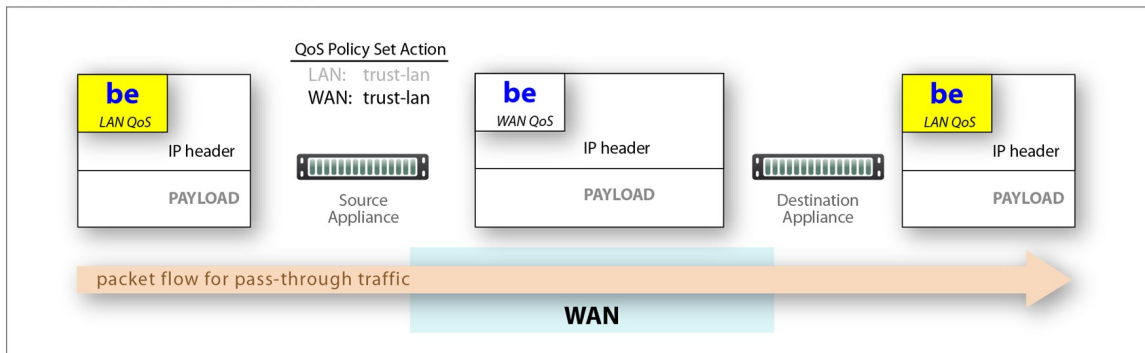
LAN setting changed, WAN setting changed



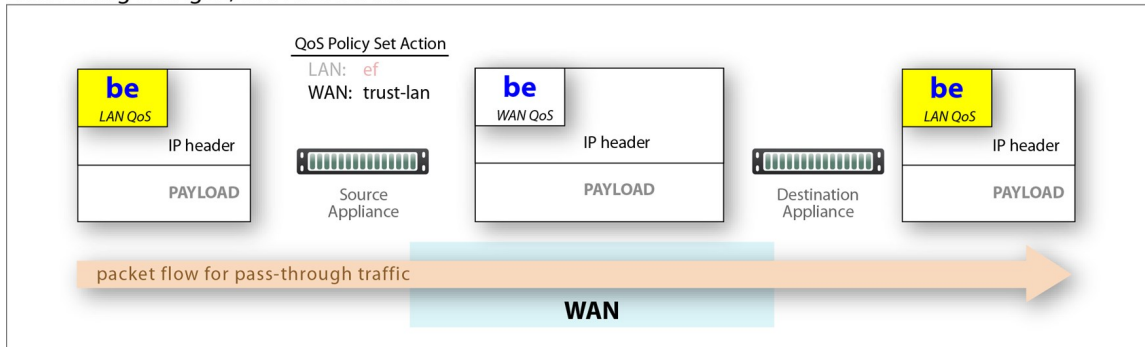
Apply DSCP Markings to Pass-through Traffic

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

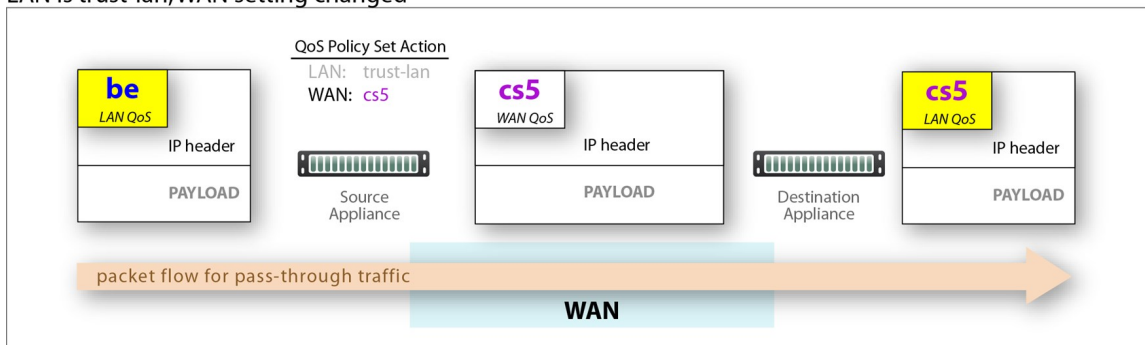
LAN and WAN set to trust-lan



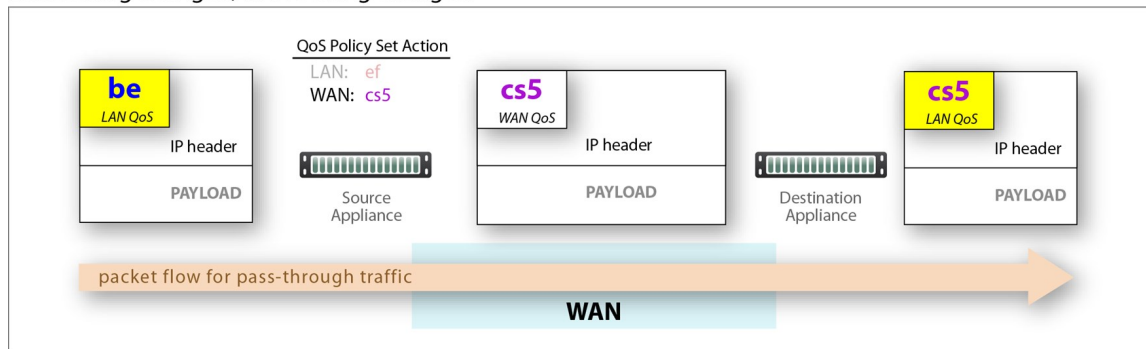
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



LAN setting changed, WAN setting changed



Routes Template

Select the following check boxes if you want to globally apply them to your routes in Orchestrator.

- **Automatically advertise to local LAN subnets:** Enable if you want the system created LAN subnets of your appliance advertised to your peers.
- **Automatically advertised local WAN subnets:** Enable if you want the system created local WAN subnets of your appliance advertised to your peers.
- **Redistribute learned BGP routes to Silver Peak Peers:** Advertises BGP routes that your appliance has learned to Silver Peak peers.

Enter specific values for the following:

Field	Description
Metric for automatically added routes	50 (default value).
Route Map name to Redistribute route to SD-WAN Fabric	Name of the route map being redistributed to the SD-WAN.
Include BGP Local ASN to routes sent to SD-WAN Fabric *	Select: Don't Apply , Yes , or No .
Filter Routes From SD-WAN Fabric with Matching Local ASN	Select: Don't Apply , Yes , or No .

Field	Description
Tag BGP communities to routes	<p>Apply these communities to locally configured static routes and SD-WAN origin routes (remote local, remote eBGP, and remote OSPF).</p> <p>Routes tagged with BGP community are also advertised to:</p> <ul style="list-style-type: none"> LAN-side BGP peers SD-WAN peers <p>Select: Don't Apply, Yes, or No.</p> <p>If Yes is selected, enter the BGP communities you want to be tagged in the field.</p> <hr/> <p>NOTE A community must be a combination of two numbers (0 to 65535) separated by a colon. For multiple communities, use a comma to separate them.</p>

* SD-WAN fabric is sometimes referred to as *subnet sharing*.

INFO If you select **Don't apply**, Orchestrator ignores this field when applying this template to appliances.

BGP Template

Use the BGP template to apply BGP configurations per segment to all appliances in the SD-WAN fabric.

- Click the edit icon next to the segment for which you want to modify the configuration.
- Configure the following elements as needed:

Field	Description
AS Path Propagate	Select Yes to enable this appliance to send the full AS path associated with a prefix to other routers and appliances, avoiding routing loops. This will provide the learned path from an external prepend between a remote BGP site to local BGP peers.
Graceful Restart	Select Yes to enable receiver-side graceful restart capability. Silver Peak retains routes learned from the peer and continues to use them for forwarding if a BGP peer goes down. Retained routes are considered stale routes. They will be deleted and replaced when new routes are received.
Max Restart Time	If Graceful Restart is enabled, specifies the maximum time in seconds to wait for a capable peer to come back after a restart or peer session failure.
Stale Path Time	If Graceful Restart is enabled, specifies the maximum time in seconds following a peer restart before removing stale routes associated with a peer.

Field	Description
Next-Hop-Self	Advertised route connected to a CE router that an EdgeConnect appliance learns from a PE router.
Keep Alive Timer	This is the interval, in seconds, between keep alive signals to a peer.
Hold Timer	When availability to a peer is lost, this value specifies how long to wait before dropping the session.
Enable MD5 Password	If applied, adds a password to authenticate TCP sessions with peers.
Password / Confirm Password	If the MD5 password is enabled, use these fields to specify the password.

3. Click **Update**.

OSPF Template

Use the OSPF template to apply OSPF configurations per segment to all appliances in the SD-WAN fabric.

1. Click the edit icon next to the segment for which you want to modify the configuration.
2. Configure the following elements as needed:

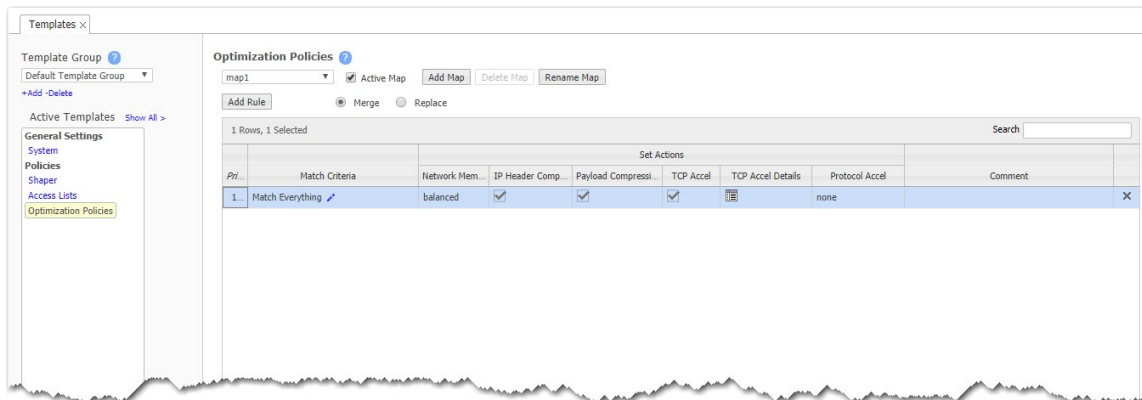
Field	Description
Enable OSPF	indicates whether the segment can access OSPF protocol. If you select Don't apply , Orchestrator ignores this field when applying this template to appliances.
Route Map name to Redistribute routes to OSPF	<p>Name of the route map being redistributed to the SD-WAN.</p> <p>The OSPF template is used in conjunction with the Route Redistribution Maps template. OSPF route maps are configured in the Route Redistribution Maps template, and then applied in the OSPF template. The default OSPF route map name is "default_rmap_to_ospf".</p> <hr/> <p>NOTE Leave this field blank to preserve the current setting on the appliance.</p>
Admin Status	Indicates whether the interface admin status is up or down. If you select Don't apply , Orchestrator ignores this field when applying this template to appliances.
Hello Interval	Length of time (in seconds) that must transpire between hello packets that a router sends on an OSPF interface.
Dead Interval	Length of time (in seconds) that must transpire before neighbors that have not detected a router's hello packets can declare the OSPF router down.

Field	Description
Transmit Delay	Length of time (in seconds) that must transpire before transmitting a link state update packet. Specify a value from 1 to 65535.
Retransmit Interval	Length of time (in seconds) that a router that has received no acknowledgment must wait before resending transmissions.
Authentication Type	Type of authentication to use for requests. Select one of the following drop-down list options: <ul style="list-style-type: none"> ■ Don't apply – Orchestrator ignores this field when applying this template to appliances. ■ None – Authentication not performed. ■ Text – Simple password authentication, which allows a key (password) to be configured per area. ■ MD5 – Message Digest cryptographic authentication. A key ID and key (password) are configured on each router. The router uses an algorithm based on the OSPF packet, the key ID, and the key to generate a message digest that gets appended to the packet.
Authentication Key	Key (password) to use for authentication of requests. This field is available only if Authentication Type is set to Text.
MD5 Key	Key ID to use for MD5 authentication of requests. This field is available only if Authentication Type is set to MD5.
MD5 Password / MD5 Confirm Password	Password for the MD5 key. These fields are available only if Authentication Type is set to MD5. Specify and confirm the password.

3. Click **Update**.

Optimization Policies Template

Optimization templates apply Optimization policies to appliances.



Priority

- With this template, you can create rules with a priority from **1000 – 9999**. When the template is applied to an appliance, Orchestrator will delete all rules having a priority in that range before applying its policies.
- If you access an appliance directly, you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

NOTE The priority range from **20000** to **24999** is reserved for Orchestrator.

- When adding a rule, the priority is incremented by ten from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

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 & Policies > ACLs > 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** check box.

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	Description
Network Memory	<p>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> <ul style="list-style-type: none"> ■ Maximize Reduction – Optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP, where bandwidth savings are the primary concern. ■ Minimize Latency – Ensures that Network Memory processing adds no latency. This might come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It is also appropriate when the primary objective is to fully utilize the WAN pipe to increase the LAN-side throughput, as opposed to conserving WAN bandwidth. ■ Balanced – Is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types. ■ Disabled – Turns off Network Memory.
IP Header Compression	<p>Process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end. It is 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.</p>
Payload Compression	<p>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.</p>

Set Action	Description
TCP Acceleration	<p>Uses techniques such as selective acknowledgments, window scaling, and maximum segment size adjustment to mitigate poor performance on high-latency links.</p> <hr/> <p>INFO Slow LAN alert goes off when the loss has fallen below 80% of the specified value configured in the TCP Accel Options window.</p> <hr/> <p>For more information, see TCP Acceleration Options.</p>
Protocol Acceleration	<p>Provides explicit configuration for optimizing CIFS, SSL, SRDF, Citrix, and iSCSI protocols. In a network environment, it is possible that not every appliance has the same optimization configurations enabled. Therefore, the site that initiates the flow (the client) determines the state of the protocol-specific optimization.</p>

Route Policies Template

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

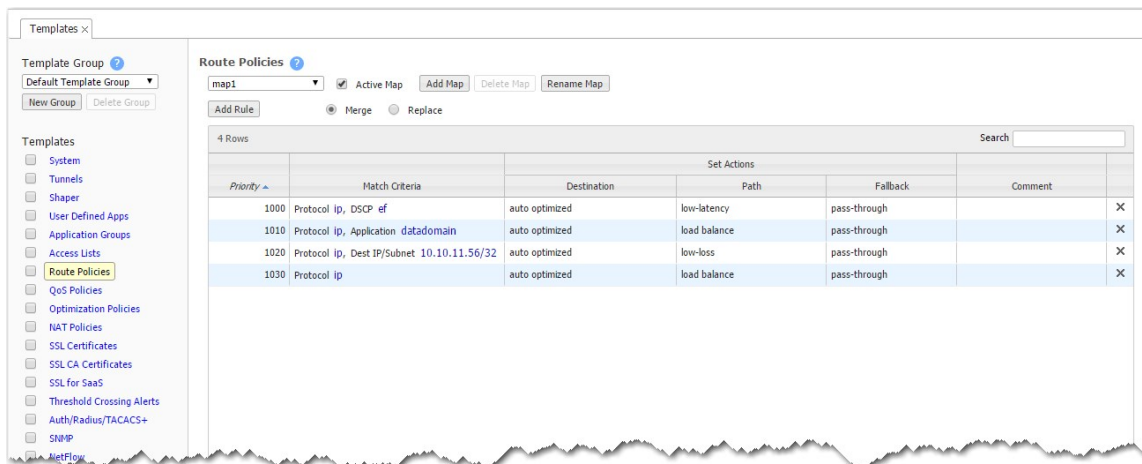
If you are creating a conventional WAN optimization network, there might be occasions when you need to directly configure Route Policies. Then, the following applies.

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

- Load balancing
- Lowest loss
- Lowest latency
- A preferred interface
- A specific tunnel



Why?

Each appliance's default routing behavior is to auto-optimize all IP traffic, automatically directing flows to the appropriate tunnel. **Auto-optimization** strategies reduce the need to create explicit route map entries for optimization. The three strategies 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 a priority from **1000 – 9999**. When the template is applied to an appliance, Orchestrator will delete all rules having a priority in that range before applying its policies.
- If you access an appliance directly, you can create rules with higher priority than Orchestrator rules (**1 – 999**) and rules with lower priority (**10000 – 19999** and **25000 – 65534**).

NOTE The priority range from **20000** to **24999** is reserved for Orchestrator.

- When adding a rule, the priority is incremented by ten from the previous rule. The priority can be changed, but this default behavior helps to ensure you can insert new rules without having to change subsequent priorities.

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 & Policies > ACLs > 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** check box.

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 determine 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 window. On the left, a sidebar contains a 'Template Group' dropdown set to 'Default Template Group', an 'Add +Delete' button, and a list of 'Active Templates' including 'System', 'Policies', 'Shaper', 'Access Lists', and 'NAT Policies' (which is highlighted). Below this is a 'General Settings' section with 'Save', 'Save As', and 'Cancel' buttons, and a note 'Applies to all templates in group' with an 'Apply Template Groups' button.

The main area is titled 'NAT Policies' and contains several sections:


- NAT All Inbound** and **NAT All Outbound** sections, each with checkboxes for 'NAT IP' (set to 'auto') and 'Fallback'.
- Advanced Settings** section with a dropdown menu set to 'map1', a checked 'Active Map' checkbox, and buttons for 'Add Map', 'Delete Map', and 'Rename Map'.
- An 'Add Rule' button and radio buttons for 'Merge' (selected) and 'Replace'.
- A table with 1 row and 7 columns: Priority, Match Criteria, NAT Type, NAT Direction, NAT IP, Fallback, and Comment. The first row has values: 1000, Protocol ip, no-nat, none, auto, and an unchecked Fallback checkbox. A search bar is located at the top right of the table.

When to NAT

Two use cases illustrate the need for NAT:

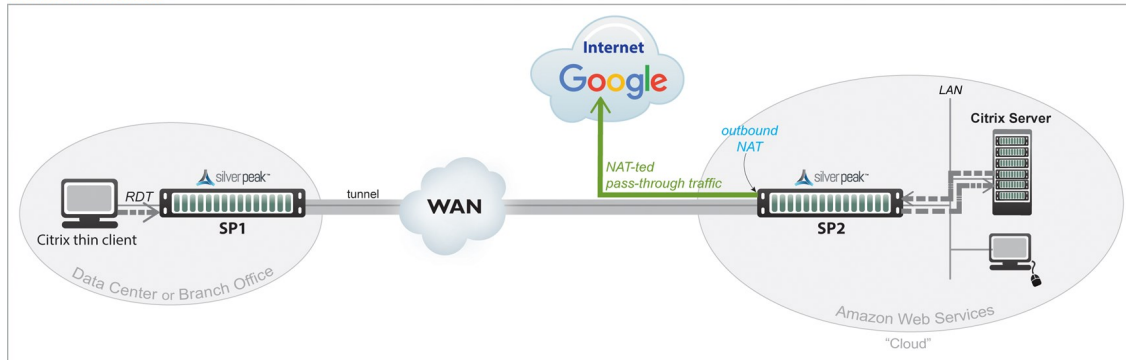
1. **Inbound NAT.** The appliance automatically creates a source NAT (Network Address Translation) map when retrieving subnet information from the Silver Peak Cloud portal. This ensures that traffic destined to SaaS servers has a return path to the appliance from which that traffic originated.

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 does not 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 does not 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 & Policies > ACLs > 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** check box.

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	Description
NAT Type	no-nat	Is the <i>default</i> . No IP addresses are changed.
	source-nat	Is the <i>default</i> . No IP addresses are changed.
NAT Direction	inbound	NAT is on the LAN interface.
	outbound	NAT is on the WAN interface.
	none	Only option if the NAT Type is no-nat .
NAT IP	auto	Select if you want to NAT all traffic. The appliance then picks the first available NAT IP/Port.
	tunnel	Select if you only want to NAT tunnel traffic. Applicable only for inbound NAT, as outbound does not support NAT on tunnel traffic.
	[IP address]	Select if you want to make NAT use this IP address during address translation.
Fallback		If the IP address is full, the appliance uses the next available IP address.

When you select a specific IP, 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 that are triggered when the specific thresholds are crossed.

Templates x

Template Group ?

Default Template Group

+ Add - Delete

Active Templates

Show All >

General Settings

System

Policies

Shaper

Access Lists

Threshold Crossing Alerts

Save Save As Cancel

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 (that is, floor and ceiling levels). For both levels, one value raises the alarm while another value clears it.

**Rules:**

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

ON by Default

- **Appliance Capacity** – Triggers when an appliance reaches 95% of its total flow capacity. It is not configurable and can be cleared only 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)

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

TCA Name	Unit	Metric
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 EdgeConnect appliance must be registered with an **Account Name** and **Account Key** for the SaaS optimization feature.

SaaS Optimization

Enable SaaS Optimization ☒

RTT Calculation Interval (1..1440) minutes

RTT Ping Interface

52 Rows Search

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	eloquatraincenter.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

Save Save As Cancel

Applies to all templates in group

Apply Template Groups

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

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

- If **Advertise** is **selected** for a service (for example, SFDC), the appliance will:
 - Ping active SaaS subnets to determine RTT/metric
 - Add subnet sharing entries locally for subnets within RTT threshold
 - Advertise subnets and their metric (within threshold) via subnet sharing to client-side appliances
 - Upon seeing an SFDC flow, generate a substitute certificate for an SFDC SSL domain (one substitute certificate per domain)
 - Auto-generate dynamic NAT rules for SFDC (but not for unchecked services)
- When **Optimize** is **selected** for a service (for example, SFDC), the appliance will:
 - Ping active SFDC subnets to determine the RTT (metric)
 - Does not advertise metric via subnet sharing (unless **Advertise** is also selected)
 - Receives subnet sharing metric (RTT) from associated appliances
 - Compares its own RTT (local metric) with advertised metric
 - If its own RTT is lower, then the packet is sent pass-through (direct to the SaaS server).
 - If an advertised RTT is lower, then the packet is tunnelized.
 - Generate a substitute certificate for an SFDC SSL domain (one sub cert per domain)
 - No NAT rules created
- When **Optimize** is **not selected** for a service (for example, SFDC), the appliance:
 - Receives subnet sharing advertisements for SFDC but does not 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 might 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 might 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*.

CAUTION If segmentation is enabled, do not use the Security Policies Template. Instead, configure Security Policies from the Routing Segmentation (VRF) tab.

- 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 also can assign a zone label to an **Overlay**. On a 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 enables 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 & Security > Security > 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.
4. You can select **More Options** to customize your rules. Select the check box next to the specific match criteria and select your desired changes from the list.
5. Click **Save**.

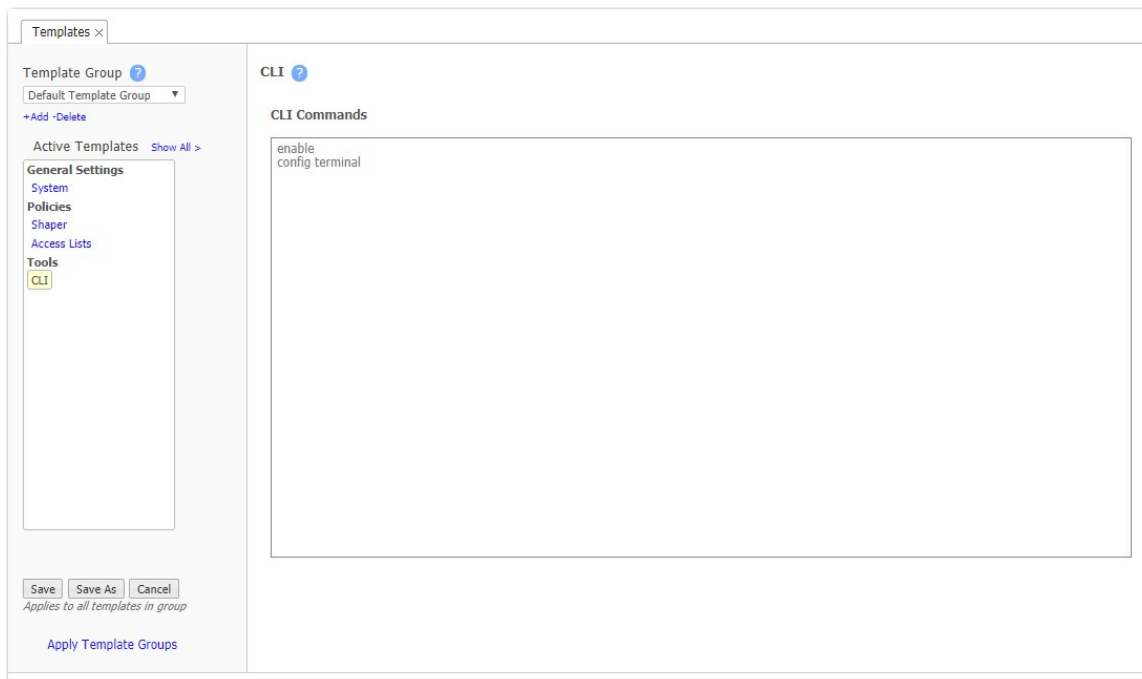
Wildcard-based Prefix Matching

- 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 settings that control access to the appliance web UI.

Templates ×

Template Group ?
Default Template Group ▾
+Add -Delete

Active Templates Show All >

General Settings
System
Policies
Shaper
Access Lists
Tools
Session Management

Save Save As Cancel
Applies to all templates in group

Apply Template Groups

Session Management ?

Auto Logout 60 (60 minutes, 0 indicates no timeout)

Max Session 50 (5-50)

OpenSSL Cipher List Type to apply Leave blank to preserve existing setting on appliances

Web Protocol ☒ HTTPS ☐ HTTP ☐ Both

Field	Description
Auto Logout	Specifies the amount of time in minutes after which an inactive session will be automatically logged out. The valid range is 0-60. Use 0 to disable automatic logout.
Max Sessions	Maximum number of active sessions on the appliance. If the maximum number of sessions is reached, users who try to log in to the appliance web UI will receive a message that the browser cannot access the appliance. On non-EdgeConnect appliance models, Orchestrator might not be able to access the appliance.
OpenSSL Cipher List	<p>List of cipher suites to enable or disable on the appliance. For details about formatting this string, see this page.</p> <p>The string can only contain the following characters: a-z, A-Z, 0-9, and +-.!_@</p> <p>WARNING: Cipher format and availability are not validated. Ciphers should be thoroughly tested in a lab environment before being applied. When ciphers are applied from a template, an improperly formatted string or unavailable ciphers can cause an appliance crash.</p>
Web Protocol	Select the web protocol to use for appliance UI sessions. HTTPS is recommended for maximum security.

Management Services Template

Use this template to globally apply the modifications made to your Management Services if segmentation is enabled or disabled. **Any** is used as the default Interface for the Source IP address; however, you can change the interface with any interfaces you have previously configured on the **Management Services** tab. To modify the interface, click **Any** in the table. For more information, refer to the **Management Services** tab.

Apply Template Groups

Configuration > Templates & Policies > Apply Template Groups

Use this tab to add or remove templates from appliances.

The screenshot shows the 'Apply Template Groups' interface. On the left, there is a 'Template Apply Order' list with the following items: Default Template Group, Demo, North America, NTP, Security1, and trial. Each item has an 'Add' button (green) and a 'Remove' button (orange). Below the list are 'Apply' and 'Cancel' buttons. On the right, there is a table with 30 rows. The table has columns for 'Hostname', 'Template Groups', and 'Changes'. The 'Template Groups' column is further divided into 'Present' and 'Changes' sub-columns. The table lists various hostnames and the templates applied to them.

Hostname	Template Groups	
	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 are applied determines which template is used. Templates applied later (lower on the apply order list) 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.

Cloud Services

This section includes the various cloud services Silver Peak offers.

AWS Transit Gateway Network Manager

Configuration > Cloud Services > AWS Network Manager

Before you begin the AWS Transit Gateway Network Manager configuration in Orchestrator, you need to create an AWS account to authenticate and authorize Orchestrator in the [AWS application](#). After you do this, complete the following prerequisites for the AWS Transit Gateway Network Manager, and then complete the Orchestrator configuration.

Prerequisites for AWS Transit Gateway Network Manager

Ensure you have completed the following tasks in AWS console before Orchestrator configuration:

- Navigate to the **Identity and Access Management (IAM)** under **Services** to create a user profile with permissions for Orchestrator.
- Navigate to the **Virtual Private Cloud (VPC) Dashboard** and configure your Transit Gateways for desired regions.
- Navigate to **Network Manager** from the **VPC Dashboard** under **Transit Gateways** to create a Global Network.
- Associate your Transit Gateways to your Global Network (Optional).

Create a User Profile in AWS

Complete the following steps to create a user profile in AWS.

1. Sign in to AWS and navigate to the **Identity and Access Management (IAM)** service from the main AWS Management Console (**Services > Security, Identity, & Compliance > IAM**).
2. Click **User** in the left menu under **Access Management**.
3. Click **Add User**.
4. Enter a user name in the **User name** field.
5. Choose the **Access Type**: Either **Programmatic Access** or **AWS Management Console Access**.

NOTE For seamless integration with Orchestrator, you will need to choose Programmatic Access to obtain the **Access Key ID** and the **Secret Access Key**.

6. Click **Next: Permissions**.
7. Set the Permissions for your user in this page. You can do this in one of three ways:

- **Adding a user to your group** – The user will inherit the permissions assigned to the group.
 - **Copying permissions from an existing user** – Copy permissions from an existing user in AWS and assign to the user you want.
 - **Attaching existing policies directly** – Attach a file containing the permissions and assign to the user.
8. Assign optional tags for your user. If you choose to add a tag, complete the steps:
- a. Enter a **key** – This represents the name of your tag.
 - b. Enter a **value** – Enter text that you want the key/tag to represent.

INFO Tags allow you to provide additional information about your user or group for tracking and organizational purposes. You can have a total of 50 tags.

9. Select **Next: Review**. This page displays the review of the profile you just created for your user. The **User Details, Permissions Summary**, and additional information such as tag, are shown.
10. Select **Create User**. The page should now show the following success message, along with **Access Key ID** and the **Secret Access Key** associated with your configured user.

Create Transit Gateways

Next, you need to create transit gateways to associate with your AWS Network Manager. Transit Gateways represent the tunnels in the various regions that allow connectivity to AWS. Complete the following steps below to create your gateways.

1. Navigate to the **Virtual Private Cloud (VPC) Dashboard (Services > Networking & Content Delivery)**.
2. Click **Transit Gateways**, under **Transit Gateways** in the left menu.
3. Click **Create Transit Gateways**.

[Transit Gateways](#) > Create Transit Gateway

Create Transit Gateway

A Transit Gateway (TGW) is a network transit hub that interconnects attachments (VPCs and VPNs) within the same account or across accounts.

Name tag ⓘ

Description ⓘ

Configure the Transit Gateway

Amazon side ASN ⓘ

DNS support ☒ enable ⓘ

VPN ECMP support ☒ enable ⓘ

Default route table association ☒ enable ⓘ

Default route table propagation ☒ enable ⓘ

Configure sharing options for cross account

Auto accept shared attachments ☐ enable ⓘ

* Required

- Complete the following fields to create your transit gateways.

Field	Description
Name Tag	Enter a name that represents your transit gateway.
Description	Enter a description to help identify your transit gateway. This is the description for the above Name Tag.
Amazon side ASN	Autonomous System Number that represents your transit gateways in AWS. You can use an existing ASN assigned to your global network or a private ASN. See the range limitations in AWS.
DNS Support	Select this check box if you want to enable Domain Name System support for your VPC within your Transit Gateways.
VPN ECMP support	Select this check box if you want to enable Equal Cost Multi-Path routing support in your Transit Gateways. This will allow for traffic with the same source and destination to be sent across the same multiple paths.
Default Route Table Association	Select this check box if you want to automatically associate other Transit Gateways to the route table that this one is using.

Field	Description
Default Route Propagation	Select this check box if you want to automatically create other Transit Gateways with this same route table.
Auto-accept shared attachments	Select this check box if you want your transit gateways to automatically accept attachments associated with different accounts.

5. Click **Create Transit Gateway**. A success message should display along with your Transit Gateway ID.

Create a Network Manager

After you have created your Transit Gateway and associated them to the required regions, you can create a Global Network in AWS. A Global Network hosts your specified transit gateways and is managed by the AWS Network Manager.

1. Navigate to the **VPC Dashboard**.
2. Click **Network Manager** under **Transit Gateways**.
3. Click **Create Global Network**.
4. Enter a **Name** and **Description** for your **Global Network**.
5. Click **Create**.

Orchestrator Configuration

When you have completed the AWS prerequisites, navigate to the **AWS Network Manager** tab in Orchestrator. There are six buttons above the table on this tab that are used to complete the AWS and Orchestrator integration: **Subscription**, **Interface Labels**, **Tunnel Settings**, **VTI Subnet Pool**, **Zone**, and **Network Manager Association**.

Subscription

1. To begin, click the **Subscription** button.
2. Enter the **Access Key ID** and the **Secret Access Key** that reflect your account in AWS.
3. Click **Save** after you have entered the information in the table below. The **AWS Reachability** field should

reflect **Connected**.

Field	Description
AWS Reachability	Connection status of the AWS Network Manager to Orchestrator: Connected or Not Connected .
Access Key ID	Access Key given to you in AWS to log in to the AWS console.
Secret Access Key	Secret Access Key given to you in AWS to log in to the AWS console.
Polling Interval	Indicates how often Orchestrator should check for configuration changes in the AWS transit gateways or Network Manager. The default polling interval is ten minutes.

4. Click **Save**.

You now should have an established connection with Orchestrator to your AWS VPC.

Interface Labels

The Build Tunnels Using These Interfaces dialog box enables you to select the interfaces to build your tunnels to AWS.

1. Click the **Interface Labels** button. The **Build Tunnels Using These Interfaces** dialog box opens.
2. Drag the interface labels you want to apply from the column on the right into the **Primary** columns.
3. Click **Save**.

Network Manager Association

In this window, you can choose which appliances you want to connect or disconnect to the AWS Gateway Network Manager through the transit gateways you previously configured in AWS. The appliances used with the transit gateway appear on the right under **Transit Gateway**.

1. Select or clear the check box next to the configured transit gateway you want to connect to the Network Manager.
2. See the table below for the fields.

Field	Description
Hostname	Host name of the appliance you want to connect or disconnect from the Network Manager.
Transit Gateways Present	Lists the gateways in association with the Network Manager and Orchestrator.
Transit Gateways Changes	Displays if the gateway has been added or removed from the Network Manager.

3. Click **Save**.

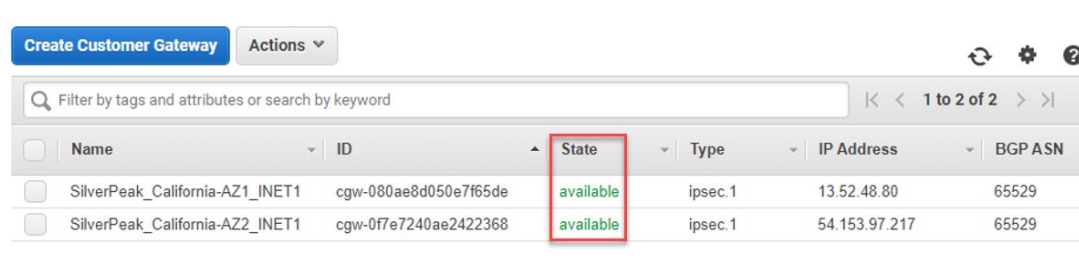
Tunnel Settings

The **Tunnel Settings** window defines the tunnels delivering traffic between AWS and Orchestrator. The tunnel settings are set using the default VPN configuration parameters received from virtual WAN APIs located in AWS.

Use the default settings for General, IKE, and IPSec tunnels and click **Save** to apply. Click the following link for information about more support tunnel options: <https://docs.aws.amazon.com/vpn/latest/s2svpn/VPNTunnels.html>

Navigate back to the **Virtual Private Network** section in AWS and select **Customer Gateways** and **Site-to-Site VPN Connections**. On these tabs, you can confirm that the IPSec tunnels you created in Orchestrator are functioning correctly.

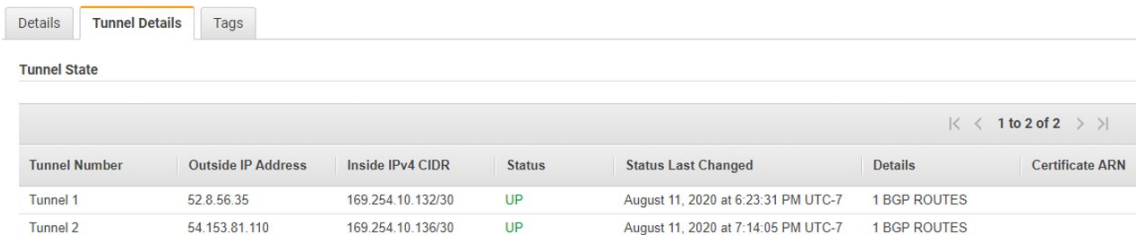
The tunnels should be in the 'available' state.



The screenshot shows the AWS Customer Gateways console. At the top, there is a 'Create Customer Gateway' button and an 'Actions' dropdown. Below is a search bar and a table with two columns: 'Name' and 'ID'. The table contains two rows of data, both with a 'State' column highlighted in green and labeled 'available'.

Name	ID	State	Type	IP Address	BGP ASN
SilverPeak_California-AZ1_INET1	cgw-080ae8d050e7f65de	available	ipsec.1	13.52.48.80	65529
SilverPeak_California-AZ2_INET1	cgw-0f7e7240ae2422368	available	ipsec.1	54.153.97.217	65529

The IPSec tunnels statuses should be 'UP'.



The screenshot shows the AWS Tunnel State console. At the top, there are tabs for 'Details', 'Tunnel Details', and 'Tags'. Below is a table with columns: 'Tunnel Number', 'Outside IP Address', 'Inside IPv4 CIDR', 'Status', 'Status Last Changed', 'Details', and 'Certificate ARN'. The table contains two rows of data, both with a 'Status' column highlighted in green and labeled 'UP'.

Tunnel Number	Outside IP Address	Inside IPv4 CIDR	Status	Status Last Changed	Details	Certificate ARN
Tunnel 1	52.8.56.35	169.254.10.132/30	UP	August 11, 2020 at 6:23:31 PM UTC-7	1 BGP ROUTES	
Tunnel 2	54.153.81.110	169.254.10.136/30	UP	August 11, 2020 at 7:14:05 PM UTC-7	1 BGP ROUTES	

VTI Subnet Pool

In this window, set the Subnet IP address and the mask for the AWS subnet pool. Enter the subnet IP address and the mask ID in the designated fields. Any updates to the subnet pool configuration results in service disruption.

You can have duplicated ASNs if you have a site with the same name.

NOTE This is an AWS specific subnet pool. Therefore, every subnet IP address must start with **169.254** to be included in this pool.

Verification

You can verify the stability and connectivity of your tunnels to the AWS Network Manager using the Connection Status column on the AWS Network Manager Tab. This column shows the BGP Peer status. Additional details can be found on the **Tunnels**, **VTI**, and **BGP** tabs.

After the tunnels and the BGP sessions are established, the TGW route table shows the routes learned from the EdgeConnect devices. To create a route table for your transit gateways, navigate to the **VPC Dashboard** in AWS and click **Transit Gateway Route Tables** under **Transit Gateways**. To create a static route, select the transit gateway from the Route Table and navigate to the **Routes** tab.

The screenshot shows the AWS Transit Gateway Route Tables console. At the top, there's a search bar and a table listing route tables. The selected route table is 'tgw-rtb-055011deaeffb4df'. Below this, the 'Routes' tab is active, showing a list of routes. The routes are filtered by attributes or search by keyword. The routes table has columns: CIDR, Attachment, Resource type, Route type, and Route state. Three routes are listed, all with a state of 'active'.

CIDR	Attachment	Resource type	Route type	Route state
10.192.0.0/16	tgw-attach-07781aa5f63142731 vpc-04ba33464fdaa4a82	VPC	propagated	active
10.30.12.0/24	2 Attachments	VPN	propagated	active
10.30.2.0/24	2 Attachments	VPN	propagated	active

Enter the values in the following table, and then click **Create Static Route**.

Field	Description
CIDR	Specified range of IPv4 addresses for your VPC.
Blackhole	Enable if you want your matched traffic to be dropped.
Choose attachment	Choose the attachment for your static route.

Check Point CloudGuard Connect

Check Point CloudGuard Connect provides network and cloud security with policies defined within Orchestrator overlays. The **Check Point CloudGuard Connect** tab has the following fields.

Field	Description
Subscription	Name of the appliance you want to connect with Check Point.
Interface Labels	Name of the interfaces you want to connect with Check Point.
Tunnel Settings	Defines the tunnels associated with Orchestrator and Check Point.
LAN Subnets	Subnets configured on the LAN side associated with Check Point.

Before you begin to configure Check Point CloudGuard Connect, you need to create a CheckPoint account. Visit the following link to make an account: <https://portal.checkpoint.com>.

After you create an account, you will need to create an API Key.

Subscription

1. After you complete the steps in the above URL to create your CheckPoint account, navigate to the **Check Point CloudGuard Connect** tab in Orchestrator.
2. Select the **Subscription** tab to get started with CheckPoint.
3. Enter your **Client ID** and the **Secret Key** you received when you created your Check Point account.
4. Select **Save** after you finish entering the information in the table below. The **Connection Status** should appear at the top of the **Subscription** window.

Interface Labels

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

Tunnel Settings

The **Tunnel Settings** tab helps you define the tunnels associated with CheckPoint and EdgeConnect. Use the CheckPoint default values for the **General**, **IKE**, and **IPsec** tunnel settings.

NOTE You also can configure specific General, IKE, and IPSec tunnel settings. The settings are automatically generated; however, you can make modifications if you choose to do so. To go back to the default settings, select **Use Default** on any of the tunnel windows.

LAN Subnets

You can select the LAN subnets for a given appliance to associate with your CheckPoint integration. By default, LAN subnets are configured on the **Deployment** tab. You also can add, import a CSV file, or export a CSV file of the configured subnets.

Enabling Check Point CloudGuard Connect

When you have completed configuration, you need to enable the Check Point service.

1. Navigate to the **Business Intent Overlay** tab in Orchestrator.
2. Go to the **Breakout Traffic to Internet & Cloud Services**.
3. Select the overlay that breaks out traffic to CheckPoint.

4. Drag **Check Point CloudGuard Connect** from the **Available Policies** column to the **Preferred Policy Order** column.

Verification

Navigate to the **Check Point CloudGuard Connect** tab in Orchestrator to verify successful deployment under **Site Status**. You also can verify successful deployment on the **Tunnels** tab.

Import and Export Subnets

Import enables you to import a Comma Separated Values (CSV) file into a pair of appliances used in Orchestrator. 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 **Choose File**.
2. Locate the file you want to import on your desktop.
3. Select **Open**.
4. Select **Import**. Orchestrator generates the CSV file. The following table represents the fields in the exported CSV file.

Appliance	Configured Subnets
<Appliance Hostname>	<Configured subnets IP addresses>

NOTE The titles and double quotes should be removed from your file before importing.

CAUTION This import overwrites previously configured imports.

Microsoft Azure Virtual WAN

Microsoft Azure optimizes routing, automates large scale connectivity from various branches to Azure workloads, and provides unified network and policy management within Orchestrator. Use Azure to deploy to a single WAN circuit or for branch to branch connectivity by configuring virtual WANs to associated hubs.

Before you begin Microsoft Azure Virtual WAN configuration in Orchestrator, you need to use the Azure Virtual WAN portal to authenticate and authorize Orchestrator in Azure. You need to create the service principal, which focuses on single-tenant application to run within only one organization. Click the following link to get started:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>.

Azure Prerequisites

1. Create an Application in Azure and get the following Subscription details from the Azure Active Directory:
 - Subscription ID
 - Tenant (Directory) ID
 - Application (Client) ID
 - Client Secret Key
2. Create a storage account in Azure and get the following:
 - Storage Account Name
 - Storage Access Key
3. Create a resource group
4. Create Azure Virtual WANs with Hubs from your resource groups

Orchestrator Prerequisites

Complete the following tasks in Orchestrator:

1. Configure a VTI IP Pool
 - Enter a valid IPv4 Subnet

NOTE This is a unique address across the network. VTI interfaces created for Azure integration will be selected from this pool.

INFO Azure VTI interface zone is set to WAN interface zone. Any change in deployment for the WAN interface zone is applied to Azure VTI as well.

WARNING Any change in the VTI pool after it is configured is networking affecting. This operation should be performed during a maintenance window as it can take several hours for some Cloud services to complete.

2. Configure BGP ASN Global Pool
 - Enter the start and end ranges for ASNs
 - Add any reserved ASNs to exclude from being applied to appliances

NOTE If not previously enabled, Orchestrator enables BGP.

Orchestrator Configuration

After the above configuration is complete, navigate to the **Microsoft Azure Virtual WAN** tab in Orchestrator. There are four icons at the top of the table that complete the Azure and Orchestrator integration: **Subscription**, **Interface Labels**, **Appliance to Virtual WAN Associate**, and **Tunnel Settings**.

To begin, select the **Subscription** icon.

Subscription

1. Enter the information in the Subscription fields that reflect your Azure portal account.
2. Select **Save** after you have finished entering the information in the table below. The Azure field should reflect **Connected**.

The following table represents the values in the **Subscription** window from the Azure portal.

Field	Description
Azure Reachability	This field displays the connection status of your account with Azure.
Subscription ID	ID of your subscription.
Tenant ID	Name of your Azure AD tenant.
Client ID	Client ID of your Azure application.
Client Secret Key	Secret key of your Azure application.
Storage Account Name	Name of your storage account.
Storage Account Key	Storage account key.
Storage URL	Storage account URL.*
Configuration Polling Interval	Amount of time set that Azure data is updated. This is defaulted every one minute.

*Storage URL

The Storage URL is present on the **Storage Accounts** tab in your Azure portal. Complete the following steps to obtain your storage account URL.

1. After your storage account is created in Azure, create a blob container.
2. Get the blob container URL.
3. Suffix the URL with a slash and add a file name in the **Storage URL** field.

NOTE Append the URL with a slash for the file name. Do not end the URL with a slash.

Interface Labels

Select the order in which you want your interface labels to be used.

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**.

Associate Appliance to Virtual WAN

Each appliance is associated with **one** virtual WAN. Use this tab to add or remove specific sites to your virtual WANs.

1. Select the **Associate Appliance to Virtual WAN** icon.
2. Select an appliance from the tree in the left menu.
3. Check the box to **Add** or **Remove** the appliance to your virtual WAN in Azure.
4. Click **Save**.

Tunnel Settings

The **Tunnel Settings** tab defines the tunnels associated with Azure and Orchestrator. The tunnel settings are set using the default VPN configuration parameters received from virtual WAN APIs located in your Azure portal account.

In your Azure Portal Account, navigate to the Azure Configuration table. This table displays the VPN site created for Orchestrator appliances associated to Azure virtual WANs. Additionally, manually associate sites to your hubs in Azure.

1. Navigate to **Azure Virtual WAN**.
2. Select **Azure VPN site**.
3. Select **New Hub Association**.

Verification



The **Tunnel** page displays that Azure and Orchestrator have an established connection with Azure by displaying a tunnel status of **up - active**.

For more information about Azure configuration, visit the following link: <https://docs.microsoft.com/en-us/azure/virtual-wan/virtual-wan-site-to-site-portal>.

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 on 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 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"> Microsoft Office 365 Optimize and Allow categories for the respective applications
CriticalApps	SharePoint Online, Office 365 Exchange		
Default	For everything	Any policy order except "Drop"	<ul style="list-style-type: none"> Matches Microsoft Office 365 Default categories Office365Common applications

NOTE Do not specify other individual Office applications in this group or overlay.

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

Zscaler Internet Access

Zscaler Internet Access is a cloud security service. EdgeConnect traffic can be service chained to Zscaler for additional security inspection.

Field	Description
Appliance	Name of the appliance you want to connect with Zscaler.
Interface Label	Name of the interfaces you want to connect with Zscaler.
VPN Credentials and Location Status	VPN credentials and location status of your subscription with Zscaler.
Gateway Options	Optional add-on that enables you to configure sub-locations and various rules for your sub-locations.
Zscaler ZENS	Zscaler Enforcement Nodes: The Zscaler endpoints where the tunnels connect. The discovered ZENS in this column are populated based on the appliance's geographical location.

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

NOTE This section represents **automated** configuration of IPSec, IKE, and GRE tunnels from EdgeConnect to the Zscaler cloud. Refer to the Zscaler-Silver Peak IPSec Integration Guide: Manual Mode and the Zscaler-Silver Peak GRE Integration Guide: Manual Mode if you want to manually configure the tunnels with the Zscaler cloud.

Subscription

1. Go to <https://help.zscaler.com/zia/sd-wan-api-integration>.
2. After 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** after you have finished entering the information in the table below. The Zscaler field should reflect **Connected**.

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

Field	Description
Zscaler	This field displays if you are connected or not connected to your Zscaler account.
Zscaler Cloud	Zscaler cloud URL. Ex: admin.zscalerthree.net.
Silver Peak Partner Username	Partner administrator user name you created when configuring Zscaler.
Silver Peak Partner Password	Partner administrator password you created when configuring Zscaler.
Silver Peak Partner Key	Partner key you created when configuring your Zscaler account. Select Silver Peak from the list of partners.
Domain	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. Use the Zscaler defaults for Tunnel Settings defined by the system.

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.

Interface Labels

Select the **Primary label** you want your traffic to go to. **Backup labels** will be used as the second option if the primary is unreachable.

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**.

WARNING This is service affecting. Any changes to the interface selection can cause previously built tunnels to be deleted and rebuilt.

ZEN Override

You can use the **ZEN Override** 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 **ZEN Override** tab.
2. Enter the appliance name, the interface label, and the Primary and Secondary IP addresses. Orchestrator will build tunnels to those ZENS.

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

Gateway Options

Use this tab to configure gateway options and rules for Zscaler sublocations. Orchestrator uses location and sub-locations to better define a branch site in the Zscaler cloud. Sub-locations are LAN-side segments within each branch and can be identified by LAN interfaces, zones, or a collection of LAN subnets. Click **Gateway Options** to begin configuration, if you choose to enable this add-on.

Zscaler Gateway Options

This window enables you to configure sub-locations and their rules.

1. Click **Add**.
2. Enter a Rule Name in the **Rule Name** field.

WARNING If two rules have the same sub-location name or IP address, Orchestrator picks the first match and considers the order of the rules.

3. Enter a location by entering an appliance name, region, or group in the **Appliances** field.
4. Enter the WAN label in the **Location Label** field.

If you check Sub-location:

1. Enter the sub-location name in the **Name** field.
2. Enter the subnet address (LAN label, Firewall Zone, or subnet) in the **Subnets** field.
3. Click **Save**.

NOTE Sub-locations can be applied to all WAN links chosen under **Zscaler Internet Access > Interface setting**.

Check **Show Sub-Locations** and the sub-locations configured in the Gateway options appear in the Zscaler table.

IP SLA

Click the IP SLA icon to configure IP SLA for Zscaler Tunnels. This configuration ensures tunnel connectivity and internet availability between Zscaler and Orchestrator. The **Zscaler IP SLA Configuration** window opens. If the tunnel cannot reach Zscaler, the tunnel is considered as DOWN.

Enable 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. After 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

ZEN Override

Gateway Options

☐ Show Sub-Locations

Pause Orchestration

2 Rows

Search

Appliance	Interface Label	Gateway Options	Bandwidth (Mbps)	VPN Credentials and Location Status	Zscaler ZENs
Feynman-Powers	INETA	Use XFF from Client Request=false, E...	Upload=OFF, Download=OFF	Deployed	Discovered: 104.129.206.161, 165.22...
Feynman-Powers	INETB	Use XFF from Client Request=false, E...	Upload=OFF, Download=OFF	Deployed	Discovered: 104.129.206.161, 165.22...

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

Tunnels

OverlayUnderlayPassthrough

Status: All

138/800 Rows

Search

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	100000(Auto)	
	EMEA1-Amsterdam-A...	ThirdParty_Zscaler_INETA...	up	up - active	172.29.100.4	165.225.88.39	IPSec	none	Zscaler_INETA_Backup	100000(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_INETB...	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)	
	AP11-Singapore-Azure	ThirdParty_Zscaler_INETA...	up	up - active	10.5.0.4	165.225.116.24	IPSec	none	Zscaler_INETA_Backup	200000(Auto)	
	AP11-Singapore-Azure	ThirdParty_Zscaler_INETB...	up	up - active	10.5.1.4	165.225.112.24	IPSec	none	Zscaler_INETB_Primary	200000(Auto)	
	AP11-Singapore-Azure	ThirdParty_Zscaler_INETB...	up	up - active	10.5.1.4	165.225.116.24	IPSec	none	Zscaler_INETB_Backup	200000(Auto)	
	AP11-Singapore-Azure	ThirdParty_Zscaler_INETA...	up	up - active	10.5.0.4	165.225.112.24	IPSec	none	Zscaler_INETA_Primary	200000(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)	
	AP12-Singapore-Azure	ThirdParty_Zscaler_INETB...	up	up - active	10.5.1.5	165.225.112.24	IPSec	none	Zscaler_INETB_Primary	200000(Auto)	
	AP12-Singapore-Azure	ThirdParty_Zscaler_INETA...	up	up - active	10.5.0.5	165.225.116.24	IPSec	none	Zscaler_INETA_Backup	200000(Auto)	

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.

Service Orchestration

Configuration > Cloud Services > Service Orchestration

Use the Service Orchestration tab to automate the integration of third-party services without an API. Service Orchestration automates the creation and deployment of IPSec tunnels and IP SLA probes and manages the lifecycle of the tunnels and probes.

Service Orchestration creates a local tunnel identifier (IKE ID) for each tunnel to the third-party service. After the tunnels are created, complete the integration on the third-party service's site by replacing the source identity values with the local tunnel identifiers (IKE IDs) that Orchestrator created for each endpoint.

NOTE By default, Service Orchestration provides the framework for Netskope integration. The instructions on this page are specific to Netskope, but you can apply the same general procedure to other third-party services.

Prerequisites

- You must have loopback interfaces configured to use the Service Orchestration feature.
- Service Orchestration supports third-party services that use IPSec IKEv2 endpoints.
- You will need the following information from the third-party service for each endpoint you want to add:
 - Endpoint name
 - IP address
 - Probe address

Remote Endpoint Configuration

Add the remote endpoints for Netskope.

You can add one endpoint at a time or add endpoints in bulk by importing the information from a CSV file.

Add Endpoints One at a Time

1. Click **Remote Endpoint Configuration**.

The Add Remote Endpoints for Netskope dialog box opens.

2. Click **+Remote Endpoint**.
3. Complete the following fields—press the **Tab** key to navigate to the next field.

Field	Description
Name	Name of the Netskope endpoint. IMPORTANT: If an endpoint name is decommissioned or modified, you must update the value in this table.
IP Address	IP address of the Netskope endpoint. IMPORTANT: If an IP address is decommissioned or modified, you must update the value in this table.
Interface Label	The Silver Peak interface labels that can be provisioned for this endpoint. Only labels in this list will be provisioned. HINT: Click Interface Label Default to reset the Interface Label for every endpoint in the table to the default value of Any .

Field	Description
Pre Shared Key	<p>The pre-shared key for the endpoint. To display the pre-shared key, click anywhere in the field. Do one of the following:</p> <ul style="list-style-type: none"> Edit this field for each endpoint. This value can be an ASCII string, a hex encoded string (if it has a 0x prefix), or a base64-encoded string (if it has a 0s prefix). Click PSK Default to create and save a pre-shared key. Every endpoint will use the pre-shared key you create. Because traffic going to these endpoints is encrypted, it will not compromise security to use the same pre-shared key for each endpoint.
Probe Address	<p>The Netskope endpoint that the IP SLA subsystem will ping. You can obtain the probe address from the third-party security provider.</p> <p>IMPORTANT: Orchestrator will prefill the Address field in the IPSLA Settings dialog box with this value. If you delete the value in the Probe Address field in this table, Service Orchestration will ping the value specified in the Address field in the IPSLA Settings for Netskope dialog box.</p>
Backup Remote Endpoint	<p>Enter the Netskope endpoint that you want to use as a backup tunnel. For example, ATL1-Atlanta could use DFW1-Dallas as a backup remote endpoint. If you leave this field empty, the endpoint will not have a backup tunnel. The BIO determines how traffic will be handled if a single or single and backup tunnel go down.</p>

- Repeat these steps for each endpoint.

TIP To delete an endpoint, click the **X** in the last column in the table.

- Click **Save**.

Updates are orchestrated immediately.

Add Endpoints in Bulk

- Click **Remote Endpoint Configuration**.

The Add Remote Endpoints for Netskope dialog box opens.

- Click **Import** to import a list of remote endpoints from a CSV file. The CSV file must contain columns for name, IP address, interface label, pre-shared key, probe address, and backup remote endpoint, in that order.

NOTE Remove any header rows before you import the file.

- Click **Choose File**.
- Navigate to the file, select the file, and then click **Open**.
- Click **Save**.

Updates are orchestrated immediately.

Bulk Edits

To make bulk edits to the table:

1. Click **Export**.
2. Open the CSV file and delete the three header rows.
3. Modify, save, and close the file.
4. Click **Import**, and then click **Choose File**.
5. Locate and select the file, and then click **Open**.

Orchestrator updates the table.

6. Click **Save**.

Interface Labels

Select the Primary and Backup interface labels for your traffic. Backup interface labels will be used if the primary interface labels are unreachable.

NOTE Netskope does not support Active – Active backup.

1. Click **Interface Labels**.

The Build Tunnels using these Interfaces for Netskope dialog box opens.

2. Drag the interface labels you want to use into the Primary area. (The Peer/Service names in the Tunnels table will be NSK_Primary_1 and NSK_Primary_2.)
3. Drag the interface labels you want to use into the Backup area. (The Peer/Service names in the Tunnels table will be NSK_Backup_1 and NSK_Backup_2.)
4. Drag the interface labels up or down to reorder the list as necessary.
5. Click **Save**.

Tunnel Settings

Click **Tunnel Settings** to configure the Netskope tunnel settings.

IP SLA Settings

1. Click **IP SLA Settings**.

The IPSLA Setting for Netskope dialog box opens.

2. If all fields are dimmed, click **Enable IP SLA rule orchestration**.

3. Complete the following fields.

Field	Description
Monitor	Ping or HTTP/HTTPS.
Address	Netskope endpoint that the IP SLA subsystem will ping. Orchestrator prefills the Address field with the value from the Remote Endpoint Configuration table. You can configure up to three addresses.
Source Interface	Select an orchestrated loopback label.

4. Accept the default values for the remaining field, and then click **Save**.

Orchestrator builds the tunnels.

Pause Orchestration (Optional)

When troubleshooting, you can click **Pause Orchestration** and then click **Save** to pause the service orchestration. To restart the service orchestration, click **Resume Orchestration**.

+BIO Breakout

By default, the tunnels associated with a third-party service will be available for BIOs. You can upload an icon to display on the Business Intent Overlays tab.

NOTE Supported file types include PNG, JPEG, SVG, and WEBP. The recommended dimensions are 60 x 20 pixels.

1. Click **+BIO Breakout**.

The Configure BIO Breakout for Netskope dialog box opens.

2. Click **Upload Service Icon**.
3. Locate and select the file, then click **Open**.
4. Click **Save**.

This icon will display next to the service name on the Business Intent Overlays tab.

If you do not want this third-party provider to be available for BIOs, do the following:

1. Click **+BIO Breakout**.

The Configure BIO Breakout for Netskope dialog box opens.

2. Clear the **BIO Breakout** check box.
3. Click **Save**.

Remote Endpoint Association

The final step to configure the integration in Orchestrator is to associate EdgeConnect appliances with remote endpoints. Use this page to add or remove endpoints from an appliance. Silver Peak recommends that you associate one remote endpoint per EdgeConnect appliance.

1. In the Orchestrator appliance tree, select one or more appliances to associate with Netskope remote endpoints.
2. Click **Remote Endpoint Association**.

The Associate an Appliance to Netskope Remote Endpoints dialog box opens.

3. Select the **Add** or **Remove** check box next to the endpoints you want to associate with the selected appliances. Be sure to add the endpoints that are geographically closest to the appliances.
4. Verify the proposed changes to remote endpoints in the table to the right, and then click **Save**.

Add Tunnel Local Identifiers to Netskope

After the Service Orchestration integration is complete on the Silver Peak side, you must add the local tunnel identifiers (IKE IDs) to Netskope. You can simplify this process by exporting the Netskope configuration to a CSV file. The exported file contains all of the configuration details in the table on the Netskope page for all selected appliances, including IKE IDs.

NOTE The tunnel local identifier value is a fixed format: *hostname_labelname@IPaddress*. For example, EAST3-AWS_INETA@192.x.x.xxx.

1. In the Orchestrator appliance tree, select all appliances associated with Netskope remote endpoints.
2. On the **Netskope** page on the Service Orchestration tab, click **Export** to save the contents of the table to a CSV file.
3. Log in to Netskope.
4. In the IPSec configuration panel, replace the Source Identity values with the corresponding Tunnel Local Identifiers (IKE IDs) created by Orchestrator.

Verification

After Netskope is configured and the Netskope policy is applied successfully in the BIO, deployment will begin automatically. Go to the **Netskope** tab and view the Connection Status column to verify that the deployment was successful.

Set Up a New Service

To set up a new third-party service:

1. Click **+Add Service** and complete the following fields.

Field	Description
Name	Name of the new service.
Prefix	A prefix to assign to all tunnels for this service. Orchestrator will use this prefix to filter tunnels and IP SLAs.

2. Click **Save**.

A new tab is created on the Service Orchestration page.

TIP To edit or delete a service, click the edit icon next to the service name.

3. Select the tab for the new service and follow the steps explained in Set Up Netskope Integration to integrate this new service.

Deploy EC-V in Cloud

Starting with Orchestrator 9.0.5, you can deploy one or more EdgeConnect Virtual (EC-V) appliances in supported platforms. At this time, Silver Peak supports AWS.

Before you begin, complete the following tasks:

1. On the AWS dashboard, create an AWS Identity and Access Management (IAM) user account with required permissions for Orchestrator to create AWS resources. Silver Peak recommends a dedicated AWS IAM user account for Orchestrator.
 - a. Create an AWS Policy that contains all permissions the Orchestrator requires to create an EC-V.
 - b. Attach the Policy to the Orchestrator's IAM user account.
 - c. Download the Security credentials of the Orchestrator's IAM user account.
2. On the EC2 dashboard, create a key pair to assign to the EC-V. You will need this key pair if you want to SSH into the EC-V after the deployment.

After creating the AWS IAM account, click **New EC-V Deployment** on the [Deploy EC-V in Cloud Tab](#) to configure and deploy one or more EC-V cloud instances.

For more information about deploying EC-V appliances in the public cloud, see:

- [Deploy EC-V in Cloud Tab](#)
- [Cloud Deployment Accounts](#)
- [AWS Account Configuration](#)
- [EC-V Deployment Configuration](#)

After deploying an EC-V in the cloud, navigate to Orchestrator's Discovered Appliances page to view the deployment status. If the EC-V is still being deployed, the status in the Approve column will indicate Configuring. It takes approximately ten minutes to deploy and configure a cloud EC-V. Click **Refresh Discovery Information** to determine whether the appliance is ready to be approved into the SD-WAN fabric.

When configuration is complete and the green Approve button appears, the EC-V is fully configured in Inline Router mode with mgmt0, wan0, and lan0 MAC addresses assigned. While adding the EC-V, the Deployment Profile page will show LAN IP address, WAN IP address, WAN interface firewall mode, and WAN bandwidth value assigned by Orchestrator.

You can upgrade the appliance software version on a cloud EC-V after approving and adding it to the SD-WAN fabric.

After a cloud EC-V has been deployed, you can add another EC-V into the same deployment. The new EC-V will use the same settings from the existing deployment configuration such as AWS account, region, VPC, key pair, and instance type. You can deploy the new instance into an Availability Zone that is already used by an existing appliance or a new Availability Zone.

Deploy EC-V in Cloud Tab

Configuration > Cloud Services > Deploy EC-V In Cloud

The Deploy EC-V In Cloud tab provides the AWS account details and EC-V deployment configuration details for all cloud EC-Vs that have been deployed.

Deploy EC-V In Cloud 14 mins

Accounts New EC-V Deployment

8 Rows Search

Name	Provider	Account	Instances	Status	Destroy	Deployment Info	Resources	Comment
WestHubs	aws (vpc:10.39.0.0/24)	SEWAN AWS	2 <a>+Add	Deployed	Destroy	<a>i	<a>i	
EastHubs	aws (vpc:10.37.1.0/24)	SEWAN AWS	1 <a>+Add	Deployed	Destroy	<a>i	<a>i	
MiddleEastHub1	aws (vpc:10.37.4.0/24)	SEWAN AWS	1 <a>+Add	Deployed	Destroy	<a>i	<a>i	
AustraliaHub2	aws (vpc:10.37.6.0/24)	SEWAN AWS	1 <a>+Add	Deployed	Destroy	<a>i	<a>i	
IndiaHubs	aws (vpc:10.37.8.0/24)	SEWAN AWS	1 <a>+Add	Deployed	Destroy	<a>i	<a>i	
SouthAmericaHubs	aws (vpc:10.37.10.0/...	SEWAN AWS	1 <a>+Add	Deployed	Destroy	<a>i	<a>i	
AWS-Lon-Lowe	aws (vpc:10.207.0.0/...	Laurence Lowe	1 <a>+Add	Incomplete <a>i	Destroy	<a>i	<a>i	
Denver-Hub	aws (vpc:10.37.9.0/24)	SEWAN AWS	1 <a>+Add	Deployed	Destroy	<a>i	<a>i	New hub for Denver area

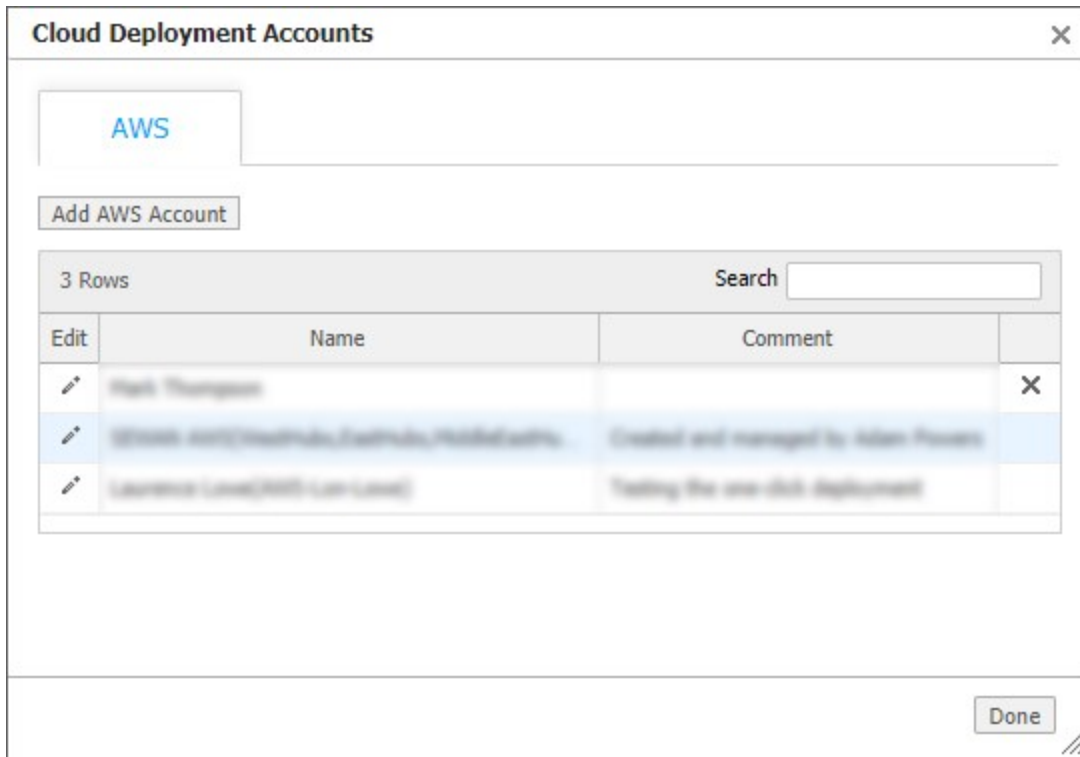
The following table describes each field on this page.

Field	Description
Name	Name given on the deployment configuration page.
Provider	Name of the cloud provider and VPC details.
Account	Name of the AWS account that was used to deploy the EC-Vs.

Field	Description
Instances	Number of EC-V instances in the deployment. To add one or more EC-Vs to the deployment, click +Add . In the New Instance on AWS dialog box, select the availability zone to use and any optional tags to apply to the new instance.
Status	Status of the deployment. If more information is available, an information icon is displayed.
Destroy	To permanently delete a deployment, click Destroy . This action deletes all resources associated with the EC-Vs, including all EC2 resources.
Deployment Info	Click the info icon in this column to view deployment and instance details.
Resources	Click the info icon in this column to view details about each AWS resource that Orchestrator created during the deployment.
Comment	Comments that were added to the deployment when the EC-V was created.

Cloud Deployment Accounts

The Cloud Deployment Accounts dialog box lists all of the AWS accounts that have been added.



- Click **Add AWS Account** to create a new account for EC-V deployments.
- Click the edit icon next to an existing account to modify that account's details.

AWS Account Configuration

Complete the following steps to create an AWS IAM user account with the required permissions for creating EC-V instances in AWS.

Create a Policy with Required Permissions

1. Log in to the AWS Dashboard.
2. On the Find Services search menu, enter **IAM** to open the Identity and Access Management (IAM) page.
3. Under Access Management, click **Policies**. The Policies page opens.
4. Click **Create policy** and click the **JSON** tab.
5. Delete the existing text.
6. Go to this [web page](#), click the link for your version of Orchestrator, and then copy and paste the JSON policy text into the editor.
7. Click **Next: Tags**.
8. (Optional) Add metadata to the policy by attaching tags as key-value pairs.
9. On the Review policy page, enter a name and optional description for the new policy.
10. Review the policy summary to see the permissions granted by your policy, and then click **Create policy** to save your work.

Attach Policy to the Orchestrator IAM User Account

1. Click **Users > Add user**. The Add user page opens.
2. Enter a user name in the User name field (for example, ArubaOrchestrator).
3. Under Access type, select **Programmatic access**, and clear the AWS Management Console access check box.
4. Click **Next: Permissions**.
5. Under Set Permissions, click **Attach existing policies**.
6. Select the Policy document you created from the list, and then click **Next: Review**.
7. Under Permissions summary, click **Add permissions**.

Download Orchestrator IAM User Account Credentials

1. On the Users page, click the **Security credentials** tab.
2. Download or copy and paste the Access key ID and Secret key ID to a secure place for later use.

Create a Key Pair to Assign to EC-Vs

Review the instructions on [this page](#) to create a key pair on the AWS region where you plan to deploy the EC-V.

Add the AWS Account to Orchestrator

Complete the following fields for Orchestrator, and then click **Save** when finished.

Field	Description
Name	Enter a unique name. If you have multiple AWS accounts, you must enter a unique name for each account.
Access Key	Enter the Orchestrator IAM user's Access Key ID that you saved earlier.
Secret Key	Enter the Orchestrator IAM user's Secret Key ID that you saved earlier.
Comment	Enter a comment that provides any additional information about the AWS account.

Orchestrator validates the account information. This takes approximately 45 seconds.

EC-V Deployment Configuration

Use the Deploy EdgeConnect in AWS page to create one or more EC-V instances in an AWS region.

Deploy EdgeConnect in AWS ?

Accounts

Name*

AWS Account*

Region*

VPC CIDR* /

SSH Key*

Boost ☐

WAN Bandwidth* Mbps

Instance Type*

AWS Tags

Comment

Horizontally Scale - 1 +

By default, Orchestrator deploys ec-v version: 8.3.0.16

Instance - 1

Availability Zone*

Appliance Tag

[Advanced Settings](#)

Field	Description
Name	Enter a name for the deployment. This name is used only for identifying the deployment. A deployment consists of one or more EC-Vs that an Orchestrator creates in an AWS Virtual Private Cloud (VPC). Only alphanumerical letters and hyphens are allowed in the deployment name. The maximum allowed length is 20 characters.
AWS Account	Select an AWS account to use for deploying the EC-V.

Field	Description
Region	Select an AWS region where you want to deploy the EC-V.
VPC CIDR	Enter a VPC Classless Inter-Domain Routing (CIDR) block. The smallest supported CIDR block is /24 and the largest supported CIDR block is /16. Orchestrator creates all AWS resources required for the EC-V deployment within this VPC. For each EC-V you deploy, Orchestrator creates three subnets that are /28 in size. In other words, if you deploy two EC-Vs, Orchestrator creates six subnets in total. This is true even if both EC-Vs are created in a single Availability Zone.
SSH Key	Select an existing AWS key pair to assign to the EC-V. A key pair must be created prior to the deployment.
Boost (Optional)	<p>Boost requires additional resources on an AWS EC2 instance. After Boost and an appropriate WAN Bandwidth value are selected, Orchestrator displays the appropriate AWS instance types for the deployment on the Instance Type drop-down menu.</p> <hr/> <p>NOTE Selecting the Boost check box does not enable Boost on the EC-V. It only allows Orchestrator to display appropriate AWS instance types that can support Boost for the selected WAN bandwidth. To enable Boost on the EC-V, go to the Deployment page and the Business Intent Overlay (BIO) page after the deployment is complete.</p> <hr/>
WAN Bandwidth	The Bandwidth drop-down list displays the current EdgeConnect license tiers. After you select a WAN Bandwidth value, Orchestrator displays the appropriate AWS instance types for the deployment on the Instance Type drop-down menu.
Instance Type	Based on your selection of Boost and WAN Bandwidth values, Orchestrator displays the appropriate AWS instance types on this drop-down menu.
AWS Tags (Optional)	Any comma-separated tags entered here are applied to all AWS resources that Orchestrator creates while deploying the EC-V. If you do not enter any tags, Orchestrator automatically creates a unique tag for each AWS resource that it creates while deploying the EC-V. This AWS tag is created to identify each resource created by Orchestrator. The tag is formatted as follows: <i>sp-automated-deployment name-instance-index-resource name</i> .
Comment (Optional)	Enter an optional comment if you want to attach any additional details for the deployment.
Advanced Settings	Custom AMI ID: Leave this field blank unless you have an EdgeConnect AMI that you want to use for the deployment. When this field is blank, the base AMI for the deployment will be obtained from the AWS Marketplace.
Horizontally Scale	You can deploy multiple EC-Vs by clicking + and selecting the Availability Zone for each EC-V. If the selected region supports multiple Availability Zones, each Availability Zone is shown on the drop-down menu. When deploying multiple EC-Vs, it is best practice to deploy each EC-V in a unique Availability Zone.
Appliance Tag (Optional)	Enter an Appliance Tag on this field if you want to assign a pre-configuration file to the deployment. If this field is left blank, Orchestrator will automatically assign an Appliance Tag for its own configuration purposes.

When you have completed all of the required fields, click **Review and Deploy**. Review the configuration summary, and click **Deploy** to create the EC-V instances.

Appliance Administration

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.

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 EdgeConnect appliance's **built-in user database** supports user names, groups, and passwords.

- Each appliance has two default user accounts, **admin** and **monitor**, that 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 by using the 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, 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	Description
Appliance Name	Name of the appliance selected.
Authentication Order	When it is possible to validate against more than one database (local, RADIUS server, TACACS+ server), Authentication Order specifies which method to try in what sequence: Authentication Order First , Order Second , and Order Third .
Authorization Map Order	Map ordering determines which server is used first. Select the map ordering from the drop-down list: Local-Only , Remote-First , and Remote-Only . The default (and recommended) value is remote-first .
Authorization Default Role	Default role assigned for authorization. The default (and recommended) value is admin .
Authentication	Process of validating that the end user, or a device, is who they claim to be.
Authorization	Action of determining what a user is allowed to do. Generally, authentication precedes authorization.
Map Order	Default (and recommended) value is Remote First .

RADIUS and TACACS+

RADIUS and TACACS+ Server Fields

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

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
✎	Tallinn	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 are not all synchronized, charts (and stats) inevitably have different timestamps for the same data, depending on which device you use to view the reports.

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

Click the edit icon to begin specifying the date and time for your appliances.

DNS (Domain Name Servers) Tab

Administration > General Settings > Setup > DNS

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

Edit	Appliance Name	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.

Field	Description
Appliance Name	Name of the appliance.
Primary DNS IP addr	IP address of the DNS the system uses first.
Secondary DNS IP addr	IP address of the DNS the system uses second.
Tertiary DNS IP addr	IP address of the DNS the system uses last.

Click the edit icon to add the three domain name servers.

SNMP Tab

Administration > General Settings > Setup > SNMP

This tab summarizes the **SNMP** configuration for each of the selected appliances.

SNMP								
Manage SNMP with Templates								
SNMP ?								
3 Rows								
Search								
	Appliance Name	Enable SNMP Agent	Enable SNMP Traps	Enable SNMP V1/V2	Enabled V3 Users	Trap Receivers		
Edit						Trap Receiver 1	Trap Receiver 2	Trap Receiver 3
	Tallinn	Yes	Yes	Yes				
	laine-vxa	Yes	Yes	Yes				
	laine-vxb	Yes	Yes	Yes				

SNMP Overview

EdgeConnect appliances support Management Information Base (MIB-II) as described in RFC 1213 for cold start traps, warm start traps, and Silver Peak proprietary MIBs. Appliances issue an SNMP trap during reset when loading a new image, recovering from a crash, or rebooting.

An appliance sends a trap every time an alarm is raised or cleared. Traps contain additional information about alarms, including severity, sequence number, a text-based description of the alarm, and the time the alarm was created. For more information, you can download a .zip archive containing supported MIBs at <https://www.silver-peak.com/download/latest/mibs.html>.

Modify SNMP Configuration

Click the edit icon to the left of an appliance row to modify the SNMP configuration.

Use this page to configure the appliance's **SNMP** agent and trap receivers.

1. Select the **Enable SNMP** check box to activate configuration options for SNMP v1/v2, SNMP v3, and **Trap Receivers** details.
2. If you select the **Enable SNMP Traps** check box, the SNMP agent on the appliance sends traps to configured receivers.

3. Use the **Default Trap Community** field to specify the string the trap receiver uses to accept traps being sent to it. The default value is **public**. You can modify this value.

SNMP v1/v2

Configure the following fields for SNMP v1 and v2c.

Field	Description
Enable SNMP	Allows the SNMP agent on the appliance to send traps to configured receivers.
Read-Only Community	The SNMP application needs to present this text string (secret) to poll the appliance's SNMP agent. The default value is public . You can modify this value.

SNMP v3

For additional security, configure SNMP v3 if you want to authenticate without using clear text. To add an SNMP v3 user, click **Add** above the SNMP v3 table and configure the following properties:

Field	Description
Enabled	Select this check box to enable the selected user. Clear this check box to disable the user and maintain the configuration.
Username	Enter the username to identify the SNMP v3 user.
Authentication Type	Select the authentication type to use for SNMP requests from the user. NOTE Authentication type is required and SHA-1 is the only supported algorithm.
Authentication Password	Enter a password that the SNMP agent can use to authenticate requests sent by the user. NOTE The password must be at least 20 characters long.
Privacy Type	Select the encryption type to use for encrypting requests from the SNMP user. NOTE Encryption is required, and AES-128 is the only supported algorithm.
Privacy Password	Enter a password (key) to use for encrypting requests sent by the user. NOTE The password must be at least 20 characters long.

NOTE To delete an SNMP v3 user, click the **X** to the right of the entry in the table.

Trap Receivers

To configure a trap receiver, click **Add** above the Trap Receivers table and configure the following properties:

NOTE You can configure up to three trap receivers per appliance.

Field	Description
Host	IP address of the host where traps should be sent.
Version	Select the SNMP version of the trap receiver.
Community/Username	For v1 and v2c, enter the community string the receiver should use to accept traps. If left blank, the default community string (public) is used. If a different community string is configured on the trap receiver, enter it here. For v3, specify the SNMP v3 user that is sending traps to the receiver.
Enabled	Select this check box to enable the receiver. Clear this check box to disable the receiver and maintain the configuration.

NOTE To delete a receiver, click the **X** to the right of the entry in the table.

Flow Export Tab

Administration > General Settings > Setup > Flow Export

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

Flow Exporting Enabled allows the appliance to export the data to collectors. The appliance exports flows against two virtual interfaces—**sp_lan** and **sp_wan**—that accumulate the total of LAN-side and WAN-side traffic, regardless of physical interface.

Select the **Edit** icon to open the **Flow Export Configuration** dialog box.

Silver Peak Custom Information Elements

See the table below for the Silver Peak Custom Information Elements.

Custom IE Name and Implementation Description	Semantics	Units	Field Length (bytes)	Enterprise ID
Data Type: ipv4Address				

clientIPv4Address:	default		4	1
<ul style="list-style-type: none"> TCP: source ipv4 address of SYN initiator is the client. UDP: source ipv4 address of the first packet is the client. 				
serverIPv4Address	default		4	2
<ul style="list-style-type: none"> TCP: destination ipv4 address of SYN initiator is the client. UDP: destination ipv4 address of the first packet is the client. 				
connectionInitiator	default		4	7
<ul style="list-style-type: none"> TCP: source ipv4 address of SYN initiator is the connection initiator. UDP: source ipv4 address of the first packet is the connection initiator. 				
Data Type: unsigned8				
connectionNumberOfConnections	totalCounter		1	9
<ul style="list-style-type: none"> Number of TCP connections (3-way handshake) or UDP sessions established. 				
connectionServerResponsesCount	totalCounter		1	10
<ul style="list-style-type: none"> Currently 1 				
connectionTransactionCompleteCount	totalCounter		1	21
<ul style="list-style-type: none"> Currently 1 				
Data Type: unsigned32				

connectionServerResponseDelay <ul style="list-style-type: none"> TCP: Round-trip time between SYN and SYN-ACK. UDP: Round-trip time between first onward and return packet. 		MS	4	11
connectionNetworkToServerDelay <ul style="list-style-type: none"> TCP, Round-trip time between SYN and SYN-ACK. UDP, Round-trip time between first onward and return packet. It is also called Server Network Delay (SND). 		MS	4	12
connectionNetworkToClientDelay <ul style="list-style-type: none"> TCP: Round trip between SYN-ACK and ACK. UDP: Round-trip time between first response and second request packet. It is also called Client Network Delay (CND). 		MS	4	13
connectionClientPacketRetransmissionCount <ul style="list-style-type: none"> Currently 1 	totalCounter		4	14

connectionClientToServerNetworkDelay <ul style="list-style-type: none"> Network Time/Network Delay is known as the round-trip time that is the summation of CND and SND. It is also called Network Delay (ND). 		MS	4	15
connectionApplicationDelay <ul style="list-style-type: none"> TCP: Round-trip time between SYN and SYN-ACK. UDP: Round-trip time between first onward and return packet. 		MS	4	16
connectionClientToServerResponseDelay <ul style="list-style-type: none"> The round-trip time that is the summation of CND and SND. 		MS	4	17
connectionTransactionDuration <ul style="list-style-type: none"> The flow displays the time difference between the first and last packet. 		MS	4	18
connectionTransactionDurationMin <ul style="list-style-type: none"> The flow displays the time difference between the first and last packet. 		MS	4	19
connectionTransactionDurationMax <ul style="list-style-type: none"> The flow displays the time difference between the first and last packet. 		MS	4	20
Data Type: unsigned64				
connectionServerOctetDeltaCount <ul style="list-style-type: none"> Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter. 	deltaCounter	octets	8	3
connectionServerPacketDeltaCount <ul style="list-style-type: none"> Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter. 	deltaCounter	packets	8	4
connectionClientOctetDeltaCount <ul style="list-style-type: none"> Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter. 	deltaCounter	octets	8	5

connectionClientPacketDeltaCount	deltaCounter	packets	8	6
<ul style="list-style-type: none"> Server initiated byte count. If flow is lan to wan, Lan-Tx byte counter. If flow is wan to lan Lan-Rx byte counter. 				

Data Type: String				
applicationHttpHost	default		variable length	8
<ul style="list-style-type: none">http destination domain name				
applicationCategory	default		variable length	27
<ul style="list-style-type: none">application group				
from-zone	default		variable length	22
<ul style="list-style-type: none">(source zone) name for the flow when ZBF is configured				
to-zone			variable length	23
<ul style="list-style-type: none">(destination zone) name for the flow when ZBF is configured				
tag	default		variable length	24
<ul style="list-style-type: none">the user-specified readable string/tag that can be specified when the ZBF rule is configured. If "tag" is not specified, an automatic tag will be created and exported. The automatic/default tag is constructed by concatenating <from-zone>_<to-zone>_<rule priority>. For example, "lan-zone_corp-zone_10000".				
overlay	default		variable length	25
<ul style="list-style-type: none">The overlay name the zone belongs to.				
direction	default		variable length	26
<ul style="list-style-type: none">The direction of the flow: outbound or inbound.				

Logging Tab

Administration > General Settings > Setup > Logging

This tab summarizes the following configured logging parameters:

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

The logs keep track of alarms, events, or any other issue involved with your appliances. The following table provides more details.

Severity Levels

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

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

- The **bolded** part of the name is what displays in Silver Peak 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, when they clear, list as the **ALERT** level in the **Event Log**.

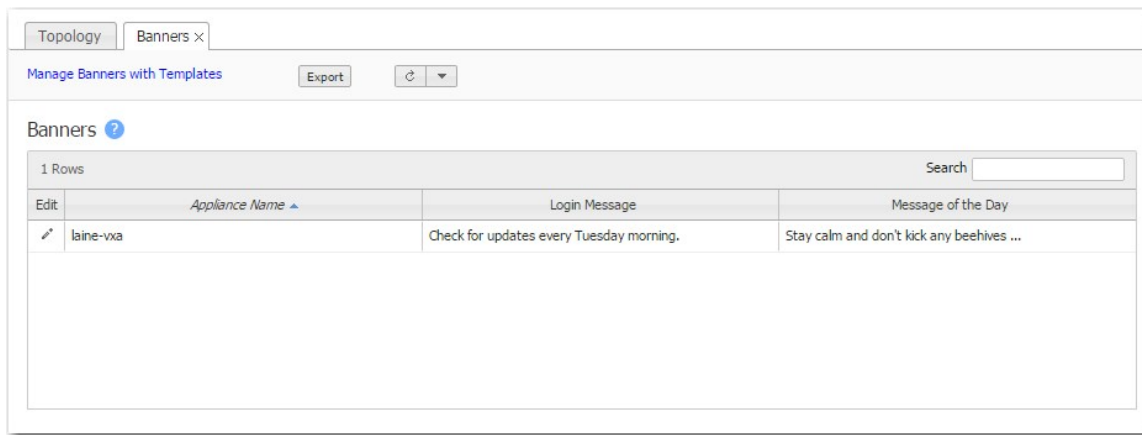
Remote Logging

- 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 might 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.



Each appliance can have two **banner messages**:

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

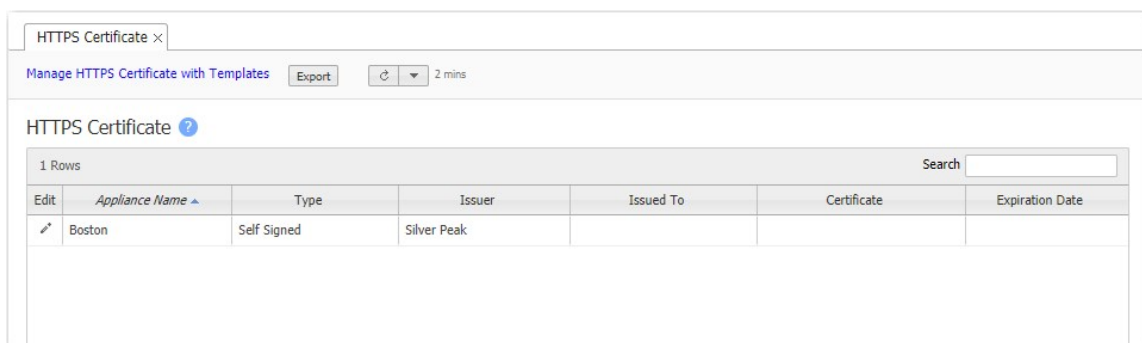
Click the edit icon to enter your banner message.

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.



To use a custom certificate 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, and so forth.

- 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, 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

10.0.185.23

Use Orchestrator Management IP

Add

1 Rows, 1 Selected		Search	
Label	Orchestrator IP or Domain Name	Priority	
MPLS	10.0.185.23	1	X

MPLS

MPLS

Internet

LTE

Internet2

Save Close

Custom Appliance Tags

Administration > General Settings > Setup > Custom Appliance Tags

Use this tab to create and assign tags to an appliance or a group of appliances. A tag acts as a filter or identity when searching for appliances. Complete the following steps to create a custom tag.

1. Click the edit icon.
2. Click the selected row in **Key**, and then enter the name of the tag you want to use.
3. Click the selected row in **Value**, and then enter a brief description of what the tag represents.
4. Click **Apply**.

NOTE You can create up to eight tags.

System Information

Administration > Software > Upgrade > System Information

You can manage system information with templates, except for Deployment Mode, which is an appliance-specific configuration. To change a Deployment Mode, navigate to **Configuration > Networking > Deployment**.

When you click the **Edit** icon next to a specific appliance, the following two screens are available:

System Summary

System Information - ecvb	
<div>System Summary System Settings ?</div>	
General	
Appliance Key	1.NE
Platform	VMware
Uptime	2d 5h 49m 18s
Active Release	9.0.3.0_89659
Appliance ID	634918
Discovery Method	PORTAL
Connection Type	WEBSOCKET
Configuration	
System Bandwidth	4000 Kbps
Mode	router
Hardware	
Appliance Model	EC-V 208001009006 Rev 76564
BIOS Version	6.00
Serial Number	00-1B-BC-09-B0-26
<div>Apply Cancel</div>	

System Settings

System Information - ecvb

System Summary
System Settings

General

Model EC-V 208001009006 Rev 76564

Serial 001BBC09B026

Site Name

Hub Site? No

Contact Name

Contact Email

Location Santa Clara, California, 95050, US

Region Default

Optimization

IP Id auto optimization

TCP auto optimization

Flows and tunnel failure fail-stick

Network Memory

Encrypt data on disk

Configured Media Type ram and disk

Media Type ram and disk

Shell Access

Shell Access Status Open Shell Access

Excess Flow Handling

Excess flow policy bypass

NextHop Health Check

Enable Health check

Retry count 4 (1..255)

Interval 10 (1..255) seconds

Hold down count 1 (1..255)

Miscellaneous

SSL optimization for non-IPSec tunnels

Bridge Loop Test

Enable IGMP snooping (Not Available in 9.0.3.0)

Auto Flow Re-Classify 60 (0..65535) seconds

Always send pass-through traffic to original sender

IPSec UDP Port 12000

Enable default DNS lookup

Enable HTTP/HTTPS snooping

Quiescent tunnel keep alive time 60 (1..65535) seconds

UDP flow timeout 120 (1..65535) seconds

Non-accelerated TCP Flow Timeout 1800 (1..65535) secs

Maximum TCP MSS 9000 (500..9000) bytes

NAT-T keep alive time 300 (0..65535) seconds

Tunnel Alarm Aggregation Threshold 5 Tunnel Alarms

Raise only 1 alarm above this threshold

Maintain end-to-end overlay mapping

IP Directed Broadcast

Allow WAN to WAN routing

Apply
Cancel

System Information Property Keys

Property Key	Description
Active Release	Specifies the software release the appliance is running.
Allow WAN to WAN routing	Redirects inbound LAN traffic back to the WAN.
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	Unique identifier for the appliance.
Appliance Key	Orchestrator assigns and uses this key to identify the appliance.
Appliance Model	Specific EC, EC-V, NX, VX, or VRX model.
Auto Flow Re-Classify	Specifies how often to do a policy lookup.
BIOS Version	Version of BIOS firmware that the appliance is using.

Property Key	Description
Bridge Loop Test	Only valid for virtual appliances. When enabled, the appliance can detect bridge loops. If it detects a loop, the appliance stops forwarding traffic and raises an alarm. Appliance alarms include recommended actions.
Configured Media Type	Is either ram and disk (VX) or ram only (VRX). Can be changed for special circumstances if recommended by Silver Peak.
Connection Type	Method that Orchestrator uses to communicate with the appliance. Options are WEBSOCKET, PORTAL, and HTTP.
Contact Email	Email address of the person to contact within your organization (optional).
Contact Name	Name of the person to contact within your organization (optional).
Discovery Method	Specifies how Orchestrator discovered the appliance: <ul style="list-style-type: none"> ■ PORTAL – Orchestrator discovered the appliance through the portal account. ■ MANUAL – The appliance was added manually. ■ APPLIANCE – Orchestrator's IP address was added to the appliance. Portal was not involved.
Enable default DNS lookup	Allows the appliance to snoop the DNS requests to map domains to IP addresses. This mapping then can 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.
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 bypass flows. Or, you can choose to drop the packets.

Property Key	Description
Flows and tunnel failure	<p>If there are parallel tunnels and one fails, Dynamic Path Control determines where to send the flows. There are three options:</p> <ul style="list-style-type: none"> ▪ fail-stick – When the failed tunnel comes back up, the flows do not return to the original tunnel. They stay where they are. ▪ fail-back – When the failed tunnel comes back up, the flows return to the original tunnel. ▪ disable – When the original tunnel fails, the flows are not routed to another tunnel.
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.
Interval	Specifies the number of seconds between each ICMP echo sent.
IP Directed Broadcast	Allows an entire network to receive data that only the target subnet initially receives.
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.
Maintain end-to-end overlay mapping	Enforces the same overlay to be used end-to-end when traffic is forwarded on multiple nodes.
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.
Mode	Specifies the appliance's deployment mode: Server, Router, or Bridge.
Model	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 to keep the mappings in the NAT device intact.
Non-accelerated TCP Flow Timeout	Specifies how long to keep the TCP session open after traffic stops flowing. The default is 1800 seconds (30 minutes).
Platform	Underlying cloud platform on which the EdgeConnect appliance runs, such as Silver Peak, Amazon EC2, Azure, Google Cloud, or VMware.

Property Key	Description
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	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 / Serial Number	Serial number of the appliance.
Shell Access Status	<p>Specifies the current shell access policy for EdgeConnect appliances.</p> <ul style="list-style-type: none"> ▪ Open Shell Access – Full access granted to the underlying Linux operating system shell. ▪ Secure Shell Access – Access denied to the shell, but Support can grant access. Contact Support for assistance. You cannot change this setting to Open Shell Access. ▪ Disabled Shell Access – Access permanently denied to the shell. You cannot change this setting to Open Shell Access or Secure Shell Access. <p>This setting is managed on the Advanced Security Settings page (Configuration > Overlays & Security > Security > Advanced Security Settings). Changes to this setting affect all appliances in your network.</p>
Site Name	Orchestrator will not build tunnels between appliances with the same user-assigned site name.
SSL optimization for non-IPSec tunnels	Specifies whether 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 EdgeConnect appliances or just to a specified group of appliances.
System Bandwidth	Appliance's total outbound bandwidth, determined by appliance model or license.
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).
Tunnel Alarm Aggregation Threshold	Specifies the number of alarms to allow before alerting the tunnel alarm.
UDP flow timeout	Specifies how long to keep the UDP session open after traffic stops flowing. The default is 120 seconds (2 minutes).
Uptime	Time elapsed since the appliance became operational and available.

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 <input type="text"/>	
	Partition 1
Appliance Name	Build Version
Tallinn	8.1.1.0_60681
laine-va	8.1.5.1_65516
laine-vb	8.1.5.1_65516

Upgrade 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** dialog box 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

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.

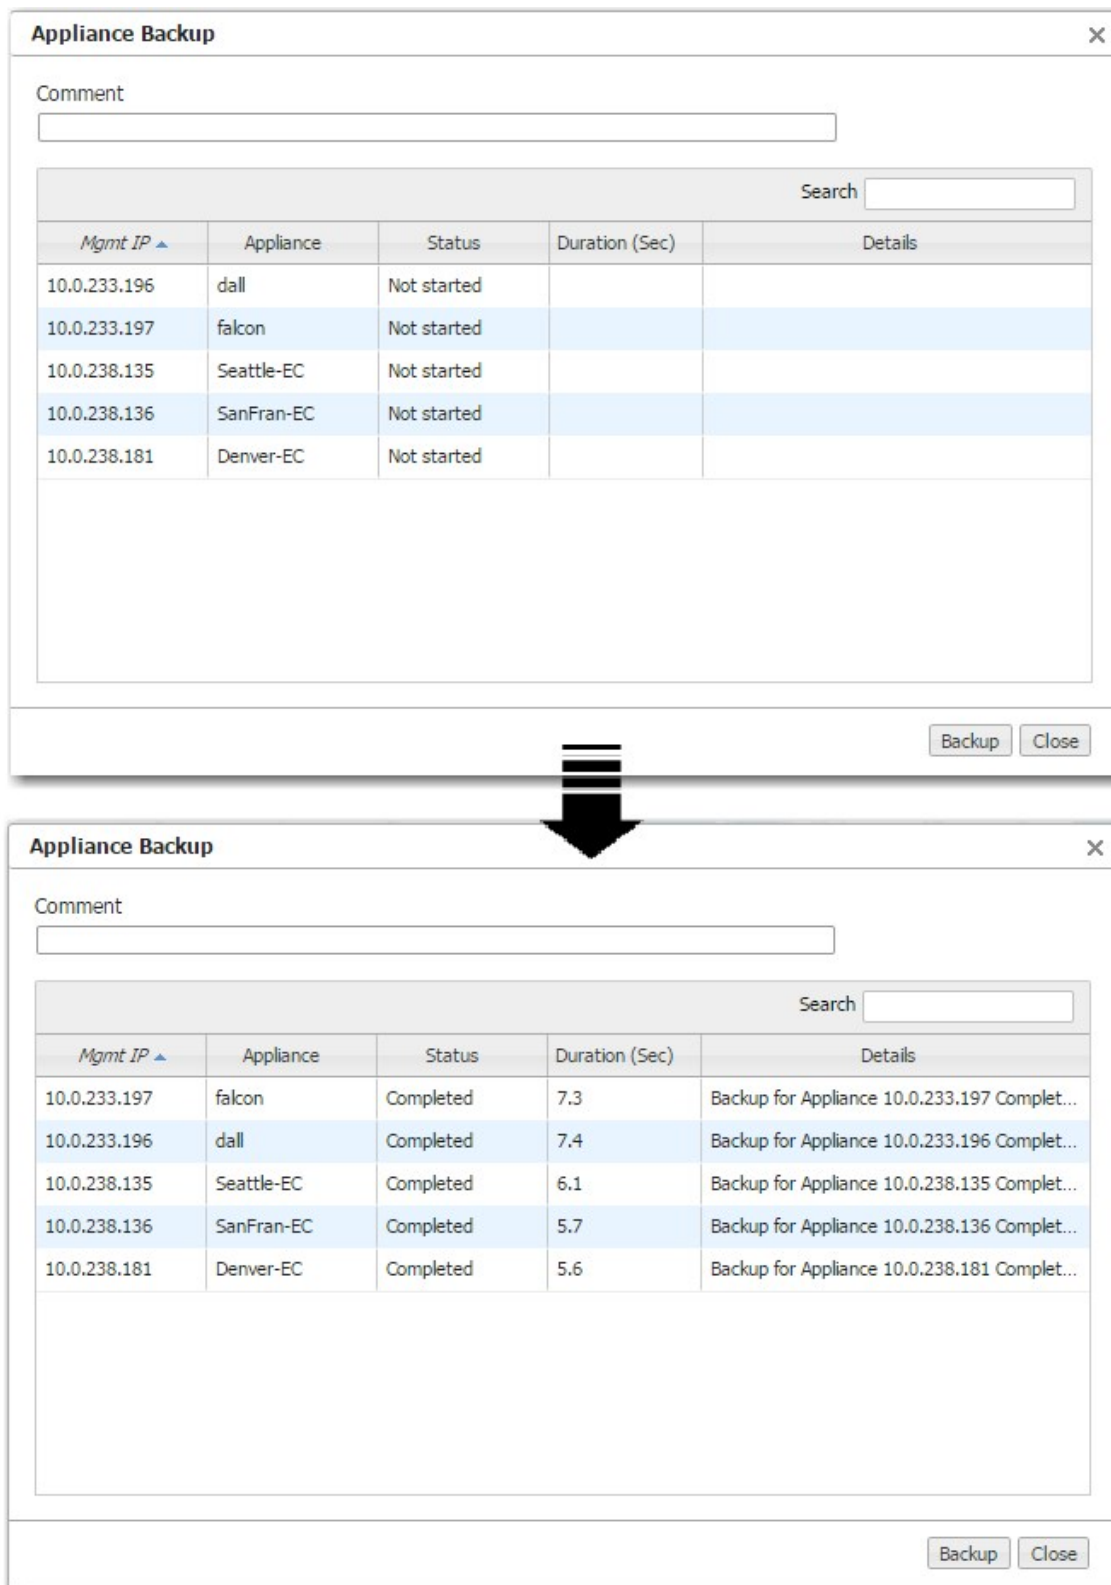
- **Install and reboot** installs the image into the appliance's inactive partition and then reboots the appliance to begin using the new software.
- **Install and set next boot partition** installs the image into the appliance's inactive partition and then points to that partition for the next reboot.
- **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 appliance tree, navigate to **Administration > Software > Backup & Restore > Backup Now**, and then click **Backup**.



The diagram illustrates the process of backing up appliances. It shows two sequential screenshots of the 'Appliance Backup' window, connected by a large downward-pointing arrow. The top screenshot shows a table with five appliances, all with a status of 'Not started'. The bottom screenshot shows the same five appliances, now with a status of 'Completed', and each with a duration and a detailed backup description.

Appliance Backup [X]

Comment

Search

Mgmt IP ▲	Appliance	Status	Duration (Sec)	Details
10.0.233.196	dall	Not started		
10.0.233.197	falcon	Not started		
10.0.238.135	Seattle-EC	Not started		
10.0.238.136	SanFran-EC	Not started		
10.0.238.181	Denver-EC	Not started		

Backup Close

Appliance Backup [X]

Comment

Search

Mgmt IP ▲	Appliance	Status	Duration (Sec)	Details
10.0.233.197	falcon	Completed	7.3	Backup for Appliance 10.0.233.197 Complet...
10.0.233.196	dall	Completed	7.4	Backup for Appliance 10.0.233.196 Complet...
10.0.238.135	Seattle-EC	Completed	6.1	Backup for Appliance 10.0.238.135 Complet...
10.0.238.136	SanFran-EC	Completed	5.7	Backup for Appliance 10.0.238.136 Complet...
10.0.238.181	Denver-EC	Completed	5.6	Backup for Appliance 10.0.238.181 Complet...

Backup Close

You cannot delete an appliance backup from Orchestrator.

View 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	View	X
laine-vxb (10.0.238.69)	initial	28-Dec-16 04:0...	8.1.1.0_60681	View	X
laine-vxa (10.0.238.71)	initial	2-Nov-16 05:00:...	8.1.1.0_60681	View	X
laine-vxb (10.0.238.69)	initial	2-Nov-16 05:00:...	8.1.1.0_60681	View	X
laine-vxa (10.0.238.71)	initial	1-Nov-16 05:00:...	8.1.1.0_60681	View	X
laine-vxb (10.0.238.69)	initial	1-Nov-16 05:00:...	8.1.1.0_60681	View	X

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:E5:96:B4
5 interface mgmt0 mac address 00:0C:29:19:53:A1	5 interface mgmt0 mac address 00:0C:29:E5:96:96
6 interface mgmt1 mac address 00:0C:29:19:53:AB	6 interface mgmt1 mac address 00:0C:29:E5:96:A0
7 interface wan0 mac address 00:0C:29:19:53:B5	7 interface wan0 mac address 00:0C:29:E5:96: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

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
```

Restore a Backup to an Appliance

Administration > Software > Backup & Restore > Restore

You can restore an appliance configuration backup from Orchestrator to any other EdgeConnect appliances 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: laine-vxa
Target Appliance: laine-vxb

select an appliance
laine-vxa
laine-vxb
Tallinn

Select the configuration to restore from the table below

File Name	Backup Time	Software Vers		Comment
initial	28-Dec-16 04:00:05	8.1.1.0_60681	View	
initial	2-Nov-16 05:00:05	8.1.1.0_60681	View	
initial	1-Nov-16 05:00:05	8.1.1.0_60681	View	
initial	15-Sep-16 05:00:05	8.1.1.0_60681	View	
initial	14-Sep-16 05:00:05	8.1.1.0_60681	View	
initial	26-Aug-16 05:00:06	8.1.1.0_60681	View	
initial	25-Aug-16 05:00:05	8.1.1.0_60681	View	
initial	12-Jul-16 05:00:06	8.1.1.0_60681	View	
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

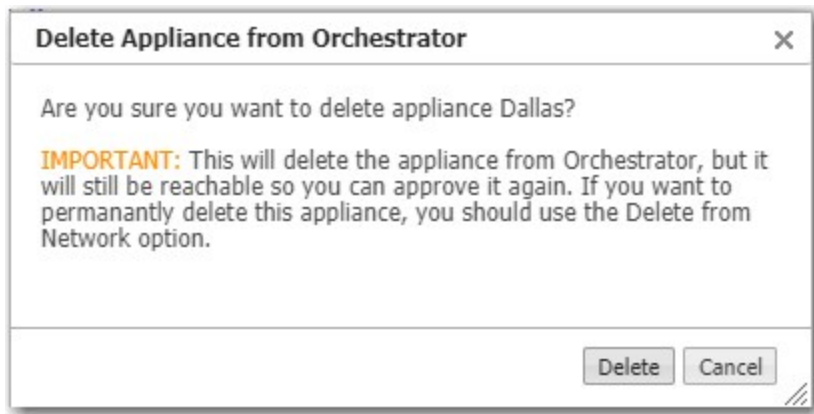
Status Log

Restore
Close

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, and so forth 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, and so forth to this device.
- It tells the Portal to "unlicense" the appliance.

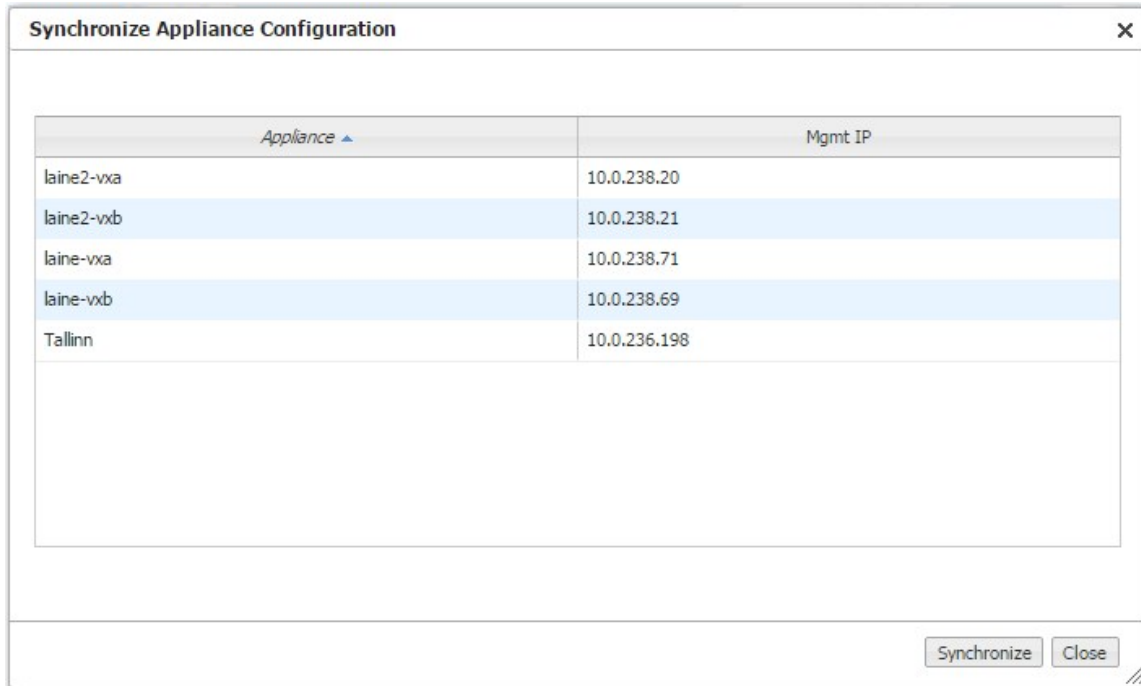
Synchronize Appliance Configuration

Administration > Tools > Synchronize

Orchestrator keeps its database synchronized with the running configurations for the appliances.

- When you use Orchestrator to make a configuration change to an appliance's 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, you can use this dialog box to manually resync to ensure that Orchestrator has an appliance's current running configuration.



Put the Appliance in System Bypass Mode

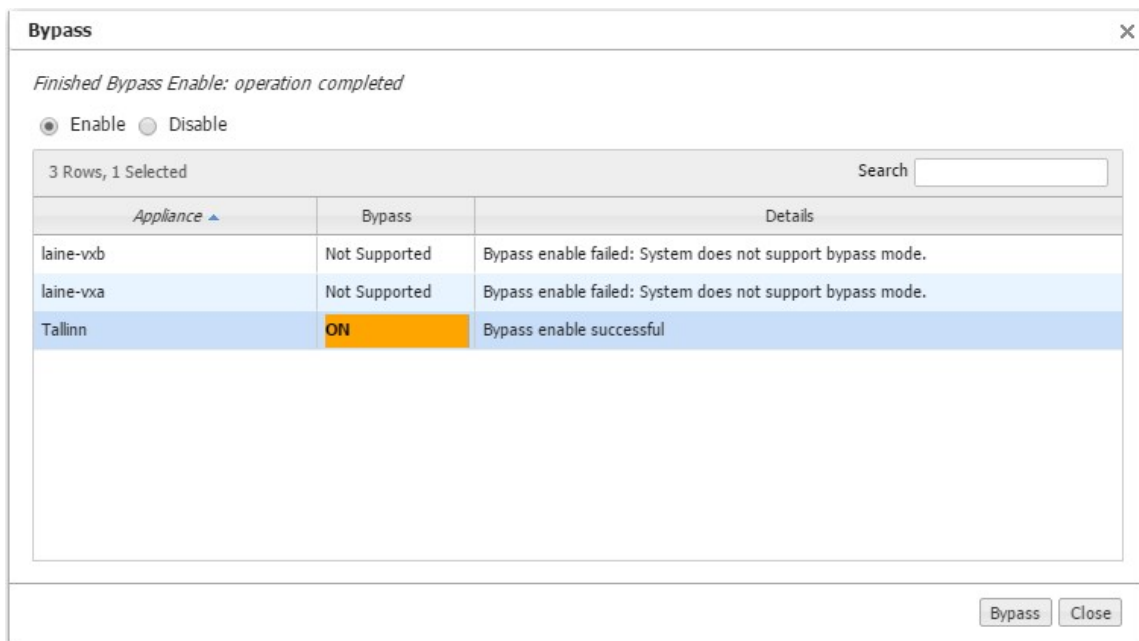
Administration > Tools > Bypass

System Bypass mode is only available for certain models of EdgeConnect physical appliances. Virtual appliances do not 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.



When the appliance is in Bypass mode, a message displays in red text in the upper-right corner of the user interface.



Broadcast CLI Commands

Administration > Tools > Broadcast CLI

You can simultaneously apply Command Line Interface (CLI) commands to multiple, selected appliances.

The dialog box automatically provides you with the highest user privilege level.



INFO 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 enables you to measure the throughput and integrity (amount of loss) of your WAN link. You can run either **iperf** or **tcppperf** (Version 1.4.8).

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>

Option	Description
-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 **-l**, *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 tcpperf to use IPv6 addresses only. Default is IPv4 addresses.
-l	<i>interface IP</i> : Specify source interface IP address. Default is any .
-o	<i>outname</i> : Output filename. Default is stdout .
-u	<i>update <secs></i> : Frequency of printed updates in seconds. Default is 1 .
-d	<i>duration <secs></i> : Set maximum test duration in seconds. Default is infinite .
-w	<i>wait <secs></i> : Wait until <secs> since 1970 before transmitting data.
-z	<i>realtime</i> : Elevate to realtime priority. Requires root privilege.
-cm	<i>cpu mask</i> : Specify CPU affinity. Requires root privilege.
-q	<i>quiet <level></i> : Suppresses detail based on level: <ul style="list-style-type: none"> ■ 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.

Option	Description
-tm	<i>tcpmss</i> : TCP mss. Default is OS default.
-tn	<i>tcpnodelay</i> : TCP nodelay option. Default is nagle enabled.
-tq	<i>tcpquickack</i> : TCP quick ack option. Default is delayed acks.
-td	<i>tcpdscp</i> <cp>: Sets IP DSCP to <cp> (decimal). Default is 0 .
-tr	<i>tcpretries</i> <n>: Sets number of times to retry TCP connections.
-tp	<p>tcppace <n> [mode]: Pace TCP connection setup rate. Limits number of half-open connections to <n>.</p> <p>Valid <mode> types are:</p> <ul style="list-style-type: none"> ▪ preestablish. All connections are established before data transmission. Default. ▪ simultaneous. Begin data transmission as soon as connection made.
-ta	<i>tcpabort</i> : Sends RSTs instead of FINs on close.
-tf	<i>tcpfindelay</i> <secs>: Time to wait after all data is 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</i> <i>: Set test ID. The same test ID produces the same data set. User different test IDs to generate unique data for each test run. Default is zero.
-n	<i>number</i> <n>: Generate <n> flows. Default is one.
-b	<i>begin</i> <byte>: First byte in transmission. Default is zero.
-e	<p><i>end</i> <byte>: End byte in transmission (number of bytes to transmit). Default is file size.</p> <p>Begin and end bytes can be greater than file size. The content is repeated to create extra bytes.</p>
-a	<p><i>antipat</i> <mode>: Antipattern mode: default is mutate:</p> <ul style="list-style-type: none"> ▪ none – Repeats same content verbatim on all flows. Repeats content if end byte exceeds content size. ▪ mutate – Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Destroys cross flow similarity. Destroys original byte code distribution. ▪ shuffle – Ensures all flows and data repeats are unique. Preserves short range patterns within flow. Preserves cross flow similarity. Preserves original byte code distribution. ▪ fast – Ensures all flows and data repeats are unique. Does not preserve short range patterns. Destroys cross flow similarity. Destroys original byte code distribution. Uses less CPU than mutate or shuffle.

Option	Description
-l	<p><i>loopback [mode]</i>: Loopback. Default is unidirectional.</p> <ul style="list-style-type: none"> ■ uni – Unidirectional client to server. ■ rev – Unidirectional server to client. ■ bidir – Bidirectional, client and server independently send data on the same TCP connection. ■ bidir2 – Bidirectional, client and server independently send data on secondary TCP connections. ■ loop – Bidirectional, server loops data back to client on the same TCP connection. ■ loop2 – Bidirectional, server loops data back to client on a secondary TCP connection. ■ bidir2 – Bidirectional, transmits one transaction at a time. Client waits for previous transaction to be echoed. Emulates transactional data. <p>NOTES:</p> <ol style="list-style-type: none"> 1. Content source for traffic originating at the server is determined by the server (not client) command line. 2. loop2 and bidir2 modes 2 x <n> TCP connections and requires that the server has even-numbered ports available.
-r	<i>rate <bps></i> : Limits aggregate transmission rate to <bps>. Default is no rate limit.
-t	<p><i>trans <min> [max]</i>: Sets size of each socket transaction. Default is 64000.</p> <p>If <min> and <max> are specified, client generates transactions with random sizes between <min> and <max>. This feature is often used with -l and -r. Set the minimum transaction size to 100000 to improve single-flow performance.</p>
-v	<p><i>verify <mode></i>: Verify integrity of received data. Default is global.</p> <ul style="list-style-type: none"> ■ none – No verification. Fastest/least CPU load. ■ global – Single global hash per flow. Fast, but cannot isolate an errored block. ■ literal – Literal comparison of data upon reception. Fast, can isolate errors to the byte level. Requires that server has same content as client. Use random data gen or same -f file at server. ■ embedded – Embedded hashes every 4096 bytes. Slower, can isolate errors to 4096 byte block.
-p	<p><i>repeat <n></i>: Repeat each content byte n times. Default is 1 (no repeats).</p> <p>Works for both random data and file content.</p>
-k	<p><i>corrupt <n> <m> <s> [<%change>[<%insert>[<%delete>]]]</i>: Corrupt 0 to n bytes of data every m bytes using seed s. Delta bytes will require 0.5*n/m percent overhead. Each corrupt can 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>

Option	Description
-x	<i>excerpts</i> <e> </> [<i>s</i>]: Send random excerpts of average <l> length bytes from content between egin and <e>nd bytes. The -b and -e options still specify total bytes to send. Uses random seed <i>s</i> .
-y	<p><i>defred</i> <s%> <m%> <l%> <sb> <smin> <smax> <mb> <mmin> <mmax> <lb> <lmin lmax> :</p> <p>Generate content based on defined reduction model.</p> <p>Content is drawn from three data sets: s, m, and l:</p> <ul style="list-style-type: none"> ▪ s% – Specifies fraction [50%] of s-type content (short term reducible). ▪ m% – Specifies fraction [30%] of m-type content (medium term reducible). ▪ l% – Specifies fraction [20%] of l-type content (long term reducible). <p>Short term content comes from data set of sb Mbytes [100MB] with excerpts uniformly distributed between smin and smax bytes [10K-1M].</p> <p>Medium term content comes from data set of mb Mbytes [100GB] with excerpts uniformly distributed between lmin and lmax bytes [10K-1M].</p> <p>Long term content comes from data set of lb Mbytes [100TB] with excerpts uniformly distributed between smin and smax bytes [10K-1M].</p> <p>The -b and -e options still specify total bytes to send.</p> <p>Performance is best if -b is 0.</p> <p>Uses random seed <i>s</i>.</p>
-ssl [param=value ...]	<p>Enable SSL on connection with optional parameters.</p> <ul style="list-style-type: none"> ▪ version=2 3 t10 t11 t12. Set the protocol version. ▪ cipher=OPENSSL-CIPHER-DESC. Set the choice of ciphers. ▪ ticket=yes no. Enable/disable session ticket extension. ▪ cert=FILENAME. Use this certificate file. ▪ key=FILENAME. Use this private keyfile. ▪ compression=none any deflate zlib rle. Set the compression method. ▪ sslcert. Print the SSL certificate in PEM format. ▪ sslkey. Print the SSL key in PEM format.

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 (Redundant Array of Independent Disks) arrays with encrypted disks.
- Disk failure results in a **critical alarm**.

- If a row indicates that a disk has failed, click the edit icon to access the appliance, and then follow the directions in the local help to replace 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' tab in the Silver Peak Unity Orchestrator. It displays a table with 10 rows of disk information. An orange circle highlights the 'Edit' icon (a pencil) in the first row, which corresponds to the 'Tallinn' appliance. A callout box points to this icon with the text: 'For example, to access Tallinn's own Disk Management page, click **any** of these Edit icons.'

The table has the following columns: Edit, Appliance Name, Appliance Model, Slot ID, Pairing Slot ID, Status, Size (GB), Serial Number, Removable, and SMART Data. The first two rows are for 'Tallinn' (NX-3700) with Slot IDs 0 and 1, both showing 'OK' status and 465 GB size. The remaining rows are for other appliances like 'laine-vxa' and 'laine-vxb'.

A pop-up window titled 'Smart data for Tallinn ID 1' is shown below the table. It contains a table with 21 rows of SMART data. The columns are Attribute, Normalized Value, Worst Value, and Raw Value. The data includes various metrics such as Read Error Rate, Spin up time, Start/stop count, Reallocated sector count, Seek error rate, Power on hours, Spin retry count, Device power cycle count, End to End error, Reported Uncorrectable Errors, Command timeout, High Fly Writes, and Temperature difference.

To replace a failed disk:

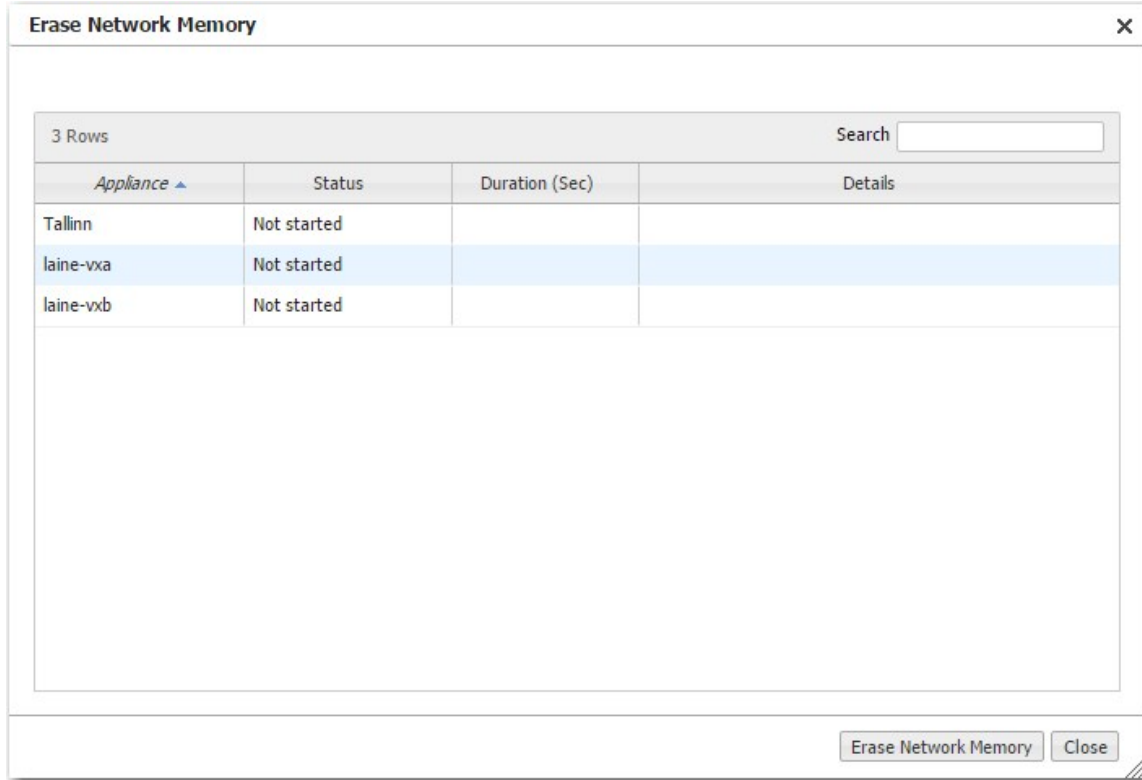
1. Log in to your Support portal account, and then click **Open a Self Service RMA** for disk replacement.
2. Complete the wizard. Use the serial number of the appliance (not the disk).
3. After you receive the new disk, access Appliance Manager by clicking any edit icon that belongs to the appliance in question.
4. Follow the instructions on that page's online help.

Erase Network Memory

Administration > Tools > Erase Network Memory

Erasing Network Memory removes all stored local instances of data.

No reboot required.

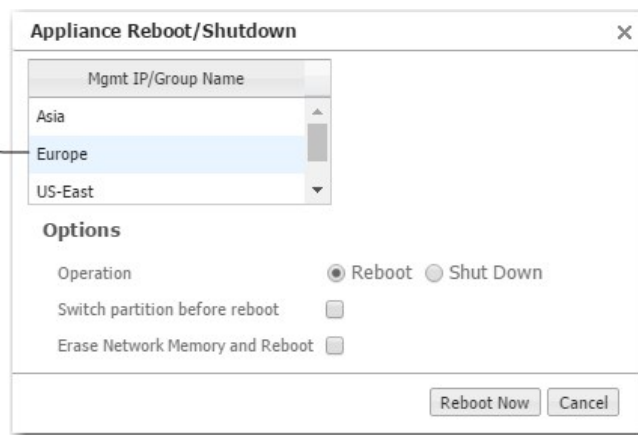


Reboot or Shut Down an Appliance

Administration > Tools > Reboot > Appliance Reboot / Shutdown

The appliance supports three types of reboot:

Displays selected appliances



- **Reboot.** Reboots the appliance gracefully. This is your typical "vanilla" restart.

Use case: You are 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 are 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:

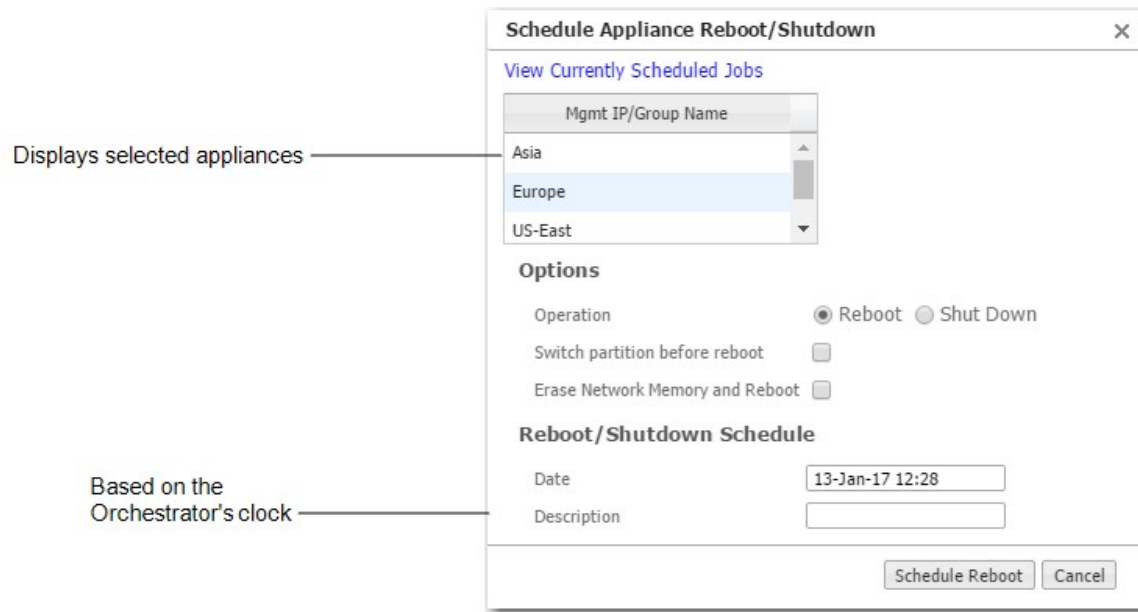
- *hardware bypass*, if deployed in-line (Bridge mode), or
- *an open-port state*, if deployed out-of-path (Router mode or Server mode).

Unless a **virtual appliance** is configured for a high availability deployment, all flows are discontinued during reboot.

Schedule 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 are 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 are decommissioning the appliance.
 - You need to physically move the appliance to another location.
 - You need to re-cable 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.

INFO To specify the timezone for scheduled jobs and reports, navigate to **Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs**.

Reachability Status Tab

Administration > Tools > Monitoring > Reachability Status

This tab summarizes the status of communications in two directions—**Orchestrator to Appliances** and **Appliances to Orchestrator**.

Orchestrator to Appliances 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

Appliances to Orchestrator 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 in to 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

This tab lists users who are logged in to Orchestrator and the appliances that Orchestrator is currently managing.

To list active user sessions, click **Orchestrator**.

Active Sessions x

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

To list active appliance sessions, click **Appliance**.

Active Sessions x

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

Role Based Access Control provides a more customized Orchestrator experience. On a per-user basis, you can assign roles that specify access levels for a user, control the menu options available in the Orchestrator UI, and grant or deny access to appliance groups.

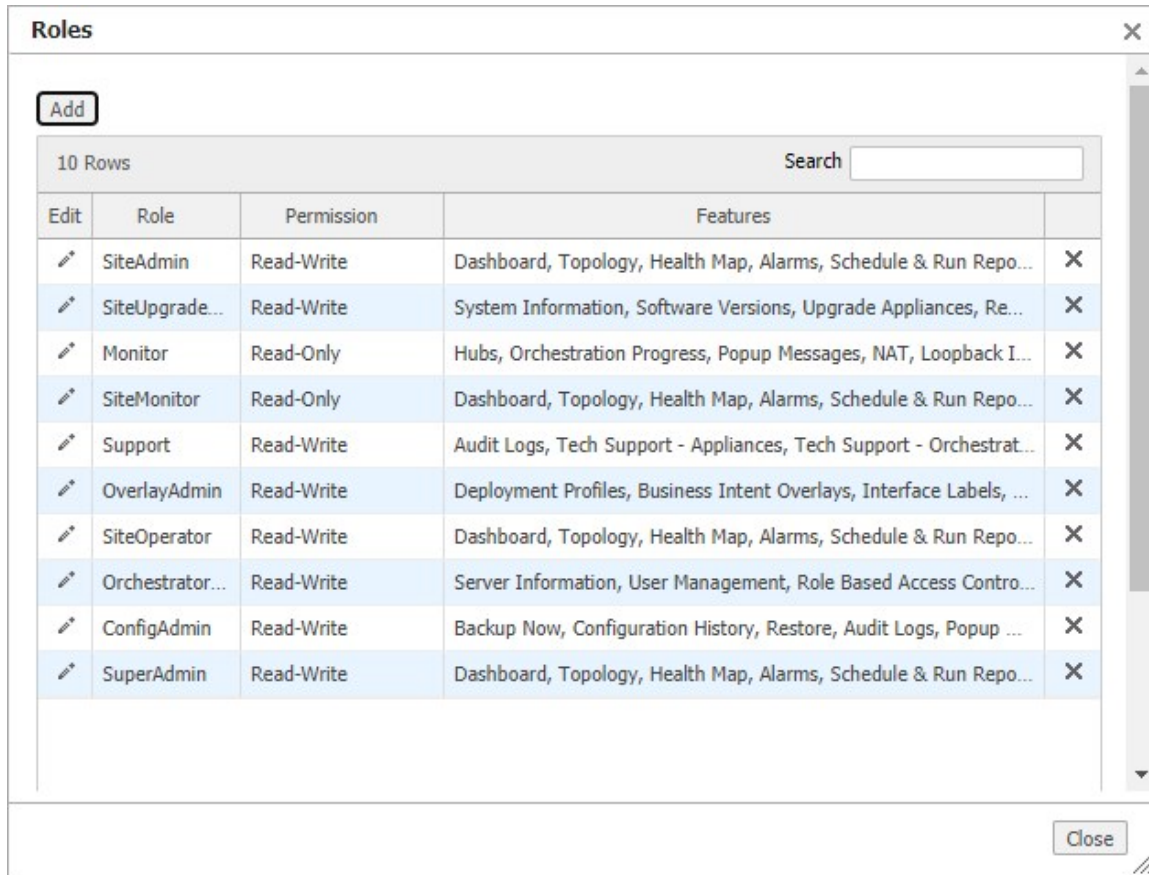
Roles

Orchestrator provides a set of default roles. You can create new roles or modify an existing role.

Field	Description
Role	Name of the role.
Permission	Overall access level assigned to the selected role, Read-Write or Read-Only .
Features	Orchestrator features available to the selected role.

To add a role:

1. Click **Create Roles**. The Roles dialog box opens.



2. Click **Add** to create a new role, or click the **Edit** icon to the left of any existing role.
3. Enter or modify the role name.
4. Select a category you want to assign to your user from the following tabs: **Monitoring**, **Configuration**, **Administration**, **Orchestrator**, **Support**, or **Miscellaneous**.
5. Select **Read Only** or **Read & Write** to assign the overall access level for the role.
6. Select the check box corresponding to the Orchestrator menu options you want to make available to the role.

NOTE You can **Select All** or **Unselect All**.

7. Click **Save**.

Appliance Access

With appliance access groups, you can restrict appliance access to one or more groups or regions. Complete the following steps to customize appliance access.

1. Click **Create Appliance Access Groups** on the **Role Based Access Control** tab. The **Appliance Access Group** dialog box opens.

Edit	Appliance Access Group	Groups / Regions	
	US WEST	-- US WEST --, US-WEST	
	APJ	APJ, AUS, WEST	

2. Click **Add** to create a new group, or click the **Edit** icon to the left of any existing group.
3. Add or modify the name of the appliance access group.
4. Choose how you want to add appliances: **Select By Groups** or **Select By Region**. You can manually select groups or regions to include, or use the buttons to select or clear all options.
5. Click **Save**.

WARNING A non-RBAC user or an RBAC user with appliance access and no assigned role will have access to the Appliance Manager, CLI Session, and Broadcast CLI. An RBAC user with any role assigned will be denied access to the Appliance Manager, CLI Session, and Broadcast CLI.

User	Appliance Access	Roles?	Menu Options
Non-RBAC User	N/A	N/A	Appliance Manager, CLI Session, Broadcast CLI
RBAC User	Yes	None assigned	Appliance Manager, CLI Session, Broadcast CLI
RBAC User	No	Any	Appliance Manager, CLI Session, and Broadcast CLI will be denied

Assign Roles and Appliance Access

Complete the following steps to assign roles and appliance access.

1. On the **Role Based Access Control** tab, select **Assign Roles & Appliance Access Groups**.
2. Click in the **User** field and enter a name of an existing Orchestrator user.
3. Click the arrow in the **Appliance** field and select the name of an existing Appliance Access Group.
4. Select the check boxes for one or more roles you want to assign to the user.
5. Click **Save**.

The following table defines the roles that are provided by default in Orchestrator (roles are listed alphabetically).

Role	Description
ConfigAdmin	Back up and restore appliance configuration and view the configuration history.
OrchestratorAdmin	Allows you to perform Orchestrator operations only , such as settings, tools, user management, and Orchestrator upgrades. Appliance operations are not allowed.
OverlayAdmin	A global role for managing SD-WAN overlays. NOTE Overlay management cannot be specific to a site or region.
SiteMonitor	Read-only permissions equivalent to SiteAdmin.
SiteOperator	Enables appliance or site specific operations, such as configuring appliance specific policies, ACLs, TCAs, and SSL certificates. You cannot upgrade an appliance or remove it from the network, or perform global SD-WAN functions such as overlay management or Zscaler orchestration.
SiteUpgradeAdmin	Upgrade appliances and remove them from the network.
SuperAdmin	Enables full read-write access to all menu items.
SiteAdmin	Enables appliance or site-specific operations, such as configuring appliance specific policies, ACLs, TCAs, SSL certificates, and upgrades. You cannot remove an appliance from the network or perform global SD-WAN functions such as overlay management or Zscaler orchestration.
Support	Enables access to all support operations.
Monitor	Provides read-only access to all menu items.

View Orchestrator Server Information

Orchestrator > Orchestrator Server > Server Management > Server Information

This dialog box provides data specific to this Orchestrator server.

A screenshot of a dialog box titled "Orchestrator Server Information" with a close button (X) in the top right corner. The dialog box contains a table with server details. The table has two columns: the left column lists system metrics and the right column lists their values. The data is as follows:

Orchestrator Server Information	
Orchestrator Hostname	DMerwin-GXV
Serial Number	10-10-10-10-10-10
Uptime	1d 5h 10m 31s
Time	Thu Jun 25 19:47:11 PDT 2015
Used disk space	26G
Number of CPUs	4
Model	GX-V
IP Address	10.10.10.10
Active users	2
Load Average	0.00, 0.01, 0.05
OS Version	2.6.35.14-106.fc14.x86_64
Free disk space	57G
Memory (MB)	3964
Revision	6.0.0.0

A "Close" button is located at the bottom right of the dialog box.

Restart, Reboot, or Shutdown

Orchestrator > Orchestrator Server > Server Management > Reboot Orchestrator

Orchestrator > Orchestrator Server > Server Management > Shutdown Orchestrator

Orchestrator provides these two convenient actions in the **Orchestrator** menu:

- **Reboot Orchestrator** reboots the Orchestrator server.
- **Shutdown Orchestrator** results in the server being unreachable. You will have to manually power on the server to restart.

Manage Orchestrator Users

Orchestrator > Orchestrator Server > Users & Authentication > User Management

The **User Management** page enables you to manage who has Read-Write or Read-Only access to Orchestrator.

User Management

Auto Logout: (1-10080 minutes)

Max Sessions: (5-10000)

10 Rows Search

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	✕
	sbheemaraj...					No	No	18-Apr-18 0...	Active	Read-Write	✕
	srinivas					No	No	30-Apr-18 1...	Active	Read-Write	✕
	syen	Shyh-Pei	Yen		syen@silver...	No	No	07-May-18 0...	Active	Read-Write	✕
	anusha-ro	anusha-ro	read-only			No	No	21-May-18 1...	Active	Read-Only	✕

Add 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

Orchestrators support Multi-Factor Authentication (MFA). This is available on all platforms of the 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.

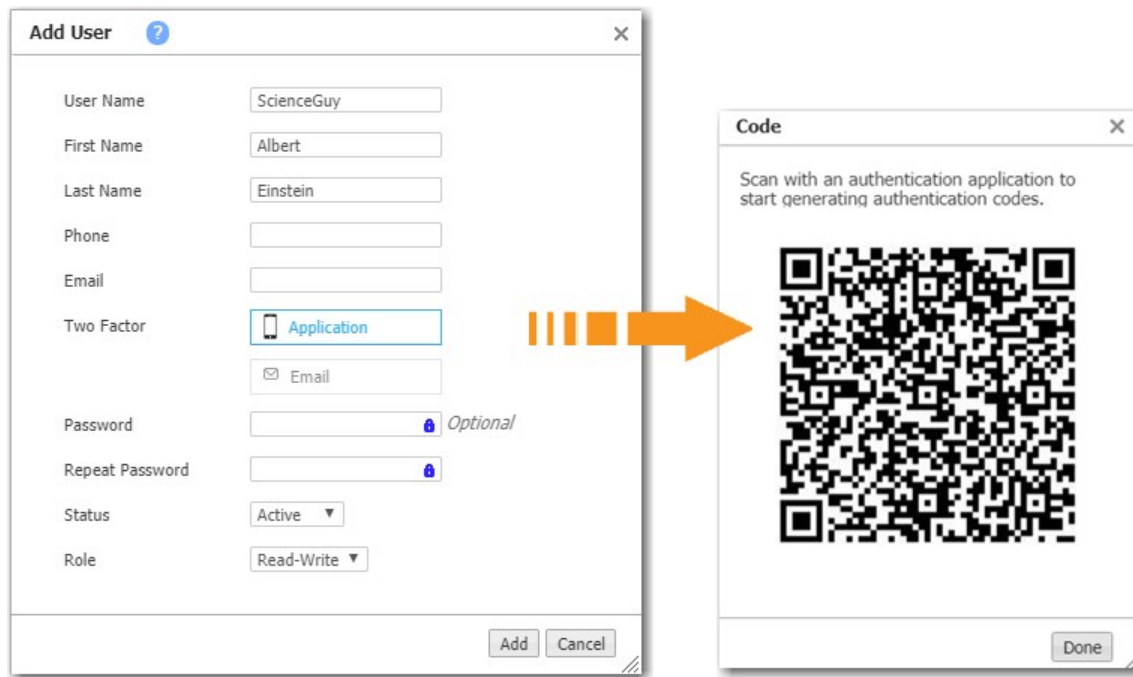
NOTE Currently, only **admin** users can only configure Multi-Factor Authentication, and only for themselves.

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.

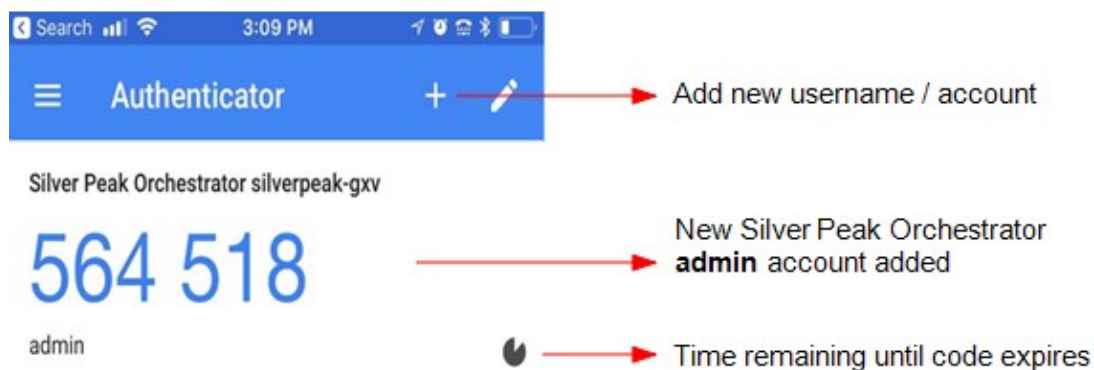
- To enable Multi-Factor Authentication, navigate to **Orchestrator > Orchestrator Server > Users & Authentication > User Management**, and then click on your username.

- For **Two Factor**, click **Application**. Orchestrator generates a time-limited QR code.



- With the Google Authenticator app, use the **Scan barcode** function to read the QR code. You also will 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

- To enable Multi-Factor Authentication, navigate to **Orchestrator > Orchestrator Server > Users & Authentication > User Management**, and then 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 box, Orchestrator sends a time-limited authentication code to your email address. To verify your email address, click that link.

Orchestrator then opens a browser window telling you that your email address has been verified.

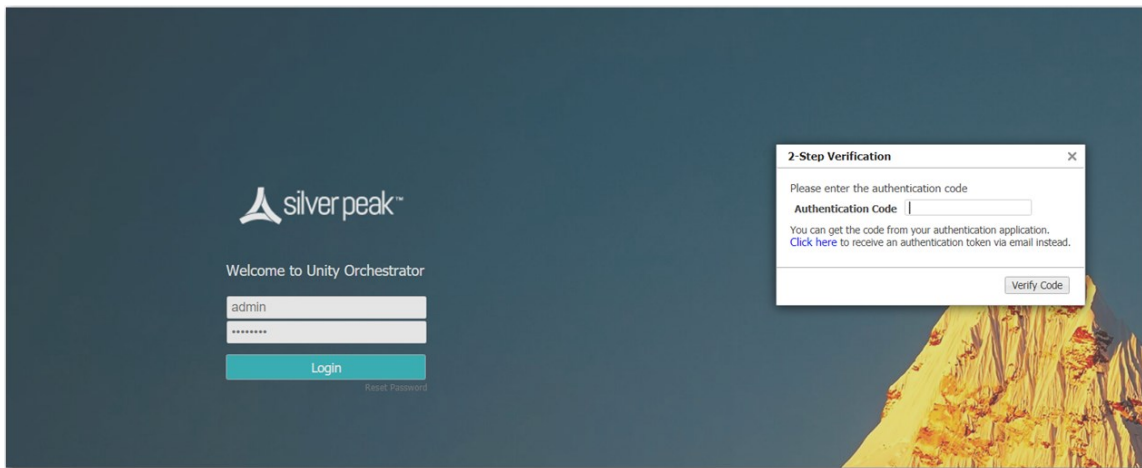
Using Multi-Factor Authentication

After Multi-Factor Authentication is configured, every login requires two steps—entering the username/password and entering the current token.

Based on the 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 > Orchestrator Server > Users & Authentication > User Management > Edit > Modify User

The 'Modify User' dialog box is shown with the following fields and values:

- User Name: test
- First Name: (empty)
- Last Name: (empty)
- Phone: (empty)
- Email: spondugula@silver-peak.
- Two Factor: Application (selected), Email (unchecked)
- Password: (masked with dots)
- Repeat Password: (masked with dots)
- 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 that requires an authentication code.

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.

API Key

Orchestrator > Orchestrator Server > Users & Authentication > API Keys

Use this page to allow your applications to utilize REST APIs without session authentication and management. You can specify permissions, status, name, and IP allow list for your API keys.

An API key can be passed either in the HTTP request header field "X-Auth-Token" or as a query parameter "apiKey".

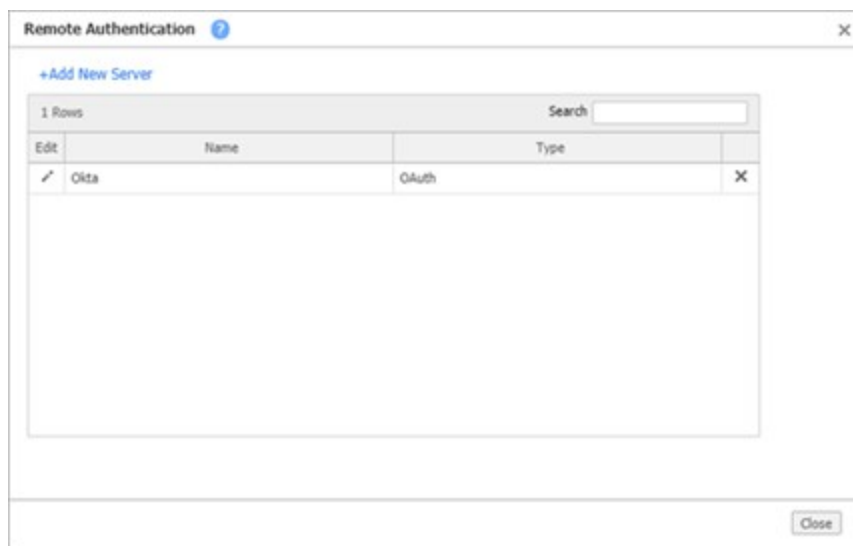
NOTE It is recommended to use different keys for different applications and users.

Click the edit icon to add and define a new API key by entering the fields below.

Field	Description
Key Name	Name of the key you are creating.
Key	Text you cut and paste and insert into your client code.
Permission	Read-Only or Read-Write.
Description	Enter details in this field that describe the purpose of the key you are configuring.
Expiration	Set the expiration date if you want a certain application or script to access the key for a fixed amount of time.
Active	Select Yes or No to display if the key is active or inactive.
IP Allow List	Filters traffic to your private resources through this specified IP range. Traffic is able to pass through with the IP addresses defined in this field.

Remote Authentication

Orchestrator > Orchestrator Server > Users & Authentication > Authentication



Use the **Remote Authentication** dialog box to manage different remote authentication methods for Orchestrator users.

- To add a new remote authentication method, click **+Add New Server**.
- To view or modify the settings for an existing remote authentication method, click the edit icon in the row of the existing method.

Orchestrator supports the following for remote authentication:

- RADIUS
- TACACS+
- OAuth
- JWT
- SAML

Configure a RADIUS or TACACS+ Server

You will need to configure the following when adding or modifying a RADIUS or TACACS+ server:

Field	Description
Read-Write Privilege	RADIUS only: 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	RADIUS only: 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	Select the authentication type that matches what is configured on the RADIUS or TACACS+ server.
Default Role	If RBAC is enabled, you must specify a default role.
Primary/Secondary Server	For each server in use, enter the IP address or hostname, port, and secret key of the RADIUS or TACACS+ server.

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 Remote first .
Primary/Secondary Server	IP address or hostname of the RADIUS or TACACS+ server.
Secret Key	String defined as the shared secret on the server.
Read-Write Privilege	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	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 is configured on the TACACS+ server.

Configure an OAuth Server

Orchestrator supports remote authentication via the OAuth 2.0 framework. Before configuring an OAuth server in Orchestrator, you will need to register Orchestrator as an application with your OAuth provider.

Prerequisites

- The OAuth server must support OAuth 2.0 authorization codes, ID tokens, and optionally refresh tokens.
- The ID token is used to get username, RBAC roles, and RBAC appliance access groups.
- The refresh token can be checked periodically to ensure the user is still authorized/valid.
- Depending on the OAuth server configuration, refresh tokens can be permanent or they can expire. If a token is revoked or expires, the user will be forced to authenticate again.

Register Orchestrator as an App

Before adding an OAuth server in Orchestrator, register a new app on your OAuth server for Orchestrator. You will need to provide the following details when registering the app:

Application Type	Register Orchestrator as a web app.
Allowed Grant Types	Authorization code (required) Refresh token (optional)
Redirect URL	Orchestrator endpoint to which the user will be redirected after successful authentication, which should be <code>https://<Orchestrator_domain_or_IP_address>/gms/rest/authentication/oauth/redirect</code>

Configure OAuth Server Properties in Orchestrator

When adding a new OAuth server or modifying an existing server, you will need to configure the following fields in the Remote Authentication Server dialog box:

Field	Description
Name	Name to identify the server. This name will be displayed on a button on the Orchestrator login page as an alternative method of authentication.
Client ID	Client ID for the Orchestrator application that you created in your OAuth provider.
Client Secret	Client secret for the Orchestrator application that you created in your OAuth provider.
Scopes	OAuth 2.0 uses scope values, as defined in RFC6749, to specify which access privileges are being requested for in Access Tokens. The default scopes for Orchestrator are openid , offline_access , and email .
Authentication URL	This is the Issuer Identifier URL with the authentication request path appended. For example: <code>https://<your-oauth-domain>/oauth2/v1/authorize</code> .

Field	Description
Token URL	This is the Issuer Identifier URL with the token path appended. For example: <code>https://<your-oauth-domain>/oauth2/v1/token</code> .
Username key	This is the OAuth attribute to be sent as the username. Use email if username is an email address. If any other key is used, ensure that it is mapped to the correct scope in the OAuth server.
Roles key (optional) ¹	<p>This field can be left with the default value, <i>sp-roles</i>, or you can enter a new key name, but the key name must match what is configured in your OAuth provider.</p> <p>This is a user claim sent in the ID token that maps to Orchestrator roles defined in Role Based Access Control (RBAC). For example, the OAuth server attribute <i>userType</i> maps to <i>sp-roles</i>, and the OAuth user in Orchestrator has <i>userType</i> = <i>OverlayAdmin</i>.</p>
Appliance Access Group key (optional) ¹	<p>This field can be left with the default value, <i>sp-aag</i>, or you can enter a new key name, but the key name must match what is configured in your OAuth provider.</p> <p>This is a user claim sent in the ID token that maps to Orchestrator Appliance Access Groups defined in Role Based Access Control (RBAC). For example, the OAuth server attribute <i>department</i> maps to <i>sp-aag</i>, and the OAuth user in Orchestrator has <i>department</i> = <i>Asia-Admin</i>.</p>
Default role	If RBAC is enabled, you must specify a default role.

Type	OAuth
Name	
Client ID	
Client Secret	
Scopes	openid,offline_access
Authentication Url	https://authserver.com/oauth2/serve123/v1/authorize
Token Url	https://authserver.com/oauth2/serve123/v1/token
Username key	sp-name
Roles key	sp-roles (optional)
Appliance Access Group key	sp-aag (optional)
Default role	Select role (optional)

Apply Cancel

Configure a JWT Server

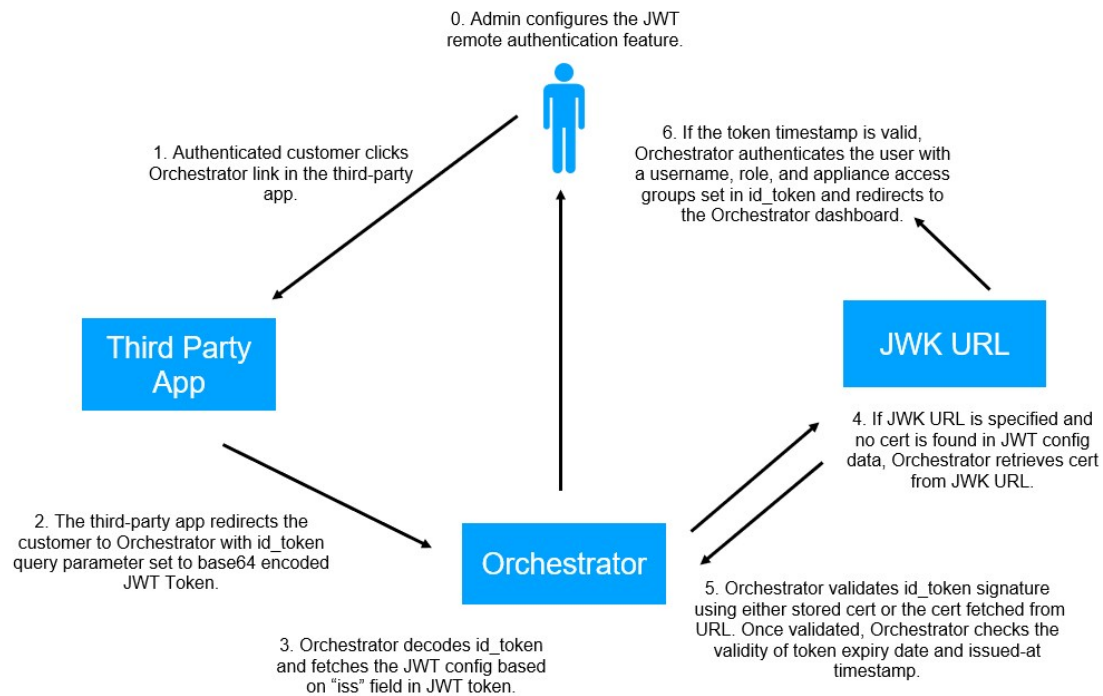
To begin JWT server configuration, the assigned admin needs to specify the following JWT configuration parameters. This includes the following:

- Issuer 'iss'
- Auditor 'aud'
- expiration 'exp'
- signature
- user, role, and AAG

NOTE See the following descriptions in the table below.

- Redirect URL based on successful authentication: `https://<orchestrator_domainName>?access_token=<token>&id_token=<token>&state=<state>&token_type=Bearer&expires_in=3596`

Review the following diagram for more details about the workflow of JWT authentication.



Then, complete the following steps in Orchestrator:

1. Navigate to the **Authentication** tab in Orchestrator.
2. Click **+Add New Server**. The **Remote Authentication Server** window opens.
3. Select **JWT** from the **Type** drop-down menu and complete the following fields.

Field	Description
Name	Name of your JWT provider.
Cert/Signing Key	HMAC or RSA public key used to verify the id_token.
JWK URL	URL that hosts the public certification.
Validation Window	Maximum amount of time in minutes that the expiration is found for the id_token, before a new id_token is created.
Issuer	Issuer claim found in the id_token.
Auditor	Auditor claim found in the id_token.
Username Key	This attribute is sent as the username. Use email if username is an email address. If any other key is used, ensure that it is mapped to the correct scope in the OAuth server.

Field	Description
Roles Key ¹	<p>This field can be left with the default value, <i>sp-roles</i>, or you can enter a new key name, but the key name must match what is configured in your JWT provider.</p> <p>This is a user claim sent in the ID token that maps to Orchestrator roles defined in Role Based Access Control (RBAC). For example, the OAuth server attribute <i>userType</i> maps to <i>sp-roles</i>, and the OAuth user in Orchestrator has <i>userType</i> = <i>OverlayAdmin</i>.</p>
Appliance Access Group Key ¹	<p>This field can be left with the default value, <i>sp-aag</i>, or you can enter a new key name, but the key name must match what is configured in your JWT provider.</p> <p>This is a user claim sent in the ID token that maps to Orchestrator Appliance Access Groups defined in Role Based Access Control (RBAC). For example, the JWT server attribute <i>department</i> maps to <i>sp-aag</i>, and the JWT user in Orchestrator has <i>department</i> = <i>Asia-Admin</i>.</p>
Default role	If RBAC is enabled, you must specify a default role.
JWT token consuming URL	URL of Orchestrator that remains the same.

Configure a SAML Server

Orchestrator supports SAML 2.0 integration, providing authentication and authorization of your credentials through an IdP (Identity Provider), SP (Service Provider), and a Principal. Refer to the list below for the represented meanings:

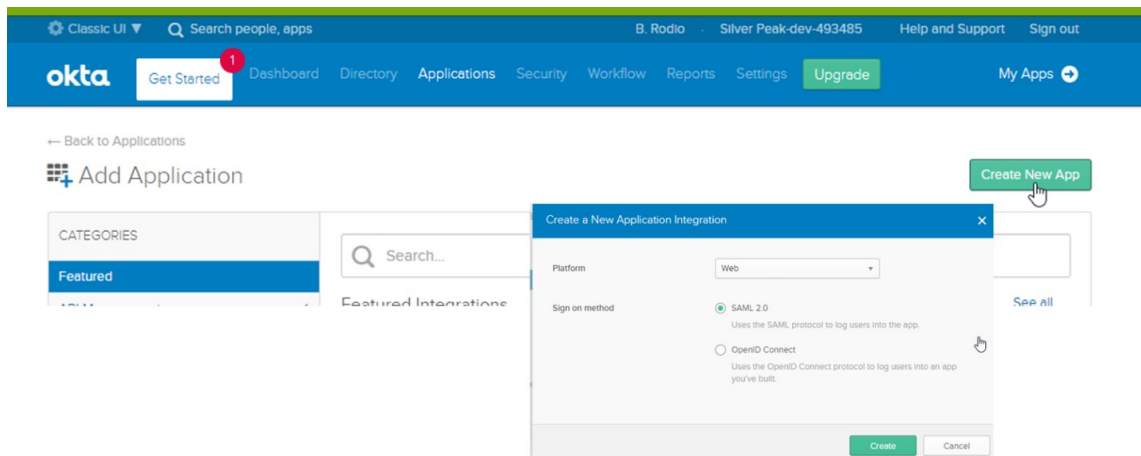
- IdP: Okta
- SP: Orchestrator
- Principal: Principal end user

SAML and Orchestrator Configuration

Complete the following instructions to complete SAML and Orchestrator integration.

TIP It is recommended to have Orchestrator open next to your Okta window while completing these instructions.

1. Sign in to your Okta account.
2. Select **Add Application** and select **SAML 2.0**.
3. Click **Create New App**.



4. Sign in to Orchestrator and navigate to the **Authentication** tab (**Orchestrator > Users & Authentication > Authentication**).
5. Click **+Add New Server**.
6. Select **SAML** from the **Type** field.
7. In Orchestrator, copy the **ACS URL** and the **SP SLO Endpoint** by clicking the icon next to the fields.
8. Navigate back to your SAML application configuration window.
9. Enter the copied URLs in the following fields in the **Step 2: Configure SAML** section:
 - a. Paste the **ACS URL** in the **Single Sign On URL** and **Audience URL** (SP Entity ID) fields.
10. Specify the attributes and their corresponding values on the SAML Settings page. These are configured and assigned on the **RBAC** tab in Orchestrator.
 - a. sp-name: user.email
 - b. sp-role: user.usertype
 - c. sp-aag: user.department
11. Click **Next**.
12. Click **Finish**.
13. Click the **View Setup Instructions** box on the completed **SAML Application Settings** page and enter the following URLs in the corresponding Orchestrator fields:

SAML Field	Orchestrator Field
Identity Provider Single Sign-On URL	SSO Endpoint
Identity Provider Issuer	Issuer URL
X.509 Certificate	IdP X.509 Cert


The following table provides more details about the fields in Orchestrator.

Field	Description
Name	Any text value for your SAML account for identification purposes.
Username Attribute	This attribute is used to retrieve the username from the SAML XML response.
Issuer URL	Unique identifier of the issuer (for example: Okta, OneLogin).
SSO Endpoint	Unique endpoint for the SAML application created on the IdP server.
IdPX.509 cert	Certificate issued by IdP to verify and validate the response received from the IdP (Okta) server.
ACS URL	Orchestrator endpoint needed for configuration on the IdP server. This is provided as a redirect URL after you are authenticated on the IdP server.
SP SLO Endpoint (Optional)	This endpoint is used by IdP to initiate the logout request from Orchestrator to the IdP server.
IdP SLO Endpoint (Optional)	This endpoint is used by IdP to initiate the logout request from Orchestrator to the IdP server. The endpoint used by Orchestrator to initiate the logout request to IdP.
SP X.509 Cert SLO (Optional)	Certificate used by IdP to verify the Single Logout request initiated by Orchestrator to logout the IdP.
Roles Attribute (optional) ¹	This field can be left with the default value, <i>sp-roles</i> , or you can enter a new key name, but the key name must match what is configured in your SAML provider. This is a claim sent to Orchestrator that maps to roles defined in Role Based Access Control (RBAC).
Appliance Access Group key (optional) ¹	This field can be left with the default value, <i>sp-aag</i> , or you can enter a new key name, but the key name must match what is configured in your OAuth provider. This is a claim sent to Orchestrator that maps to Orchestrator Appliance Access Groups defined in Role Based Access Control (RBAC).
Default role	If RBAC is enabled, you must specify a default role.

Cloud Portal

Configuration > Overlays & Security > 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 appliances.
- The cloud portal populates the **Contact** field from information included in your purchase order.
- Use of these services requires that your appliances can access the cloud portal via the Internet.

Audit Logs

Orchestrator > Orchestrator Server > Tools > Audit Logs

The **Audit Logs** tab lists 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

All	Completed	In Progress	Queued	Log Level	Info	Auto Refresh	Pause	Record Count	500	(Max 10000)	Appliance	Type to select	From	To	User	Type to select	Export
-----	-----------	-------------	--------	-----------	------	--------------	-------	--------------	-----	-------------	-----------	----------------	------	----	------	----------------	--------

Audit Logs										
500 Rows										
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	Description
User Name	You can filter/search for an audit log by the user name of the appliance.
IP Address	IP address of the selected appliance.
Host Name	Host name of the appliance that the audit log is coming from.
Action	What you want the audit log to do.
Task Status	Status of the audit log task.
Results	Results of the audit log being searched.
Start Time	Time when the search of the audit log started.
End Time	Time when the search of the audit log ended.
Queued Time	Time when the process/task was requested or scheduled in the queue.
% Completed	Percent completed of the audit log task.
Completion Status	Whether the task has been completed.

Orchestration Settings

Orchestrator > Orchestrator Server > Tools > Orchestration Settings

Orchestration Settings manage Business Intent Overlays (BIOs) and the properties used to control them. It builds new tunnels and fixes existing ones.

Field	Description
Orchestrate Appliances by Applying and Updating Overlays	When selected, updates all associated appliances when overlay changes are saved. NOTE Tunnels are rebuilt only if this field is enabled.
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 Changes	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	Amount of time Orchestrator sleeps or is idle between checking for any configuration changes. For normal size networks, the recommended idle time is 60 seconds. For smaller networks, 30 seconds is the recommended idle time.
Auto Flow Re-Classify	Specifies how the Overlay Manager waits before surveying the network when configuration changes are not being made.

IPSec UDP Settings Section

Field	Description
Default Port	By default, Business Intent Overlays create IPSec UDP tunnels. Default Port is 10002 . If necessary, you can configure this for an individual appliance on its System Information page, under System Settings . This is accessible from the appliance's context-sensitive menu in Orchestrator's navigation pane.
Increment Port By	Referenced when configuring an Edge HA (High Availability) pair. When the value is 1000, the second appliance's default port would become 11002.

Tunnel Settings Tab

Orchestrator > Orchestrator Server > Tools > Tunnels Settings

Use this tab to manage the properties for those tunnels created by Orchestrator. It provides tunnel settings for General, IKE, and IPSec for MPLS, Internet, and LTE WAN Interface labels.

Tunnel Settings for Overlays

Field	Description
General	
Mode	Indicates whether the tunnel protocol is ipsec , ipsec_udp , udp , or gre . If you select IPSec, you can specify the IKE version on the IKE tab.
Auto Max BW Enabled	Allows the appliances to auto-negotiate the maximum tunnel bandwidth.
Auto Discover MTU Enabled	Allows the appliances to auto-negotiate the maximum tunnel bandwidth.
MTU	Maximum Transmission Unit (MTU) 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.
Packet	
Reorder Wait	Maximum time the appliance holds an out-of-order packet when attempting to reorder. The packets can come from either the same or a 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.
FEC	Forward Error Correction (FEC) can be set to enable , disable , and auto .
FEC Ratio	When FEC is set to auto , this specifies the maximum ratio. The options are 1:2, 1:5, 1:10, or 1:20.
Tunnel Health	
Retry Count	Number of failed keep-alive messages that are allowed before the appliance brings the tunnel down.
DSCP	Determines the DSCP marking that the keep-alive messages should use.
FastFail Thresholds	

Field	Description
FastFail Thresholds	<p>Fastfail thresholds determine how quickly to disqualify a tunnel from carrying data when multiple tunnels are carrying data between two appliances.</p> <p>The Fastfail connectivity detection algorithm for the wait time from receipt of last packet before declaring a brownout is:</p> $T_{wait} = Base + N * RTT_{avg}$ <p>where Base is a value in milliseconds and N 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 brownout if it does not see a reply packet from the remote end within 300mS of receiving the most recent packet.</p> <p>In the Tunnel Advanced Options, Base is expressed as Fastfail Wait-time Base Offset (ms), and N is expressed as Fastfail RTT Multiplication Factor.</p> <ul style="list-style-type: none"> ■ Fastfail Enabled – This option is triggered when a tunnel's keep-alive signal does not receive a reply. The options are disable, enable, and continuous. If the disqualified tunnel subsequently receives a keep-alive reply, its recovery is instantaneous. <ul style="list-style-type: none"> • 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 one per second to ten per second. Failover occurs after one second. • When set to continuous, keep-alives are continuously sent at ten per second. Therefore, failover occurs after one tenth of a second. ■ Thresholds for Latency, Loss, or Jitter are checked once every second. <ul style="list-style-type: none"> • Receiving three 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 three 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	Amount of latency measure in MS.

Field	Description
Loss	Amount of data lost measured in percent.
Jitter	Amount of jitter measured in MS.
Fastfail Wait-Time Base Offset	Base time used when calculating the fastfail timeout.
Fastfail RTT Multiplication Factor	Multiplier in the formula used to calculate the fastfail timeout.

Field	Description
IKE	
Authentication Algorithm	This is for setting tunnel authentication. Choose from SHA-1 , SHA2-256 , SHA2-384 , or SHA2-512 .
Encryption Algorithm	Specifies the encryption algorithm used for the Phase 1 negotiation. Choose from AES-256 , AES-128 , or auto .
Diffie-Hellman Group	Diffie-Hellman group used for IKE SA negotiation.
Lifetime	Lifetime of IKE SA.
Dead Peer Detection	<p>Delay time: The amount of time, in seconds, to wait for traffic from the destination IKE peer.</p> <p>Retry Count: The number of times to retry the connection before determining that the connection is dead.</p> <p>NOTE Dead Peer Detection is supported only on EdgeConnect appliances running VXOA software version 8.2.1 and higher.</p>
Phase 1 Mode	Defines the exchange mode for Phase 1. The options are Main or Aggressive . If IKEv2 is selected, the default mode is aggressive.
IKE Version	The IKE major version. Select either IKEv1 or IKEv2 .

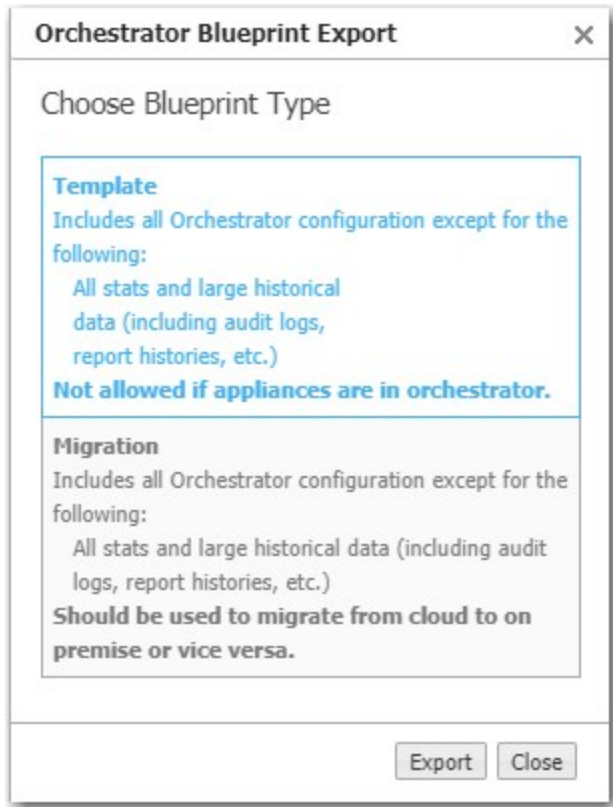
Field	Description
IPSec	

Field	Description
Authentication Algorithm	Authentication algorithm used by IPsec SA. Choose from SHA-1 , SHA2-256 , SHA2-384 , or SHA2-512 .
Encryption Algorithm	Specifies the encryption algorithm used for the Phase 1 negotiation. Choose from AES-256 , AES-128 , or auto .
Enable IPsec Anti-replay Window	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.
Lifetime	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

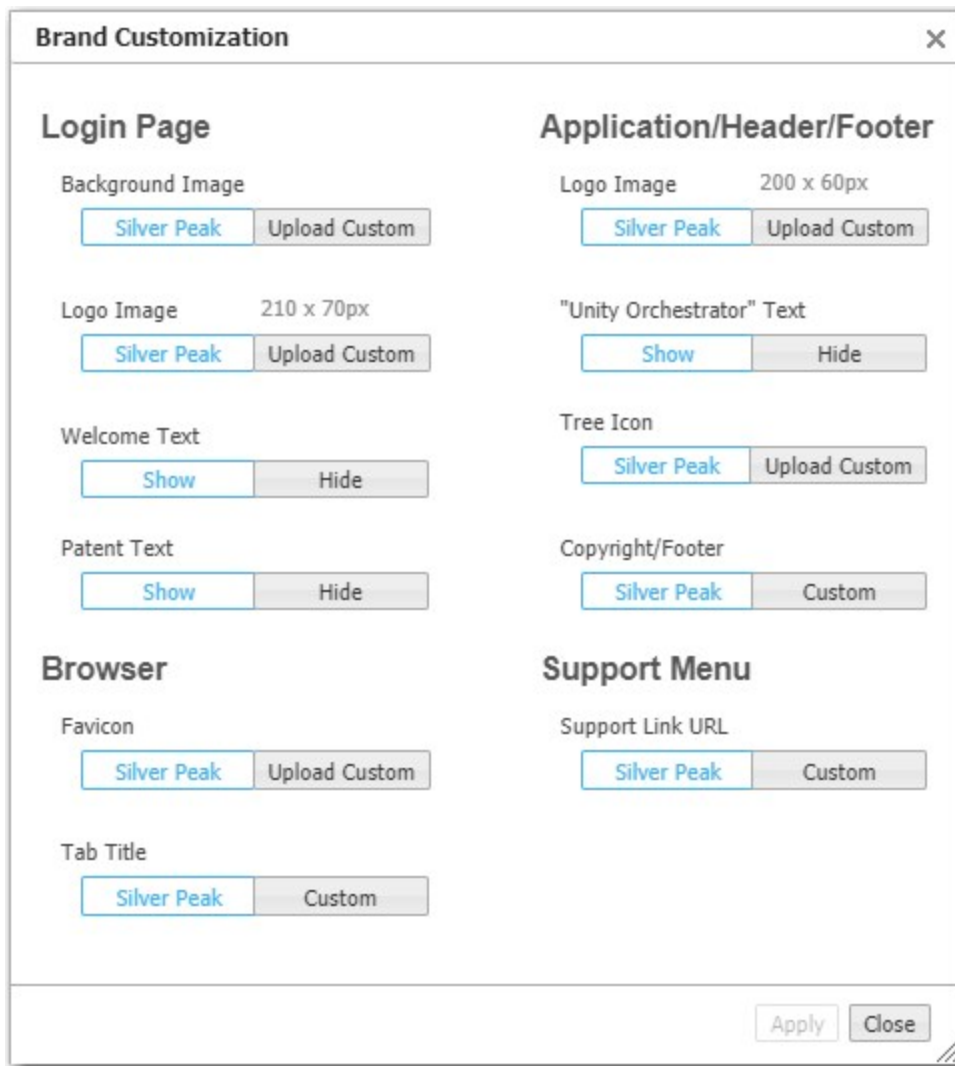
Use this dialog box to create and export a configuration that Orchestrator-SP can use as a template for other Orchestrators.



Brand Customization

Orchestrator > Orchestrator Server > Tools > Brand Customization

Use this dialog box to customize the branding elements of the Orchestrator user interface.



The **Brand Customization** dialog box is divided into four main sections: **Login Page**, **Application/Header/Footer**, **Browser**, and **Support Menu**. Each section contains options to either select the **Silver Peak** default branding or upload a custom one.

- Login Page**
 - Background Image**: Buttons for **Silver Peak** and **Upload Custom**.
 - Logo Image** (210 x 70px): Buttons for **Silver Peak** and **Upload Custom**.
 - Welcome Text**: Buttons for **Show** and **Hide**.
 - Patent Text**: Buttons for **Show** and **Hide**.
- Application/Header/Footer**
 - Logo Image** (200 x 60px): Buttons for **Silver Peak** and **Upload Custom**.
 - "Unity Orchestrator" Text**: Buttons for **Show** and **Hide**.
 - Tree Icon**: Buttons for **Silver Peak** and **Upload Custom**.
 - Copyright/Footer**: Buttons for **Silver Peak** and **Custom**.
- Browser**
 - Favicon**: Buttons for **Silver Peak** and **Upload Custom**.
 - Tab Title**: Buttons for **Silver Peak** and **Custom**.
- Support Menu**
 - Support Link URL**: Buttons for **Silver Peak** and **Custom**.

At the bottom right, there are **Apply** and **Close** buttons.

Maintenance Mode

Orchestrator > Orchestrator Server > Tools > Maintenance Mode

You can put one or more appliances in maintenance mode by selecting the specific appliance in the tree. When approved, the appliances are added to the maintenance list. You also can put an appliance in maintenance mode by searching for **"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 **Orchestrator > Orchestrator Server > Tools > Maintenance Mode**.
2. Click **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. Click **OK**.
6. Click **Save**.

NOTE The appliance goes into maintenance mode if you pause orchestration and/or suppress all alarms.

Field	Description
Host Name	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	IP address of the appliance in maintenance mode.
Version	Current version of the appliance.

Upgrade Orchestrator Software

Orchestrator > Software & Setup > Upgrade > Upgrade Orchestrator

If you are already using Orchestrator 8.6.0 or later and want to upgrade to a newer version, complete the following procedure.

WARNING An upgrade that fails can put Orchestrator into a corrupt state. Be sure to back up Orchestrator before you start the upgrade process.

1. Open an SSH session to the Orchestrator.
2. Log in as **admin** or a user with administrative privileges.
3. Switch to root:

```
su - root
```
4. Enter the root password when prompted. Contact Silver Peak TAC if you do not know your root password.

5. Change to the /home directory:

```
cd /home
```

Depending on your environment, you can upgrade Orchestrator in one of two ways:

- Upgrade via HTTP
- Upgrade via SCP

Upgrade via HTTP

If you have an HTTP URL to the Orchestrator installation file, enter the following in the existing SSH console to run the install script and point it to the hosted installation file:

```
/home/gms/gms/setup/install_orchestrator.sh <HTTP URL of the Orchestrator Installation File>
```

NOTE The upgrade process can take several hours to complete.

Upgrade via SCP

If you do not have an HTTP server, copy the installation file to Orchestrator by using SCP, run the install script, and point it to the local installation file:

1. From your local PC console, enter the following:

```
scp <Orchestrator Installation file> admin@<orchestrator_ip_address>:/home/gms
```

2. From the Orchestrator SSH console, enter the following:

```
/home/gms/gms/setup/install_orchestrator.sh /home/gms/<Orchestrator Installation file>
```

NOTE The upgrade process can take several hours to complete.

Check 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	Download
99.99.99.35870	BETA	13-May-18 00:00	Test for orchestrator feature GMS-11402	Download
99.99.99.36894	BETA	13-May-18 00:00	Test for orchestrator feature GMS-11402	Download

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

Back Up on Demand

Orchestrator > Software & Setup > Backup > Backup Now

Use this dialog box 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.

Schedule Orchestrator Backup

Orchestrator > Software & Setup > Backup > Schedule Backup

Use this dialog box to schedule backups of the Orchestrator database and optionally schedule backups of the Orchestrator Stats Collector using the same destination and schedule.

Field	Description
View Currently Scheduled Jobs	Click to open the Scheduled Jobs tab.
Protocol	Protocol to apply: FTP , SCP , HTTP , HTTPS , or SFTP .
Hostname	Host name of the backup server.
Username	Username that the Orchestrator server uses to log in to the backup server.
Password	Password for the username.
Directory	Directory name of the backup server.
Port	Port number of the backup server.
Max backups to retain	Maximum number of backups to retain.
Test	Click Test to verify that Orchestrator can reach the destination.

Field	Description
Schedule	<ol style="list-style-type: none"> 1. Click Add to create a schedule or Edit to modify a schedule. The Schedule dialog box opens. 2. Select Daily, Weekly, Monthly, or Yearly. 3. Complete the remaining fields, and then click OK. <p>TIP To specify the timezone for scheduled jobs and reports, navigate to Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs.</p>
Description	(Optional) Description for the backup schedule.
Stats Collector	<p>Do one of the following:</p> <ul style="list-style-type: none"> ■ Select the Use Orchestrator configuration check box to back up the Orchestrator Stats Collector on the same schedule and to the same destination. ■ Clear the Use Orchestrator configuration check box to specify a different backup destination and set a different schedule for the Orchestrator Stats Collector. <p>CAUTION If you clear the Use Orchestrator configuration check box, and you do not complete the Schedule Stats Collector Backup dialog box, the Stats Collector will not be backed up. For more information, see Schedule Stats Collector Backup.</p>

Schedule Stats Collector Backup

Orchestrator > Software & Setup > Backup > Schedule Stats Collector Backup

Use this dialog box to schedule backups of the Orchestrator Stats Collector.

Schedule Stats Collector Backup

[View Currently Scheduled Jobs](#)

☐ Use Orchestrator backup configuration

Destination

Protocol

SCP

Hostname

10.99.217.23

Username

admin

Password

Directory

/home/gms

Port

22

Max backups to retain

4

Test

Schedule

Schedule

Every day at 20:08 starting 12-Aug-21 23:58 PDT

Edit

Description

sss

Save

Close

Field	Description
View Currently Scheduled Jobs	Click to open the Scheduled Jobs tab.
Use Orchestrator backup configuration	Select this check box to back up the Stats Collector using the same destination and schedule set in the Schedule Orchestrator Backup dialog box. For more information, see Schedule Orchestrator Backup .
Protocol	Protocol to apply: FTP, SCP, HTTP, HTTPS, or SFTP.
Hostname	Host name of the backup server.
Username	Username that the Orchestrator server uses to log in to the backup server.
Password	Password for the username.
Directory	Directory name of the backup server.
Port	Port number of the backup server.
Max backups to retain	Maximum number of backups to retain.
Test	Click Test to verify that Orchestrator can reach the destination.

Field	Description
Schedule	<ol style="list-style-type: none"> 1. Click Add to create a schedule or Edit to modify a schedule. The Schedule dialog box opens. 2. Select Daily, Weekly, Monthly, or Yearly. 3. Complete the remaining fields, and then click OK.
TIP To specify the timezone for scheduled jobs and reports, navigate to Orchestrator > Software & Setup > Setup > Timezone for Scheduled Jobs .	
Description	(Optional) Description for the backup schedule.

SMTP Server Settings

Orchestrator > Software & Setup > Setup > SMTP Server Settings

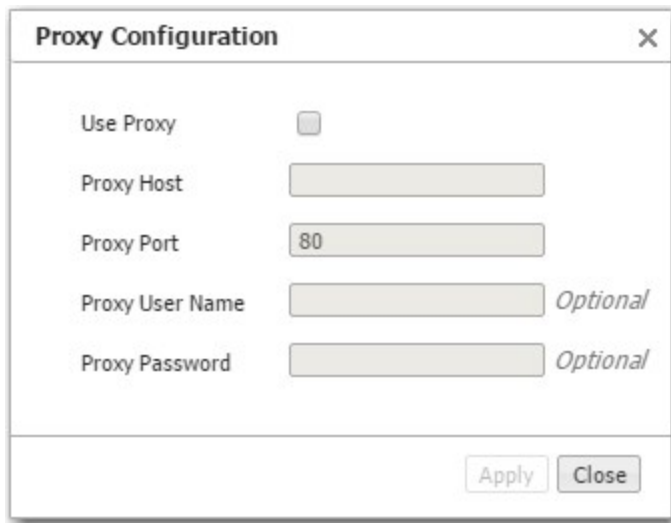
For permanent and private email delivery, change the SMTP (Simple Mail Transfer Protocol) server and settings to your company's SMTP settings.

- If a test email does not arrive within minutes, check your firewall.
- After configuring the SMTP settings, you can specify email recipients for:
 - **alarms** (Monitoring > Alarms > Alarm Email 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.



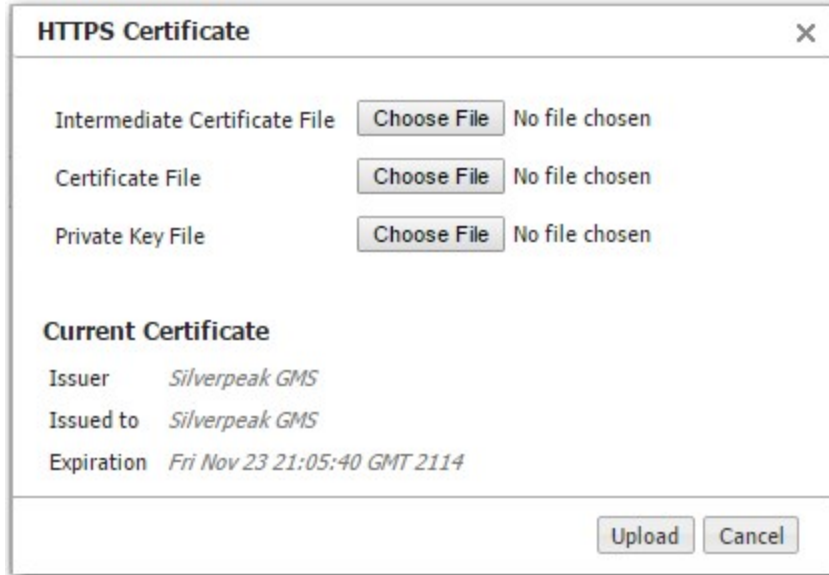
The Proxy Configuration dialog box contains the following fields and controls:

- Use Proxy:** A checkbox that is currently unchecked.
- Proxy Host:** A text input field.
- Proxy Port:** A text input field 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".
- Buttons:** "Apply" and "Close" buttons at the bottom right.

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.



The HTTPS Certificate dialog box contains the following sections and controls:

- Intermediate Certificate File:** A "Choose File" button followed by the text "No file chosen".
- Certificate File:** A "Choose File" button followed by the text "No file chosen".
- Private Key File:** A "Choose File" button followed by the text "No file chosen".
- Current Certificate:** A section displaying certificate details:
 - Issuer:** Silverpeak GMS
 - Issued to:** Silverpeak GMS
 - Expiration:** Fri Nov 23 21:05:40 GMT 2114
- Buttons:** "Upload" and "Cancel" buttons at the bottom right.

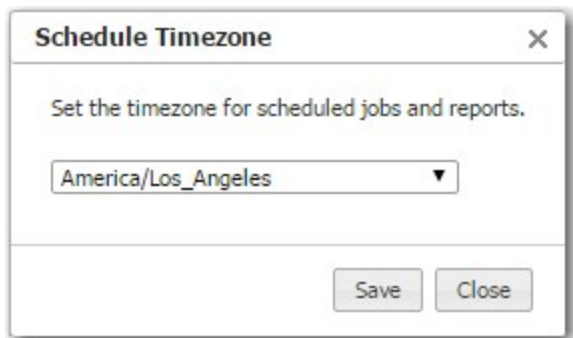
To use a custom certificate with Orchestrator:

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, and so forth.
 - 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, 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 dialog box to set the timezone for scheduled jobs and reports.



Orchestrator Advanced Properties

Orchestrator > Software & Setup > Setup > Advanced Properties

IMPORTANT: Changing the default 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

Property Name ▲	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


Apply/Restart
Restore Defaults
Close

Change the Orchestrator 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	Description
ERROR	An error. This is a non-urgent failure.
WARNING	A warning condition. Indicates an error will occur if action is not taken.
INFORMATIONAL	Informational. Used by Silver Peak for debugging.
DEBUG	Used by Silver Peak for debugging.

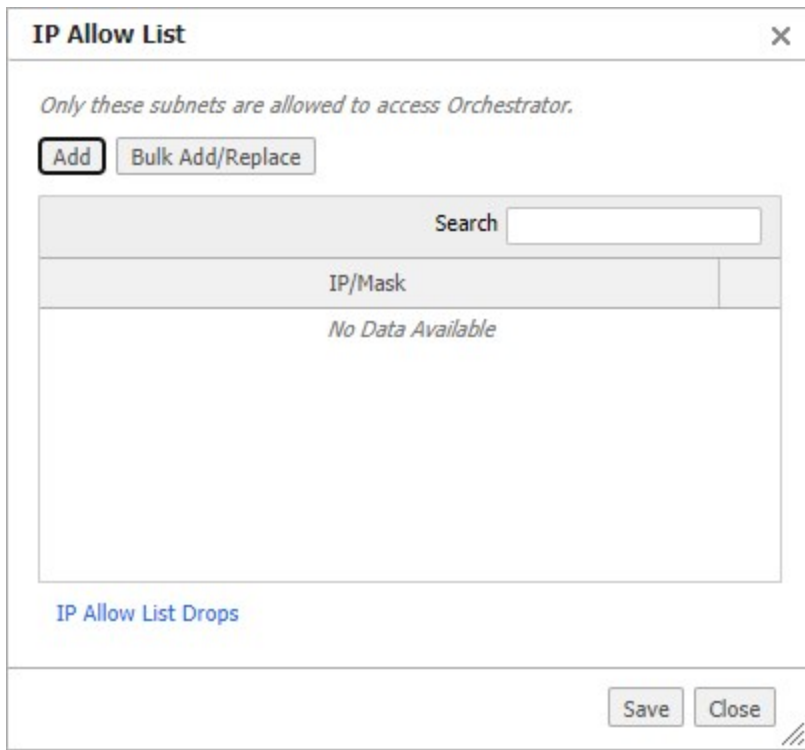
- The bolded part of the name is what displays in Silver Peak's logs.
- If you select **INFO** (the default), the log records any event with a severity of INFO, WARNING, and ERROR.
- These are purely related to event logging levels, **not** alarm severities, even though some naming conventions overlap.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

IP Allow List

Orchestrator > Software & Setup > Setup > IP Allow List

IP Allow List 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), users can get locked out of Orchestrator.



The IP Allow List dialog box is titled "IP Allow List" and includes a close button (X) in the top right corner. Below the title bar, a message states: "Only these subnets are allowed to access Orchestrator." There are two buttons: "Add" (highlighted with a red box) and "Bulk Add/Replace". Below these buttons is a search bar with the label "Search" and an input field. Underneath the search bar is a table with a single header row labeled "IP/Mask". The table body contains the text "No Data Available". At the bottom left of the dialog, there is a link labeled "IP Allow List Drops" in blue. At the bottom right, there are "Save" and "Close" buttons.

IP/Mask
No Data Available

To view a list of traffic that has been dropped because of these restrictions, click **IP Allow List 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 have assigned it, Orchestrator's **Getting Started Wizard** opens.

The wizard guides 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 are 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 is synchronized across Orchestrator and the appliances.
- **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 do not **Apply** the configuration after you complete the last page, the Orchestrator wizard reappears at your next login.

To access the Orchestrator wizard again after initial configuration, navigate to **Orchestrator > Software & Setup > Setup > Configuration Wizard**.

Statistics Retention

This tab displays all the statistics Orchestrator collects from appliances. Orchestrator saves the statistics data in a database with the retention policies defined on this tab. Complete the following steps to begin.

1. Click the edit icon in the table next to the statistic you want Orchestrator to collect.
2. Select the **Collect this statistic in Orchestrator** check box to enable or disable statistics collection.
3. Enter how long you want Orchestrator to retain the statics for Minute Granularity, Hourly Granularity, and Daily Granularity before it collects data and stores in the partition.

TIP If you click **More Options**, you can enter values for the Database Duration.

4. Click **Apply**.

Refer to the table below for more detail.

Field	Description
Statistic	The selected statistic of which you want Orchestrator to collect data.
Enabled	If you have enabled or disabled statistics retention.
Minute Granularity (hours)	Amount of times in one minute Orchestrator stores data.
Hourly Granularity (days)	Amount of times in one hour Orchestrator stores data.
Daily Granularity (months)	Amount of time in one day Orchestrator stores data.
Estimated Disk Space	Estimated amount of disk space the selected statistic uses. At the bottom of the screen, you can get an estimated disk space required for a number of appliances, overlays, and tunnels.

To display the default settings for appliance properties, click **Advanced Properties**. **IMPORTANT:** Changing the default values of these settings is not recommended without consulting Silver Peak.

Stats Collector Configuration

Orchestrator > Software & Setup > Setup > Stats Collector Configuration

Orchestrator collects statistical data from your appliances to monitor performance, network traffic, and appliance status. Before Orchestrator release 9.1.0, the process of collecting, storing, and retrieving this data impacted performance due to the amount of data stored on and requested from the database.

To improve Orchestrator performance, Orchestrator 9.1.0 includes a new Stats Collector feature that eliminates the use of Orchestrator resources for monitoring your appliances. This new architecture allows you to scale your network with greater performance.

The new Stats Collector feature collects statistics from appliances and provides the information to Orchestrator. When enabled, the new Stats Collector runs in parallel with the legacy stats collector to collect the necessary historical statistical data. After collecting that data, you can discontinue the legacy stats collection. You will not experience performance improvement until you discontinue the legacy stats collection.

Prerequisites

- Upgrade all appliances to version 9.1.0 before enabling the new Stats Collector feature.
- Create at least one remote stats collector for every 150 appliances—if you have less than 150 appliances, you can use the predefined local stats collector. Each remote stats collector must meet the following minimum requirements:
 - CPU: 4 GHz
 - RAM: 16 GB

Before You Begin

Before you can configure the new Stats Collector feature in Orchestrator, you must:

1. [Create a Remote Stats Collector](#).
2. [Authenticate the Remote Stats Collector](#).

Create and authenticate as many remote stats collectors as needed.

Create a Remote Stats Collector

To create a remote stats collector, you will use the CLI to run an Orchestrator on a virtual machine (VM) in Stats Collector Mode only.

1. Open an SSH session to the Orchestrator you want to use as a remote stats collector.
2. Log in as **admin** or a user with administrative privileges.
3. Switch to root:
`su - root`
4. When prompted, enter the root password. If you do not know your root password, contact Silver Peak TAC.

5. Change to the gms directory:
`cd gms`
6. Enter `orch-setup`, and then press Enter.
7. Enter `-m`, and then press Enter.
8. Enter the root password, and then press Enter.
9. At the prompt, enter `s`.
10. To proceed, enter `y`.

This VM is now a remote stats collector. Note the DNS name. You will need the DNS name when you configure the remote stats collector in Orchestrator.

Authenticate the Remote Stats Collector

After you create a remote stats collector, you must authenticate by copying the Orchestrator public key and pasting it into the same folder on the new remote stats collector.

1. Open an SSH session to the Orchestrator.
2. Log in as **admin** or a user with administrative privileges.
3. Go to:
`cd /home/gms/sc/publickeys`
4. To list the public key, enter `ls`, and then press Enter.
5. Copy the public key.
6. Open an SSH session to the remote stats collector.
7. Log in as **admin** or a user with administrative privileges.
8. Go to:
`cd /home/gms/sc/publickeys`
9. Paste the public key, and then press Enter.

Configure the New Stats Collector Feature

After the remote stats collectors are created and authenticated, you must configure the new Stats Collector feature in Orchestrator. Complete the following tasks:

1. Back up Orchestrator. For more information on backing up Orchestrator, see “Backing Up on Demand.”

Before you enable the new Stats Collection feature and discontinue the legacy stats collection feature, Silver Peak recommends that you back up the Orchestrator database. Discontinuing the legacy stats collection is permanent. To return to your previous configuration, you must restore the Orchestrator configuration backup.

2. [Add Remote Stats Collectors](#). You need at least one remote stats collector for every 150 appliances. If your network contains less than 150 appliances, you can use the predefined local stats collector.
3. [Associate Appliances with a Remote Stats Collector](#) or [Associate Appliances with the Predefined Local Stats Collector](#)
4. When the necessary historical data has been collected, [Discontinue the Legacy Stats Collection](#).

Add Remote Stats Collectors

You must add at least one remote stats collector for every 150 appliances in your network.

To add a remote stats collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

The Stats Collector Configuration tab opens.

2. Click **Edit Remote Stats Collectors**.

The Edit Stats Collectors dialog box opens.

3. Click **Add Remote Stats Collector**.

The New Stats Collector dialog box opens.

4. Configure the following elements as needed:

Field	Description
Name	Name of the remote stats collector.
DNS Name	Name of the DNS of the remote stats collector.
Port	Port number the remote stats collector is running on.
Protocol	HTTPS.

5. Click **Save**.

Delete a Remote Stats Collector

To delete an existing remote stats collector, click the delete icon (X) in the last column of the entry in the table.

Associate Appliances with a Remote Stats Collector

To associate appliances with a remote stats collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

The Stats Collector Configuration tab opens.

2. In the Orchestrator appliance tree, select one or more appliances to associate with a specific remote stats collector. You can associate up to 150 appliances with each remote stats collector.

IMPORTANT: The statistics for an appliance are tied to the remote stats collector it is associated with. If you associate an appliance with a different remote stats collector, you lose all statistical data associated with that appliance.

3. Select the **Add** check box next to the remote stats collector you want to associate the selected appliance(s) with.
4. Click **Apply**.

The Apply Changes dialog box opens.

5. Click **Apply Changes**.

Associate Appliances with the Predefined Local Stats Collector

If you are installing Orchestrator version 9.1.0 or upgrading to version 9.1.0 or later, Orchestrator provides a default stats collector called "local." You cannot edit or delete the local stats collector. You can associate up to 150 appliances with the local stats collector.

NOTE If you are upgrading to Orchestrator 9.1.0, Silver Peak automatically associates all appliances with the local stats collector.

NOTE If you run Orchestrator in Orchestrator Only mode (`orch-setup -m o`), the local stats collector will be disconnected.

To associate appliances with the local stats collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

The Stats Collector Configuration tab opens. This tab displays the stats collector configuration for all appliances selected in the appliance tree to the left.

2. In the Orchestrator appliance tree, select one or more appliances to associate with the local stats collector.
3. Select the **Add** check box next to the local stats collector.
4. Click **Apply**.

The selected appliances are associated with the local stats collector. The Changes column indicates the stats collectors that were added and removed.

Enable the New Stats Collector

After you associate appliances with either the local stats collector or new remote stats collectors, you must enable the new Stats Collector feature to begin collecting data.

NOTE The legacy status collector continues to collect stats in parallel with the new Stats Collector feature until you discontinue the legacy stats collection. For more information, see [Discontinue the Legacy Stats Collection](#).

IMPORTANT: You cannot disable the new Stats Collector after it has been enabled. Silver Peak recommends that you back up Orchestrator before you enable the new stats collector. For more information on backing up Orchestrator, see "Backing Up on Demand."

To enable the new stats collector:

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

The Stats Collector Configuration tab opens.

2. Click **Enable New Stats Collection**.

The Enable New Stats Collection dialog box opens.

Before you can enable the new Stats Collector feature, you must upgrade all appliances to version 9.1.0. The Enable New Stats Collection dialog box lists appliances that must be upgraded to support the new stats collection.

3. Click **Enable New Stats Collection Now**.

Discontinue the Legacy Stats Collection

IMPORTANT: Do not discontinue the legacy stats collection until you have collected sufficient historical data with the new Stats Collector feature. For example, if you need 30 days of statistical data, enable the new Stats Collector, wait 30 days, and then disable the legacy stats collection.

1. Navigate to **Software & Setup > Setup > Stats Collector Configuration**.

The Stats Collector Configuration tab opens.

2. Click **Discontinue Legacy Stats Collection**.

The Discontinue Legacy Stats Collection dialog box opens.

IMPORTANT: This step permanently disables the legacy stats collection. All legacy stats will be deleted.

3. Click **Discontinue Legacy Stats Collection**.

Notification Banner

You can add a notification in the header of your Orchestrator UI if you are conducting downtime or for maintenance reasons. Complete the following steps to add a notification.

1. Navigate to **Orchestrator > Software & Setup > Setup > Notification Banner** in Orchestrator.

The Notification dialog box opens.

2. Enter the message you want to display in the Orchestrator header.
3. Click **Save**.

ClearPass Policy Manager

Orchestrator > Aruba Central > ClearPass Policy Manager

Orchestrator supports association with ClearPass Policy Manager, which provides role-based and secure network access for devices. This integration provides user and role information for an IP address, which you can view on the Flows and Top Talkers tabs of Orchestrator.

The ClearPass Policy Manager tab displays information about users and devices provisioned to access your network via ClearPass. The searchable information on this tab includes details such as username, IP address, and role.

You can apply the following filters to your ClearPass logs:

- Select the **All**, **Active**, or **Historical** filters to determine which actions you want to display in the table.
- Select **Auto Refresh** or **Pause** to refresh or pause the table. By default, the table refreshes automatically.
- To limit the filtering criteria, enter a value in the **Record Count** field. The default value is 500, and the maximum number you can filter is 10,000.
- To filter by date and time, enter values in the **From** and **To** fields.
- To search for a specific username, enter a value in the **User** field. You can search a wild card character (*) as a username using the following schema:
 - x* = anything that starts with the entered value
 - *x = anything that ends with the entered value
- To search for a specific IP address, enter a value in the **IP** field.

To export a .csv file of your table, click **Export**.

Field	Definition
Start Time	Time when the device began its network session.
End Time	Time when the device ended its network session.
CPPM	ClearPass Policy Manager server used to authenticate.
IP Address	IP address authenticated to the network.
Username	Username authenticated to the network.
Role	Role assigned to the user that authenticated to the network.
Device Type	Device type used to connect to the network.
MAC Address	MAC address of the system connecting to the network.

Field	Definition
Posture	Security health posture of the connected device.
Location ID	Location ID of the user connecting to the network.
Protocol	Type of authentication server used to connect to the network.
Details	All user information sent from CPPM but not required by Orchestrator. Values are in JSON format.

Manage ClearPass Policy Manager Accounts

Click **Accounts** on the ClearPass Policy Manager tab to view and manage ClearPass accounts that are associated with Orchestrator.

NOTE Before you begin the ClearPass Policy Manager (CPPM) configuration in Orchestrator, you must have a ClearPass account to authenticate and authorize Orchestrator. If you do not have these credentials, contact your system administrator.

View ClearPass Policy Manager Accounts

The ClearPass Policy Manager Accounts dialog box displays the following information about ClearPass accounts that are already associated with Orchestrator:

Field	Definition
Edit	Click the icon to edit your CPPM instance.
Name	Name of your CPPM instance.
Domain/IP	Domain or URL of your CPPM instance.
Connectivity	Status of the connection between Orchestrator and your CPPM instance. The status may appear as Connected, Connecting, Auth Failed, and Unreachable.
Service Status	Status of your CPPM instance. A status other than Connected could indicate a problem with your CPPM configuration. To troubleshoot, click the Info icon, and then reset any service that is not currently connected.
Pause	To pause the connection for your CPPM instance, click this toggle.

Add a ClearPass Policy Manager Server

Follow the steps below to add a new ClearPass Policy Manager account.

1. If not already opened, click **Accounts** to open the ClearPass Policy Manager Accounts dialog box.
2. Click **+Add New Server**.

The ClearPass Policy Manager Server Configuration dialog box opens.

3. Enter the following information:

Field	Definition
Name	Name of your CPPM instance.
Domain/IP	Domain or URL of your CPPM instance.
Client ID	Client ID generated from your CPPM account.
Secret Key	Secret key generated from your CPPM account.
Verify server certificate	If you are using cloud instances of both CPPM and Orchestrator, or if you are using an on-premise instance of CPPM with a valid certificate, select this check box. If you are using an on-premise instance of Orchestrator or an on-premise instance of CPPM without a valid certificate, clear this check box.

4. Click **Save**.

Your CPPM instance now appears in the ClearPass Policy Manager Accounts dialog box. The Connectivity and Service Status fields should both appear as Connected.

Edit a ClearPass Policy Manager Server

1. If not already opened, click **Accounts** to open the ClearPass Policy Manager Accounts dialog box.
2. Click the **Edit** icon next to the instance you want to edit.

The ClearPass Policy Manager Server Configuration dialog box opens.

3. Edit the information in the dialog box, and then click **Save**.

Pause ClearPass Policy Manager Integration

To pause the integration between CPPM and Orchestrator, click **Pause Orchestration** from the ClearPass Policy Manager tab.

NOTE Clicking **Pause Orchestration** pauses the connection between all instances of CPPM configured in Orchestrator. To pause an individual instance, click **Accounts**, and then click the toggle under Pause for the instance you want to pause.

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 to troubleshoot network issues.

Tech Support - Appliances

Support > Technical Assistance > Tech Support - Appliances

Use this tab to create a new case, generate a system dump, upload files to an existing case, or download selected files to Orchestrator.

By default, the table displays all files available on the selected appliances. Click the appropriate button to filter files by type (Logs, Sys Dump, Snapshot, TCP Dump). The table includes the following details for each file:

Field	Description
Appliance Name	Name of the appliance on which the file is available.
File type	Specific file type (log, sys dump, snapshot, or TCP dump).
File Name	Name of the file.
Last Modified	Date when the file was last modified.
File Size	Size of the file.

Download to Orchestrator

Complete the following steps if you want to download one or more files to Orchestrator.

1. Select one or more files in the table (use Ctrl or Shift to select multiple files).
2. Click the **Download to Orchestrator** button above the table.
3. When prompted, click **Download** to confirm or click **Close** to cancel.

The Monitor Transfer Progress window appears, showing the status of current and previous downloads.

4. To stop any downloads that are not yet finished, click **Cancel**.

NOTE To access any files that have been downloaded, open the **Tech Support - Orchestrator** tab under the Support menu. After selecting one or more files, you can create a new case, upload files to an existing case, or download files to your local machine.

Tech Support - Orchestrator

Support > Technical Assistance > Tech Support - Orchestrator

This tab displays a list of Orchestrator log files and system dump files, as well as support files that have been downloaded from appliances. You can use these files to create or update support cases, or you can download files to your local machine from Orchestrator.

By default, the table displays all files available on Orchestrator. Click the appropriate button to filter files by type (logs, system dumps, or appliance files). The table includes the following details for each file:

Field	Description
Source	Source of the selected file (Orchestrator or a specific appliance).
File Type	Specific file type (log, sys dump, snapshot, or TCP dump).
File Name	Name of the file.
Last Modified	Date when the file was last modified.
File Size	Size of the file.

Take Action with Files

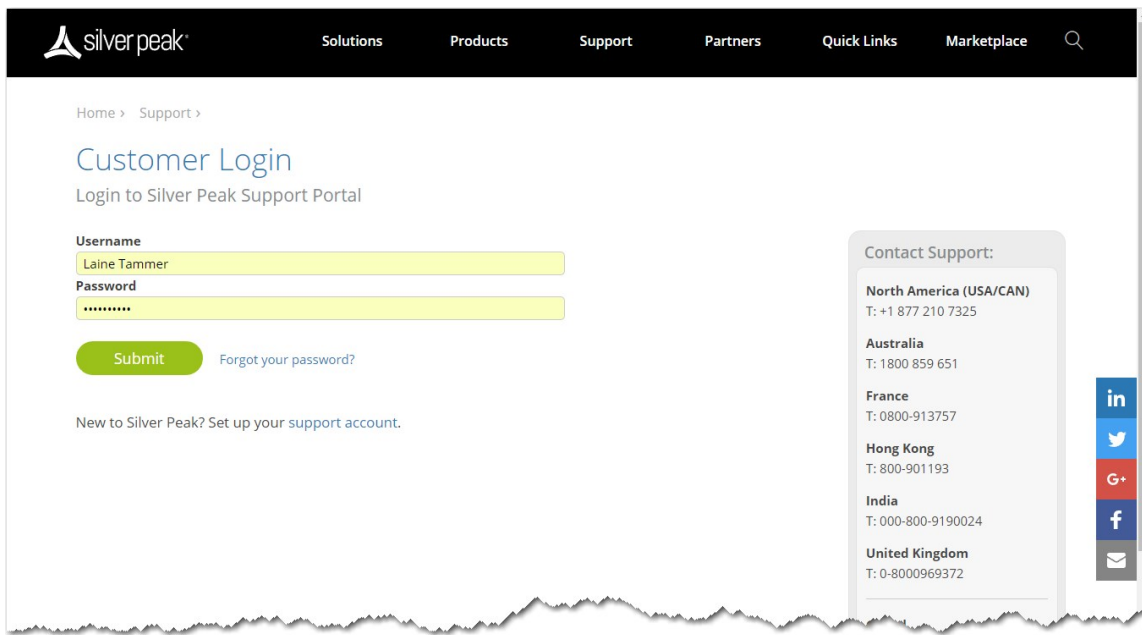
With one or more files selected, you can create a new support case, add files to an existing case, or download files to your local machine.

- Click **Create Case** to open a new support case. Fill in a few additional details and the selected files will be attached to a new support case.
- Click **Upload Selected Files** to attach files to an existing support case. You will need to know the case number when using this option.
- Click **Download selected Files** to download files to your local machine. Confirm the download and select a location where you want to save the files.

Log In to 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 > Technical Assistance > Support Portal Log-in**. The following page opens in a separate browser tab.



You also can access this page directly by going to Silver Peak’s web page and selecting **Support > Customer Login** from the menu bar.

Monitor Transfer Progress

Support > Technical Assistance > Monitor Transfer Progress

This table displays the current status of any files being uploaded to Support.

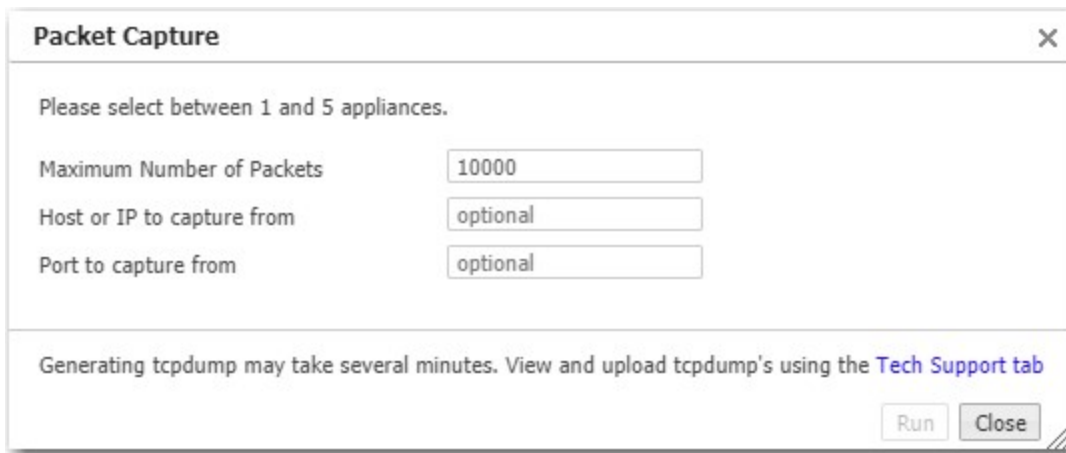


The **Monitor Uploads** dialog box features a table with the following headers: Source, Files, Start Time (with a dropdown arrow), End Time, Uploaded, Status, and Cancel. The table body is currently empty. A **Close** button is located at the bottom right of the dialog.

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 appliance tree.



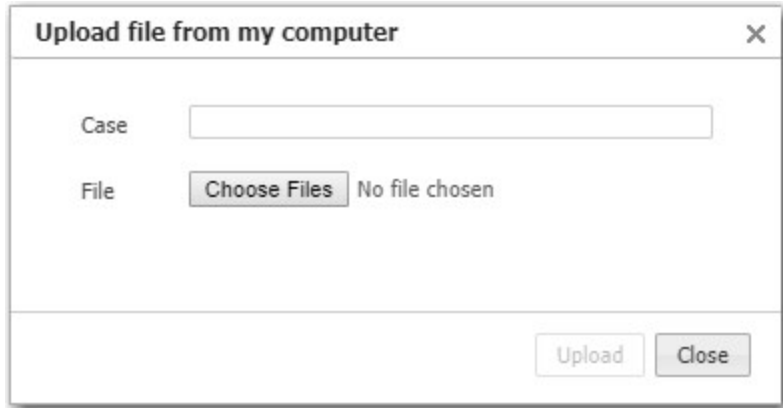
The **Packet Capture** dialog box contains the following fields and controls:

- Instruction: Please select between 1 and 5 appliances.
- Maximum Number of Packets:
- Host or IP to capture from:
- Port to capture from:
- Footer text: Generating tcpdump may take several minutes. View and upload tcpdump's using the [Tech Support tab](#)
- Buttons: **Run** and **Close**

Upload Local Files

Support > Technical Assistance > Upload Local Files

Use this dialog box to upload files related to your Support case from your computer.



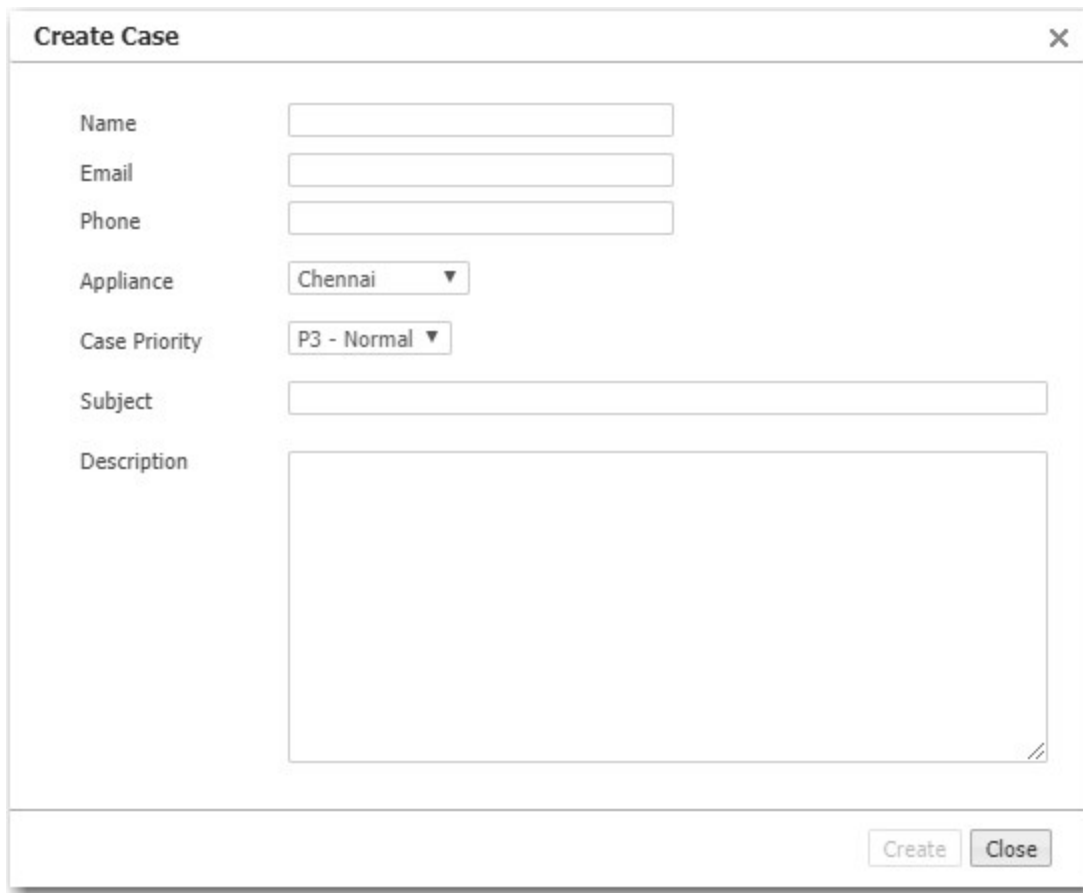
The screenshot shows a dialog box titled "Upload file from my computer" with a close button (X) in the top right corner. Inside the dialog, there is a "Case" label followed by a text input field. Below that, there is a "File" label followed by a "Choose Files" button and the text "No file chosen". At the bottom right of the dialog, 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 will receive a case number and instructions for what to do next.



The image shows a 'Create Case' dialog box with a title bar containing a close button (X). The dialog contains several input fields and dropdown menus. The fields are: Name (text input), Email (text input), Phone (text input), Appliance (dropdown menu with 'Chennai' selected), Case Priority (dropdown menu with 'P3 - Normal' selected), Subject (text input), and Description (text area). At the bottom right, there are two buttons: 'Create' and 'Close'.

Name	<input type="text"/>
Email	<input type="text"/>
Phone	<input type="text"/>
Appliance	Chennai ▼
Case Priority	P3 - Normal ▼
Subject	<input type="text"/>
Description	<input type="text"/>

Create Close

Remote Access

Support > Technical Assistance > Remote Log Receiver

When working with Support to troubleshoot, you might be asked to allow access to your EdgeConnect devices during the online support session.

Change remote access information

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.

Enabled ☐

Start date

End date

Save Close

Partition Management

Support > Technical Assistance > Partition Management

Use this tab to regain Orchestrator disk space by selectively eliminating stats you no longer need.

Partition Management

Partition Management

807 Rows

Search

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 between the following different types of receivers: HTTP, HTTPS, KAFKA, SYSLOG, and WEBSOCKET. 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.

Complete the following instructions to add 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. Click **Save**.

HTTP Receiver Settings

Field	Description
Enable Receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are going to.
Log Type	Select the type of log from the list you want to apply.
URL	URL served by HTTP/HTTPS log server that Orchestrator will send log data with POST REST calls.
User Name	User name used in Basic Authentication when making REST calls (Optional).
Password	Password used in Basic Authentication when making REST calls. (Optional).
Repeat Password	Your password repeated.

HTTPS Receiver Settings

Field	Description
Enable Receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are going to.
Log Type	Select the type of log from the list you want to apply.
URL	URL of the HTTPS Receiver.
User Name	User name used in Basic Authentication when making REST calls (Optional).
Password	Password used in Basic Authentication when making REST calls (Optional).
Repeat Password	Your password repeated.

KAFKA Receiver Settings

Field	Description
Enable Receiver	Click this slider to toggle between enabled and disabled state.
Name	Name of the receiver the logs are going to.
Log Type	Select the type of log from the list you want to apply.
Topic	Topic name on KAFKA Receiver.
Bootstrap Servers	Domain name served by KAFKA Receiver. For example, "xxx.com:9092", "1.1.1.1:9092".
Acks	<p>Defines the amount of KAFKA servers that acknowledge a message before considering the message delivered.</p> <ul style="list-style-type: none"> acks=0: Expect no acknowledge. acks=1: Only leader server must acknowledge. ack=all: All servers must acknowledge.
Retries	Amount of times KAFKA will try before returning an error.
Batch Size	Multiple messages KAFKA will produce until the batch size is exceeded.
Buffer Size	Maximum memory size that can be used for buffering messages. When buffer size is exceeded, a message will be blocked.
Linger Time	Amount of time that KAFKA will wait before sending next message batch.

SYSLOG Receiver Settings

Field	Description
Enable Receiver	Click this slider to toggle between enabled and disabled state.

General Settings Section

Field	Description
Log Type	Type of log being sent to the SYSLOG receiver.
Protocol	Protocol being used between devices.
Hostname	Hostname of the SYSLOG receiver to identity the device.
Port	Port number of the SYSLOG receiver that accepts incoming events.
Custom Data	Custom data embedded inside the SYSLOG message.

Facility Settings Section

Field	Description
Audit Log	Type of audit log.

Audit Log Severity Settings Section

Field	Description
Error	Severity level of the error; select from the drop-down menu.
Info	Severity level of the information; select from the drop-down menu.
Debug	Severity level of the debug; select from the drop-down menu.

WEBSOCKET Receiver Settings

Provides a reliable streaming mechanism for alarms and Orchestrator audit logs across all appliances. It is initiated from the client side and sent to Orchestrator for authentication. When authenticated by Orchestrator, asynchronous notifications are sent in JSON objects.

Field	Description
Enable	Click this slider to toggle between enabled and disabled state.
Name	Name of the WebSocket receiver.
Log Type	Type of log being sent to the WebSocket receiver.
IP Allow List	List of source IP addresses that are allowed WebSocket access to Orchestrator.

WebSocket Receiver Configuration

You need the following items to establish connectivity from Orchestrator to the WebSocket receiver:

- Key generated by Orchestrator after the above configuration is completed
- ID created by Orchestrator when it is configuring the WebSocket server

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	Description
Appliance Name	Name of the appliance.
Peer ID	ID of the peer.

Field	Description
Peer Name	Name of the peer.
Role	Whether the hub or spoke topology is being used for the specified peer.
Last Transmission Count	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	Last transaction count from the peer that was received.
Time since Last Received	Amount of time since the last received update.
MainVer and Region	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

The RMA (Return Merchandise Authorization) Wizard automates the RMA process for an exchange or replacement of your appliance, if needed. It includes appliance discovery, the version of the appliance, and a backup selection. Use this screen as instructed by Support to prepare an RMA.

Note the following before you begin the RMA process.

- Upgrade or downgrade the new appliance to the same software version before shipping to the site. This will save time.
- Perform a backup of the Orchestrator and EdgeConnect appliances.
- Install the new EdgeConnect appliance onsite.
- When Orchestrator discovers the new device, do not approve it. Start the RMA process to move the license to the new EdgeConnect appliance.

Run the RMA Wizard

Complete the following steps to RMA your appliance.

1. Navigate to the **RMA** tab in Orchestrator.
2. Select the appliance you want to replace from the menu.

NOTE The IP address, appliance model, hostname, serial number and software version will auto-populate after you select the appliance.

3. Select the newly discovered appliance that will replace the current appliance.

NOTE The IP address, appliance model, hostname, serial number and software version will auto-populate after you select the appliance.

4. Click **Next** >.
5. If you are adding a backup appliance, proceed to the next section. Otherwise, click **Apply**.

The Applying Configuration dialog box opens and displays the status of the upgrade and restore.

Add a Backup Appliance

If you choose to add a backup appliance from the table, complete the following steps.

1. Select the backup appliance from the table.
2. Select the version you want the backup appliance to have from the drop down menu.

NOTE If your selection results in a software downgrade, a backup must be provided.

Upgrade and Downgrade

If the software version you selected for your backup appliance is **higher** than that of the discovered appliance, you will need to do the following:

- Upgrade to the new version using Orchestrator.
- Back up the appliance from a restore, if applicable.

If the software version you selected for your backup appliance is **lower** than that of the discovered appliance, you will need to do the following:

- Install the desired version as a next boot on the appliance.
- Restore from backup.

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 ×

Built-in Policies ⓘ 7 mins

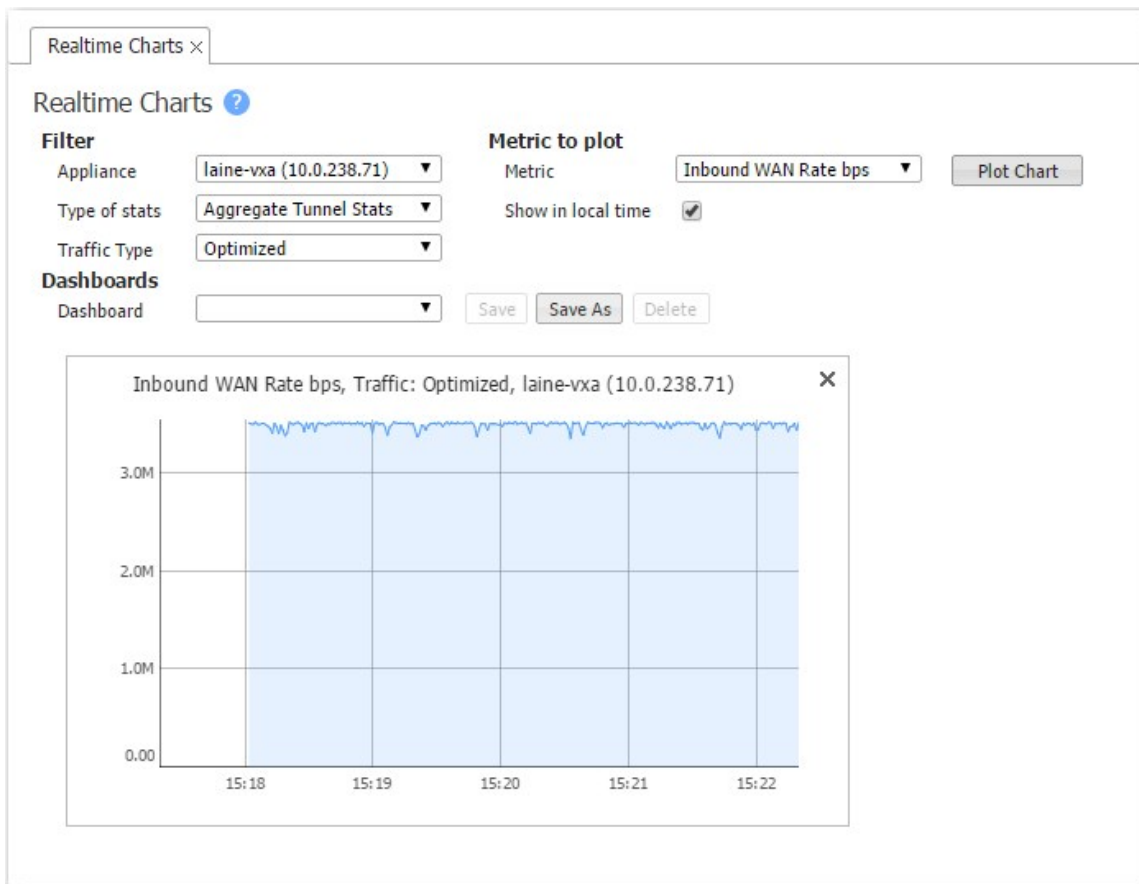
660 Rows

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.



1. Select the filters you want, and then click **Plot**.

The chart appears at the bottom of the page.

2. To save as a dashboard, click **Save As**, and then enter a name for your dashboard. Do not include spaces in your name. Click **Save**.

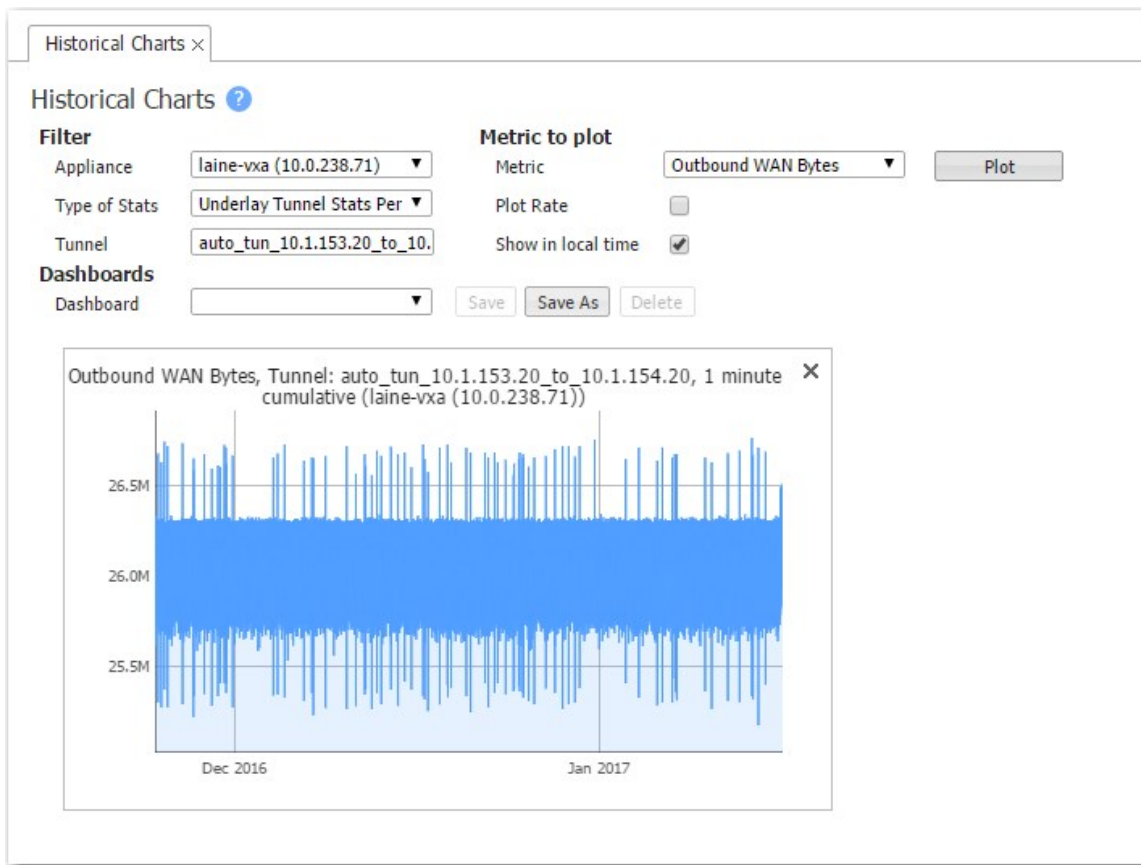
If successful, a green Success bar appears and the dashboard name shows up in the **Dashboard** field.

To retrieve it later, go to this tab and choose the dashboard from the drop-down list.

Historical Charts

Support > Reporting > Historical Charts

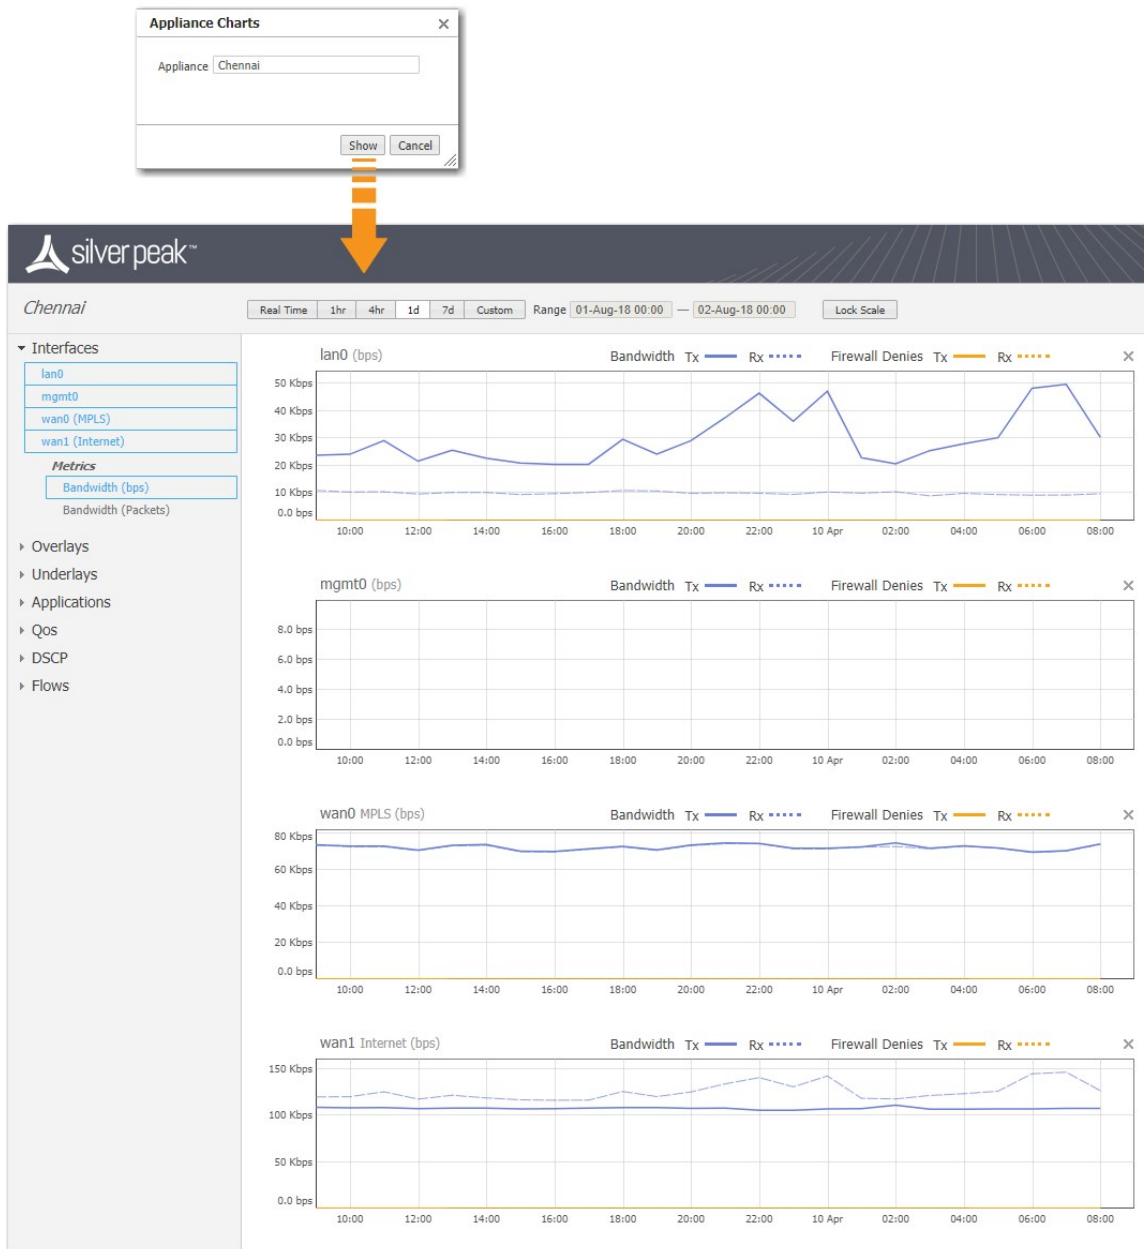
As an aid to troubleshooting, **Historical Charts** are useful for reviewing the performance of individual appliances. You can save sets of charts as dashboards.



Appliance Charts

Support > Reporting > Appliance Charts

Use this dialog box to access an individual appliance's realtime and historical charts.

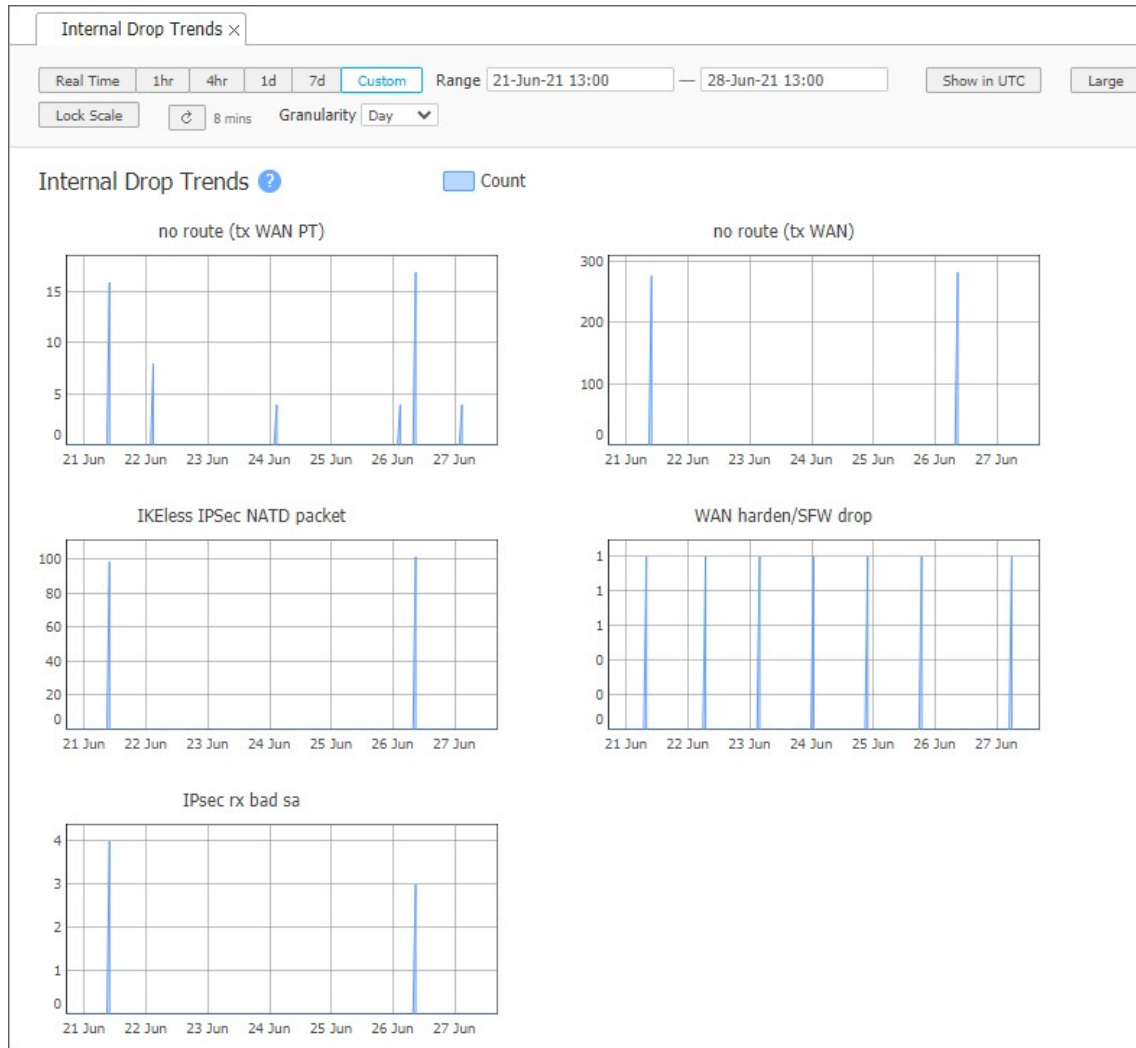


Internal Drop Trends

Support > Reporting > Dropped Packet Trends


The **Internal Drop Trends** report shows internal packet drop trends for a single selected appliance. The charts that are displayed will vary according to the cause of the drop.

Charts are available in real time or for a specific time period. Real time charts show drops over the last five minutes and refresh every five seconds.



You can customize the chart settings using the controls at the top of the tab, as follows:

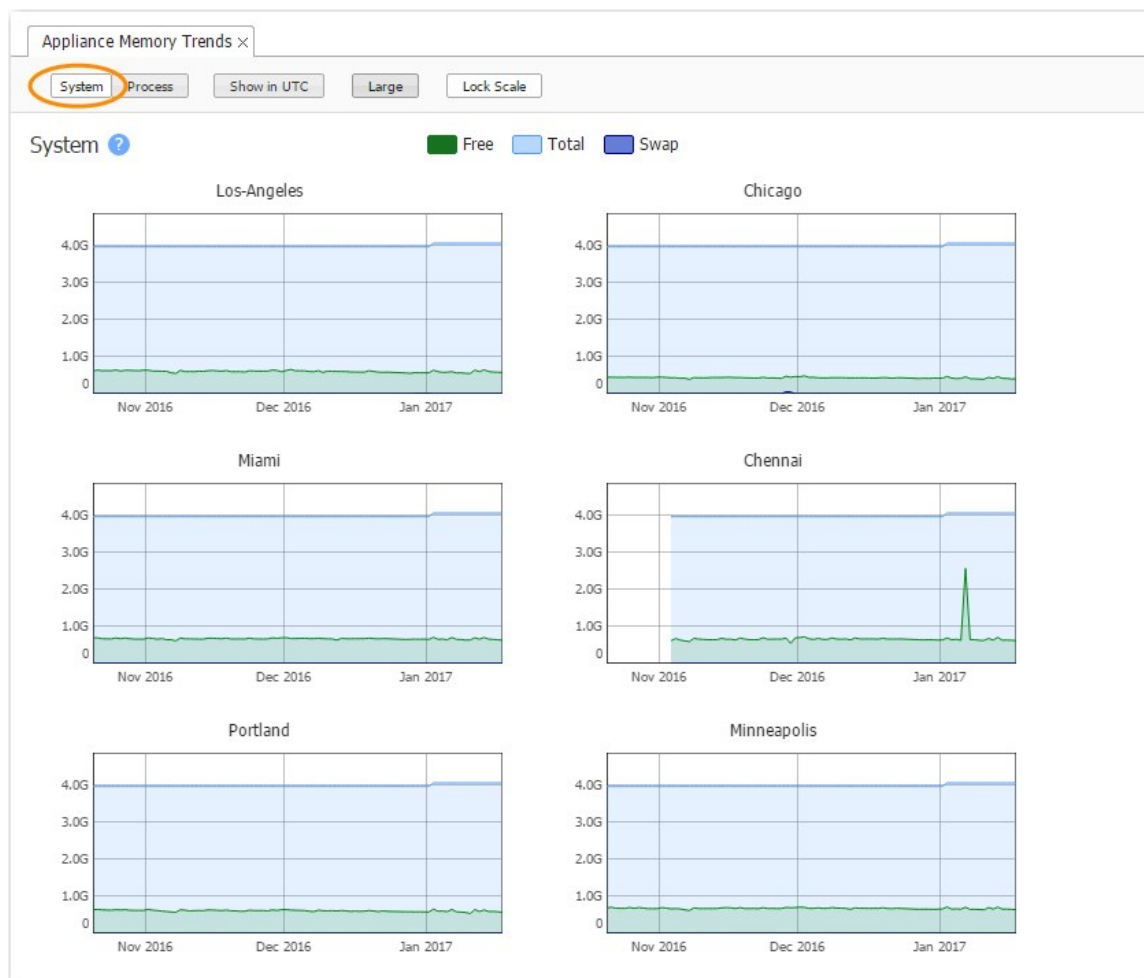
Option	Description
Time period	<ul style="list-style-type: none"> Click Real Time to enable live statistics for all available interfaces. Click a predefined time period (1h, 4h, 1d, 7d) to display statistics over the last hour, four hours, day, or seven days. Click Custom and set your own custom time range to display statistics for that time period.
Show in UTC	Click this option to toggle chart times between local appliance time or UTC.
Large	Click this option to toggle the size of the charts between smaller (default) and large.

Option	Description
Lock Scale	By default, each chart uses its own scale that is relative to the data displayed. Click this option to apply and lock the same scale to each chart.
Refresh 	Click the Refresh button to fetch data again for the selected time period.
Granularity	When a custom time period is used, select the granularity level to be applied to charts (Minute , Hour , or Day).

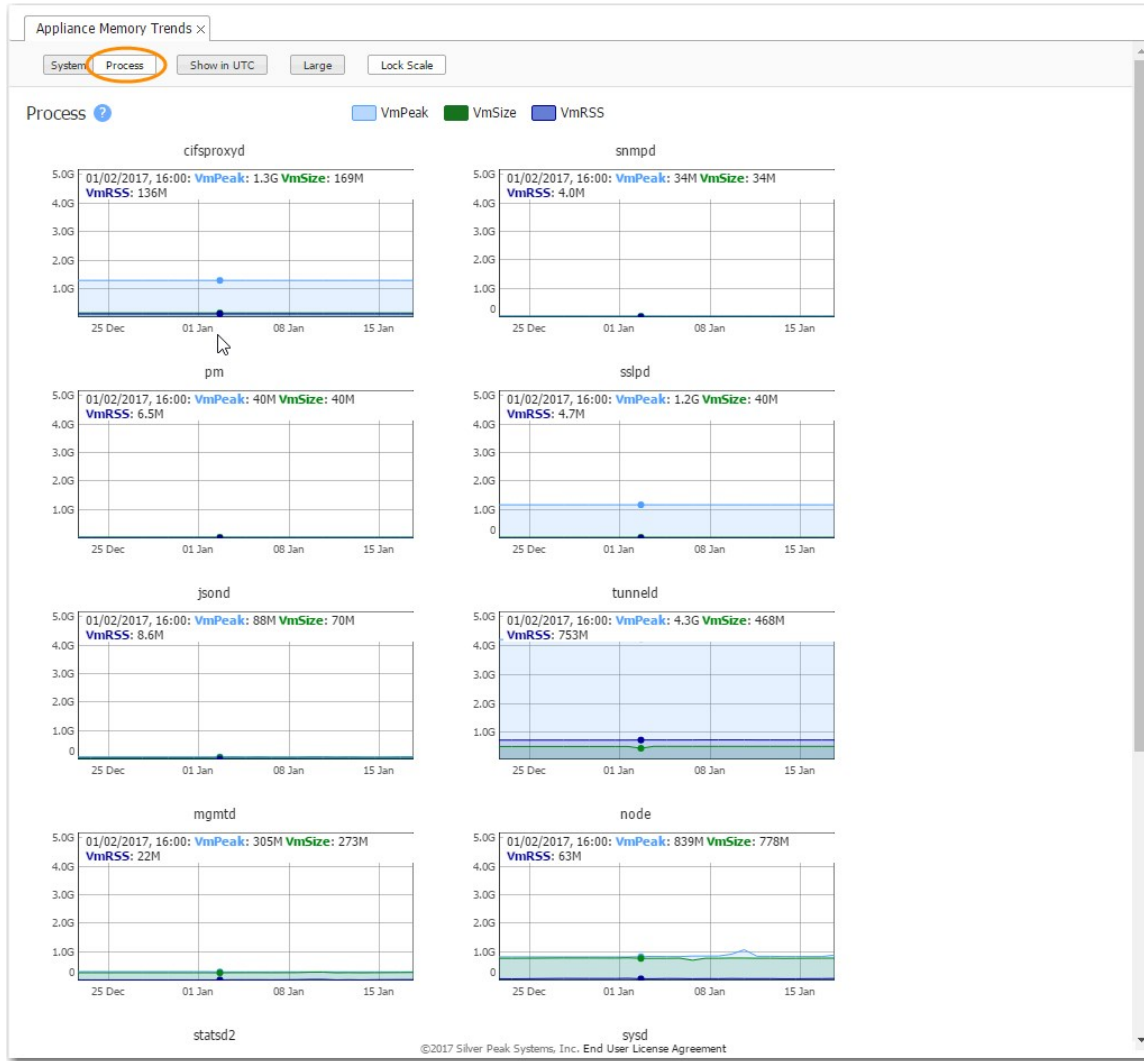
Appliance Memory Trends

Support > Reporting > Appliance Memory Trends

The **System** view shows appliance daily memory usage.



The **Process** view is for individual appliances.

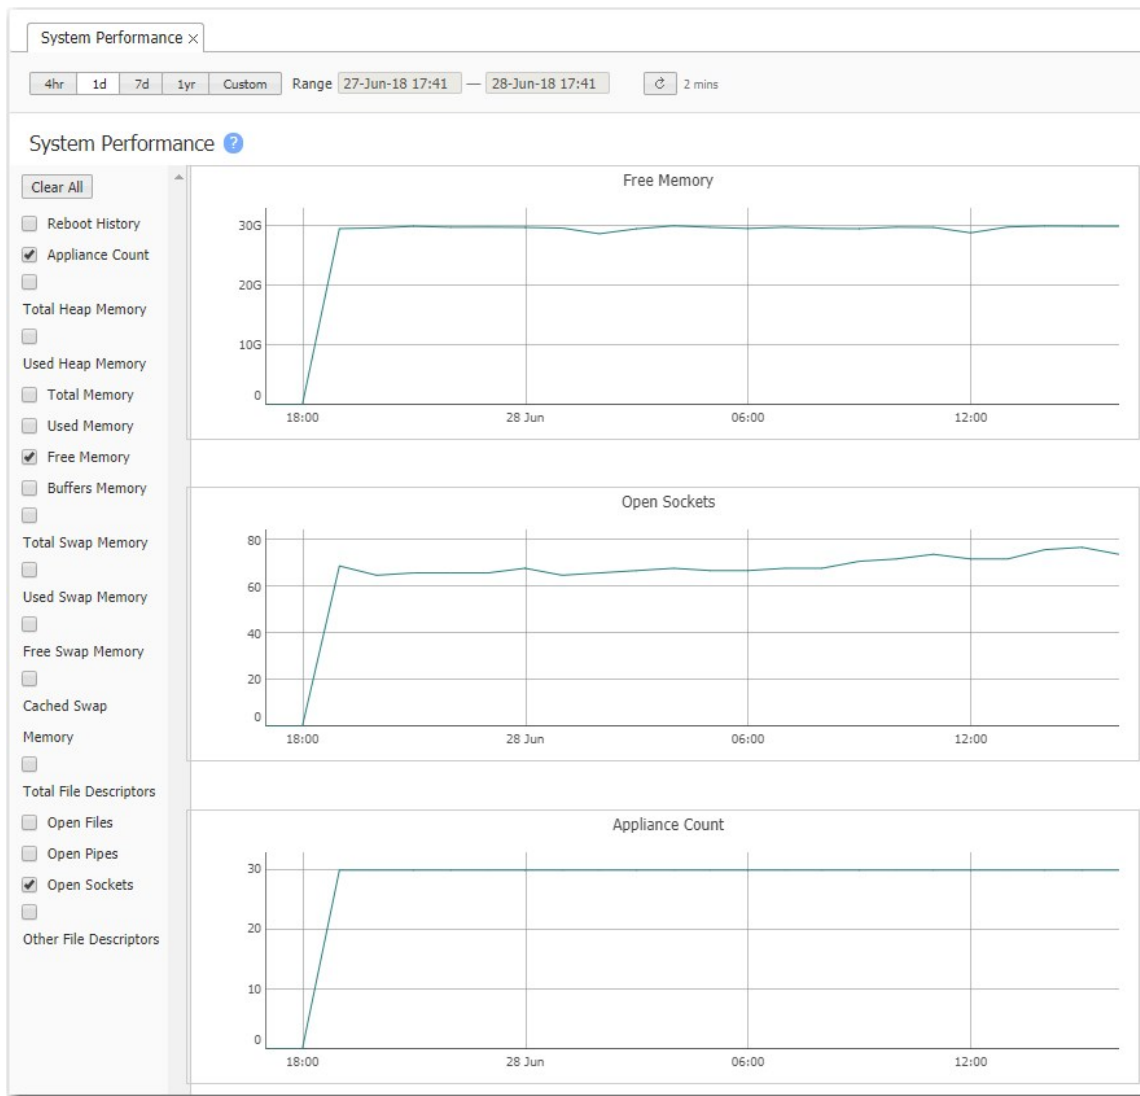


System Performance

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 CPU Usage

Support > Reporting > Appliance CPU Usage

The charts on this page provide real-time views of combined and individual CPU usage statistics for a single selected appliance. Charts show the past five minutes of usage and refresh every five seconds. By default, only total utilization is displayed on the charts. You can toggle the available statistics on or off by clicking the sample indicator line next to each statistic name.

NOTE On appliances with Boost enabled, it is common for non-CPU0 cores to run at or close to 100%. CPU0 will show occasional spikes of high usage when stats are rolled up and archived.



Appliance Crash Report

Support > Reporting > Appliance Crash Report

This report lists appliance crashes, which you can forward to Silver Peak.

Appliance Crash Report ×

3 Rows

Search

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

Send To Silver Peak

Close

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

Search

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	59e632f26cf525dea4dbd0	e0672c9a-a8ef-4d01-8c71-0235d7c1a158
1.NE	San-Jose		10.0.185.47	PORTAL	0.0.0.0_71872	c3aa83d6-8a73-4ca9-a6c4-e7c98779f8c7	59e6344526cf525dea4dbf4	c212c6c4-e02c-40f3-ac1f-cfa52bfdd5f3
10.NE	Minneapolis		10.0.185.28	PORTAL	0.0.0.0_71872	59e88c6e-681a-4c49-a874-5ae31f7e50af	59e632e526cf525dea4dbcc	f0b9043a-fd29-4a88-9520-8e23e46a3a9a
11.NE	Mumbai		10.0.185.44	PORTAL	0.0.0.0_71872	282c3a7a-c364-4524-8542-b244882b96e3	59e6343826cf525dea4dbf2	6ed5bd3c-0fc3-4876-a955-91794d28ca14
12.NE	Los-Angeles		10.0.185.25	PORTAL	0.0.0.0_71872	3082126c-fd60-49da-aba4-25b7e540253e	59e632f26cf525dea4dbcc	70540b55-903e-41e6-a207-6e130e886ed5
13.NE	Portland		10.0.185.26	PORTAL	0.0.0.0_71872	6c83b38e-0218-4cfc-aa0a-ed93da713fa8	59e632b726cf525dea4dbcc	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	59e632af26cf525dea4dbcc	8c843b37-a6a7-4537-af4a-e690de617f4a
15.NE	Singapore		10.0.185.30	PORTAL	0.0.0.0_71872	542eefcf-4117-4dd7-bf50-77cabeb3837e	59e632db26cf525dea4dbca	669c74ac-8b1b-4b41-bcda-8e979d2880ac
16.NE	Milan		10.0.185.31	PORTAL	0.0.0.0_71872	6db04d99-0a5f-467d-b2fe-7d19ef6d096e	59e632f826cf525dea4dbcc	d4b6146e-5997-4f80-9e4f-e4138b8ab69
17.NE	Paris		10.0.185.33	PORTAL	0.0.0.0_71872	97a2b896-5e60-4403-a35a-205cae11d18b	59e6331526cf525dea4dbd5	147a7381-c133-4a5a-8038-ac6418db1c1f
18.NE	Tokyo		10.0.185.34	PORTAL	0.0.0.0_71872	52901e40-41a6-43cd-84f2-87f7c69f298f	59e6332d26cf525dea4dbd7	3f317d1-60e9-41bf-84e1-52514312133f
19.NE	London		10.0.185.37	PORTAL	0.0.0.0_71872	fd25bb96-8d0b-479f-8f8a-c0ce36ce632d	59e6337c26cf525dea4dbde	cb8d24e-348b-4a5c-b3d6-e54eac883a10
2.NE	Geneva		10.0.185.45	PORTAL	0.0.0.0_71872	ec3ce87e-c2e6-4edf-bcfa-c186e59805fc	59e6344f26cf525dea4dbf6	3b734dd4-b1f6-4805-bc0b-b21848263f35
20.NE	Frankfurt		10.0.185.38	PORTAL	0.0.0.0_71872	e4c333f3-f872-4ae3-9ef2-c25ed38c2f59	59e6339326cf525dea4dbd0	1f6c1c70-f235-497a-85fc-14d68e17f236
21.NE	Barcelona		10.0.185.40	PORTAL	0.0.0.0_71872	1988e63f-a999-4e58-a883-31799889d1d0	59e633bd26cf525dea4dbd8	58d0c7f8-9ee5-484f-b082-7392eda8541d
22.NE	Seoul		10.0.185.41	PORTAL	0.0.0.0_71872	6034357a-3014-4458-bf1d-96792a1c7c35	59e633d726cf525dea4dbda	4f1b3d8a-f195-4285-82d9-adf8e84f39f
23.NE	Osaka		10.0.185.42	PORTAL	0.0.0.0_71872	a5f28a69-1663-438a-a233-f2d6615eb332	59e633ee26cf525dea4dbec	e69168f7-683d-438d-a3ec-4dc56a5e6b46
24.NE	San-Antonio		10.0.185.46	PORTAL	0.0.0.0_71872	752eaf6c-c34f-43e4-8c3d-00270a5adc2c	59e6346726cf525dea4dbfa	318091a7-fbd1-4b5c-ba2d-2dd8c9141e1b
25.NE	New-Orleans		10.0.185.48	PORTAL	0.0.0.0_71872	0a6c5b3c-e013-4473-bbc7-9e0d5d2b9bf7	59e6345c26cf525dea4dbf8	3d9650b9-e9ee-4e20-b864-2214d8d9fd2c
26.NE	Toronto		10.0.185.52	PORTAL	0.0.0.0_71872	237a7b0d-9173-4c3b-a3bc-6871f5c3881d	59e634d826cf525dea4db04	00ba0f0e-1735-44a2-944f-9aaf25df753f
27.NE	New-York		10.0.185.53	PORTAL	0.0.0.0_71872	eb01d02f-14ab-430f-89bb-f6b0355498c4	59e634ee26cf525dea4db08	be23974e-8882-4d54-a15a-a129269a725a
28.NE	Pittsburgh		10.0.185.49	PORTAL	0.0.0.0_71872	9bb8440b-5964-4ff4-b388-430cd02c7160	59e6347a26cf525dea4dbfc	62df50d8-a56b-42c8-b302-7a680d9e4bae
29.NE	Albuquerque		10.0.185.50	PORTAL	0.0.0.0_71872	fbc9a141-8b09-4585-a84f-1a574e31ae89	59e634a626cf525dea4dbfe	b8d36639-2b1d-458b-a525-5d34df9ad225
3.NE	Edinburgh		10.0.185.54	PORTAL	0.0.0.0_71872	13c4f90a-d707-48f8-9a0b-468aaf6a0a66	59e634a826cf525dea4dbf6	55a1c71c-9f8f-4a11-9a7f-a7f6a179a7d3

Close

IPSec UDP Status

Support > Reporting > IPSec UDP Status

Use this tab to review and monitor the IPSec UDP key material status for all appliances in your network.

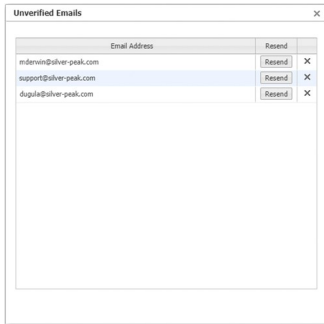
Field	Description
Host Name	Host name of the appliance.
Active Key	Indicates whether the appliance is using the active IPSec UDP key.
Active Key Pushed Time	Time when the active key was pushed to the appliance.
Active Key Activation Time	Time when the key was activated on the appliance.
Reachability	Indicates whether the appliance is reachable.
Details	Additional details about reachability or key material status.

Unverified Emails

Support > Reporting > Unverified Emails

When you add an email address to either the Alarms or the Reports email distribution list, Orchestrator sends the recipient an email that contains a link, asking them to click to provide verification.

If Orchestrator does not receive a verification, either the recipient has not responded or the email address is invalid.



- An unverified email address remains inactive and does not 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.

¹ If roles and appliance access group keys are not provided, Orchestrator inspects its own configuration to determine the role and appliance access group for the user. If it does not find that information, the user is not allowed to log in.