

Silver Peak

Appliance Manager Operator's Guide

VXOA 8.0

May 2016

PN 200030-001 Rev Q

Silver Peak Appliance Manager Operator's Guide

Document PN 200030-001 Rev Q

Date: May 2016

Copyright © 2016 Silver Peak Systems, Inc. All rights reserved. Information in this document is subject to change at any time. Use of this documentation is restricted as specified in the *End User License Agreement*. No part of this documentation can be reproduced, except as noted in the *End User License Agreement*, in whole or in part, without the written consent of Silver Peak Systems, Inc.

Trademark Notification

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

Warranties and Disclaimers

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

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

1.877.210.7325 (toll-free in USA)
+ 1.408.935.1850

www.silver-peak.com/support

Contents

Preface	v
Who Should Read This Manual?	v
Manual Organization	v
Technical Support	vi
Appliance Management Options	vi
Chapter 1 Getting Started	1
Introduction	2
Summary of Configuration Tasks	2
Getting Started with Deployment	3
Deployment Modes	4
How You Can Adjust the Basic Deployments	9
Adding Data Interfaces	10
Modifying Interface Configurations	11
Configuring Next-Hops	12
Management Next-Hops	12
WAN Next-Hops	12
LAN Next-Hops	13
Taking a Quick Look at the System Page	14
Deciding whether to build tunnels automatically	14
Manually Adding Local Subnets	15
Adding SSL Certificates and Keys for Deduplication	16
Chapter 2 Configuring the System and Network	17
System Information	18
Deployment	20
Mapping Labels to Interfaces	20
LAN-side Configuration: DHCP	20
WAN-side Configuration	21
Basic Deployments	22
How You Can Adjust the Basic Deployments	27
Adding Data Interfaces	28
Definitions	28
Interfaces	29
Routes	30
Tunnels	32
Shaper	37
Subnets	41
SSL Certificates	43
SSL CA Certificates	45
VRRP	46
WCCP	48
System Limits	51
Chapter 3 Configuring Policies	53
Route Policies	54
QoS Policies	56
Optimization Policies	60
Access Lists	62

Built-in Applications	63
User-Defined Applications	64
Application Groups	66
Flow Redirection	67
SaaS Optimization	68
NAT Policies	71
Advanced Settings	72
Threshold Crossing Alerts	74
Chapter 4 Monitoring Traffic	77
Introduction	78
Understanding Traffic Direction	78
Application View Tab	79
Network View Tab	80
Charts	81
Bandwidth	83
Reduction	83
Packets per Second	84
Flow Counts	84
Latency	85
Loss	85
Out-of-Order Packets	86
Application Statistics	87
Table View	87
Pie View	88
Realtime Charts	89
Flows	90
How the Page is Organized	91
Customizing Which Columns Display	93
Flow Details	94
Resetting Flows to Improve Performance	107
QoS Statistics	108
Tunnel Statistics	109
Internet Statistics	111
NetFlow Statistics	112
Interface Statistics	113
Bridge Interfaces	115
Route Next Hops	116
VRRP Status	117
Flow Redirection Statistics	119
SaaS Optimization	121
Orchestrator Reachability	122
Chapter 5 Monitoring Alarms	123
Understanding Alarms	124
Categories of Alarms	124
Types of Alarms	125
Viewing Current Alarms	135
Chapter 6 Administration Tasks	137
Date and Time	138
Domain Name Server (DNS)	139
SNMP	140
Loading SNMP MIBs	141

Netflow	143
License & Registration	144
Management IP / Hostname	145
Silver Peak Cloud Portal	146
Users	147
Auth/RADIUS/TACACS+	148
Banners	150
Session Management	151
Chapter 7 Maintenance & Support	153
Viewing System Information	154
Upgrading Appliance Software	155
Backing Up and Restoring Appliance Configuration Files	156
Disk Management	157
Putting the Appliance in System Bypass Mode	158
Testing Network Connectivity with Ping and Traceroute	159
Capturing Packets for Analysis	163
Erasing Network Memory	164
Rebooting or Shutting Down an Appliance	165
How to Contact Support	166
Logging In to the Support Portal	167
Viewing Logs	168
Configuring Logging Parameters	169
Managing Debug Files	171
Appendix A TCP/IP Ports Used by the Orchestrator and Silver Peak Appliances	173
Ports used by the Orchestrator	174
Ports used by Silver Peak Appliances	175
Appliance Management Plane	175
Appliance Data Plane	175



Preface

Silver Peak gives enterprises and service providers the flexibility to securely connect users to their applications via the most cost-effective source of connectivity available.

Who Should Read This Manual?

Anyone who wishes to manage Silver Peak appliances should refer to this manual. The chapters reflect the menu structure, and in most cases provide the same content available in the on-line help.

Users should have some background in Windows[®] terminology, Web browser operation, and a knowledge of where to find the TCP/IP and subnet mask information for their system.

Manual Organization

This section outlines the chapters and summarizes their content.

[Chapter 1, “Getting Started,”](#) describes the fundamentals and considerations of setting up a basic first deployment. Additionally, it describes how to work with the routing table, modify network interface parameters, configure gigabit etherchannel bonding, and add SSL certificates and keys for optimizing encrypted traffic.

[Chapter 2, “Configuring the System and Network,”](#) describes the tabs for configuring network and system-level parameters.

[Chapter 3, “Configuring Policies,”](#) describes the tabs for configuring appliance policies.

[Chapter 4, “Monitoring Traffic,”](#) describes how to view realtime and historical statistics for applications, current flows, QoS, tunnels, data reduction, bandwidth optimization, flow counts, latency, flow redirection, NetFlow, interfaces, and more.

[Chapter 5, “Monitoring Alarms,”](#) describes alarms categories and definitions. It also describes how to view and handle alarm notifications.

[Chapter 6, “Administration Tasks,”](#) describes menus related to basic appliance administration and user management.

[Chapter 7, “Maintenance & Support,”](#) describes how to use these menus to perform various system maintenance tasks. It provides tools, as well as logs and access to technical assistance.

[Appendix A, “TCP/IP Ports Used by the Orchestrator and Silver Peak Appliances,”](#) lists and diagrams the ports used by the appliances and the Orchestrator, on the management and data planes.

Technical Support

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

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

We're dedicated to continually improving the usability of our products and documentation. If you have suggestions or feedback for our documentation, please send an e-mail to **techpubs@silver-peak.com**.

For usability suggestions, questions, or issues, please send an e-mail to **usability@silver-peak.com**.

Appliance Management Options

Silver Peak provides a variety of ways for you to access and configure the appliances, as well as review statistics and events across a Silver Peak network:

- **Appliance Manager WebUI:** The Silver Peak Appliance can be managed through the web-based Appliance Manager.
- **Command Line Interface (CLI):** You can manage the Silver Peak Appliance through the CLI. You can access the full-featured CLI either locally, through the RS-232 serial (console) port, or remotely, through a Secure Shell (SSH) connection.
- **Unity Orchestrator:** This is a comprehensive platform for deployment, management, and monitoring of a Silver Peak-enabled WAN. In addition to centralizing the administration of the Silver Peak appliances, the Orchestrator provides detailed visibility into all aspects of application delivery across a distributed enterprise, including application behavior, WAN performance, Quality of Service (QoS) policies, and bandwidth utilization.
- **SNMP:** The appliances work with standard and proprietary SNMPv2c traps.



Getting Started

This chapter describes the fundamentals of setting up a basic deployment.

In This Chapter

- **Introduction** See page 2.
- **Getting Started with Deployment** See page 3.
- **Modifying Interface Configurations** See page 11.
- **Configuring Next-Hops** See page 12.
- **Taking a Quick Look at the System Page** See page 14.
- **Adding SSL Certificates and Keys for Deduplication** See page 16.

Introduction

When you first install the appliance and log in via the browser, the *Initial Configuration Wizard* appears. The wizard guides you through configuring management settings, deployment and network settings, and creating a tunnel to a remote appliance. With simpler deployments, this is enough to start optimizing traffic.

You can always access the wizard again later by going to the **Configuration** menu and selecting **Initial Config Wizard**.

In most cases, the individual wizard pages are identical to tabs you can access via the menus.

The following tasks provide a context for assessing your network with questions to consider before appliance deployment or redeployment.

Summary of Configuration Tasks

The configuration steps are as follows:

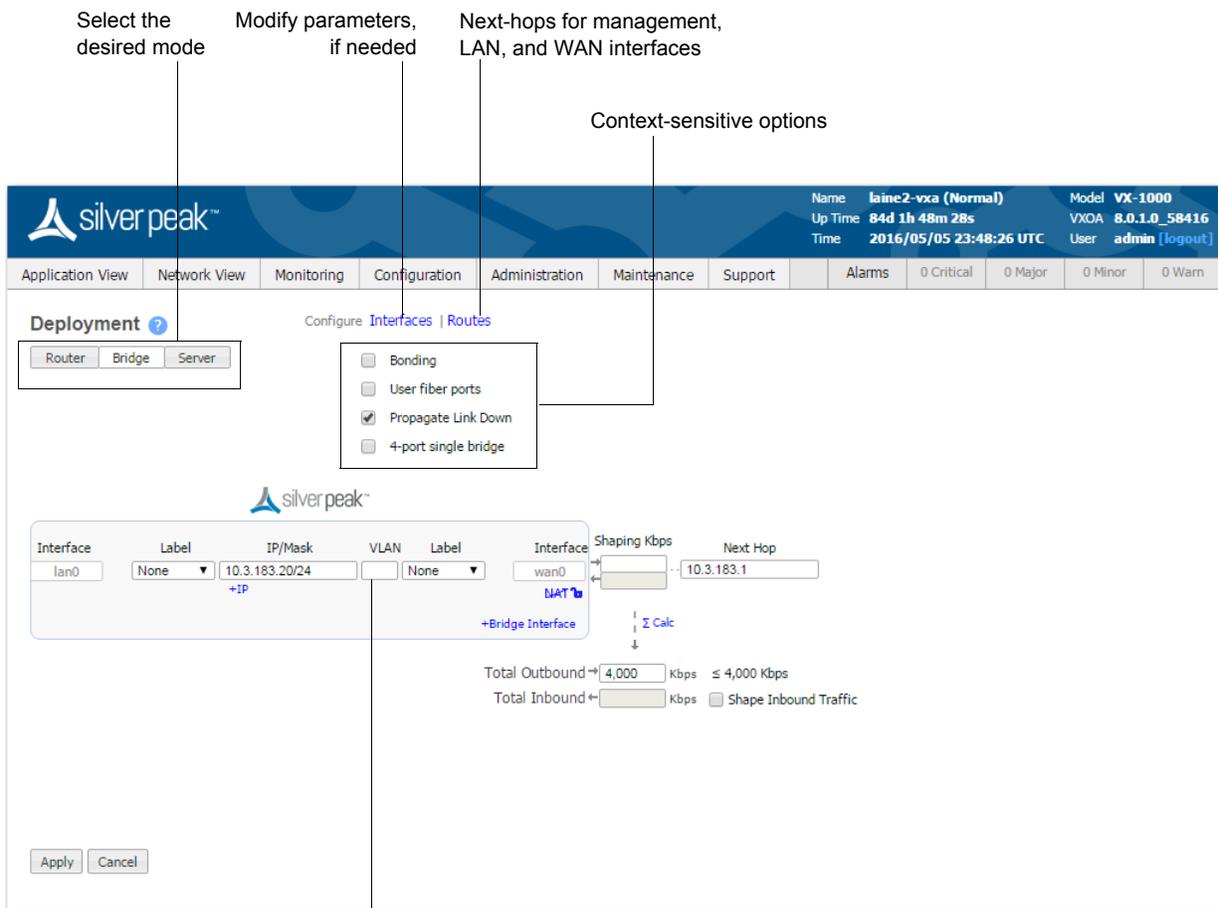
Task	Questions to consider	For detailed instructions, see...
1 Determine which deployment mode is best for your circumstances	MENU: <i>Configuration > Deployment</i> <ul style="list-style-type: none"> Is the appliance be in-path or out-of-path? Will you need to redirect traffic? If so, in which direction(s) and on which devices? If you have a 4-port physical appliance, will you use gigabit etherchannel bonding? 	"Getting Started with Deployment" on page 3.
2 Review and validate the interface configurations	MENU: <i>Configuration > Interfaces</i> <ul style="list-style-type: none"> Review the DHCP settings If in router mode, do you want to harden any of the WAN interfaces? 	"Modifying Interface Configurations" on page 11. Chapter 2, "Configuring the System and Network," "Interfaces" on page 29.
3 Configure the next-hops	MENU: <i>Configuration > Routes</i> <ul style="list-style-type: none"> Do you need to add any static routes? 	"Configuring Next-Hops" on page 12.
4 Verify system-level options	MENU: <i>Configuration > System</i> MENU: <i>Configuration > Subnets</i> <ul style="list-style-type: none"> Will you want to use auto-tunnels? Will you be using subnet sharing? Do you want to share any local subnets that aren't tunnel endpoints? 	"Taking a Quick Look at the System Page" on page 14.
5 Install SSL certificates and keys	MENU: <i>Configuration > SSL Certificates</i> MENU: <i>Configuration > SSL CA Certificates</i> <ul style="list-style-type: none"> Will you be optimizing SSL traffic? 	"Adding SSL Certificates and Keys for Deduplication" on page 16. "SSL Certificates" on page 43. "SSL CA Certificates" on page 45.

Getting Started with Deployment

This section discusses the basics of three deployment modes on the **Configuration > Deployment** page.

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

 For step-by-step examples of more complex deployments, see the *Silver Peak Network Deployment Guide*.



Select the desired mode

Modify parameters, if needed

Next-hops for management, LAN, and WAN interfaces

Context-sensitive options

Apply Cancel

A **VLAN Tag** is required if the appliance is installed on a VLAN trunk and an untagged VLAN is unavailable.

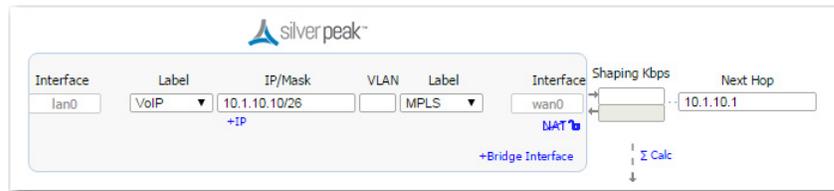
Deployment Modes

There are three basic deployment modes: **Bridge**, **Router**, and **Server**.

Bridge Mode

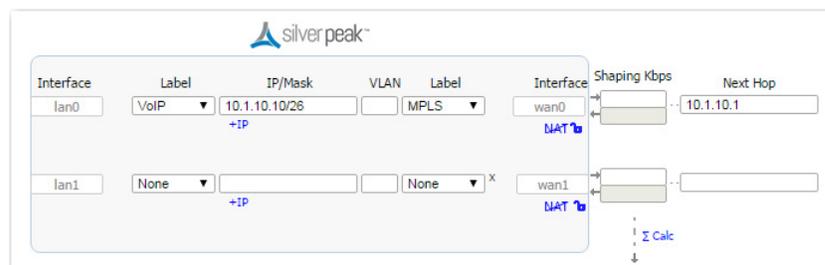
■ Single WAN-side Router

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



■ Dual WAN-side Routers

This is the most common 4-port bridge configuration.



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

■ Considerations for Bridge Mode Deployments

- Do you have a physical appliance or a virtual appliance?

A virtual appliance has no fail-to-wire, so you would need a redundant network path to maintain connectivity if the appliance fails.

- If your LAN destination is behind a router or L3 switch, you need to add a LAN-side route (a LAN next-hop).
- If the appliance is on a VLAN trunk, then you need to configure VLANs on the Silver Peak so that the appliance can tag traffic with the appropriate VLAN tag.

Router Mode

There are four options to consider:

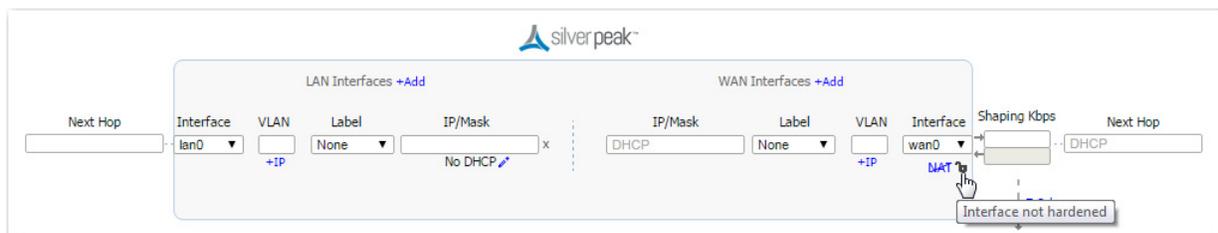
- 1 Single LAN interface & single WAN interface
- 2 Dual LAN interfaces & dual WAN interfaces
- 3 Single WAN interface sharing LAN and WAN traffic
- 4 Dual WAN interfaces sharing LAN and WAN traffic

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

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

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

■ #1 - Single LAN Interface & Single WAN Interface



For this deployment, you have two options:

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

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

■ #2 - Dual LAN Interfaces & Dual WAN Interfaces

The screenshot shows the Silver Peak Appliance Manager configuration interface. It is divided into two main sections: LAN Interfaces and WAN Interfaces.

LAN Interfaces +Add:

Next Hop	Interface	VLAN	Label	IP/Mask
	lan0	+IP	Replicati	No DHCP
	lan1	+IP	VoIP	No DHCP

WAN Interfaces +Add:

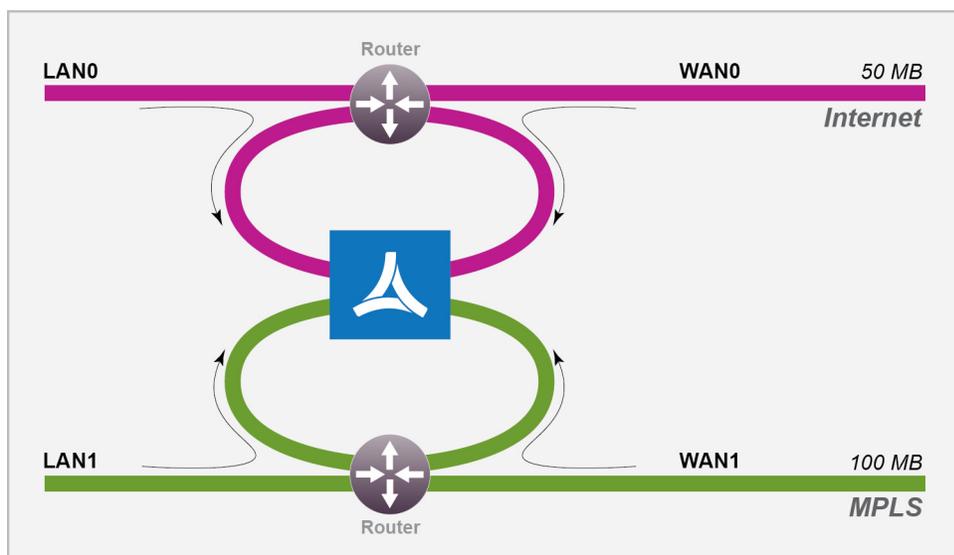
IP/Mask	Label	VLAN	Interface	Shaping Kbps	Next Hop
x DHCP	Internet	+IP	wan0		DHCP
x DHCP	MPLS	+IP	wan1		DHCP

Additional controls include a 'Σ Calc' button and a downward arrow at the bottom right.

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

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

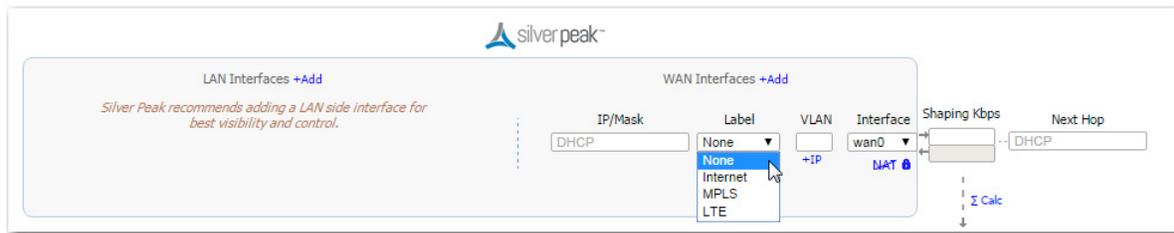
Out-of-path dual LAN and dual WAN interfaces



For this deployment, you have two options:

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

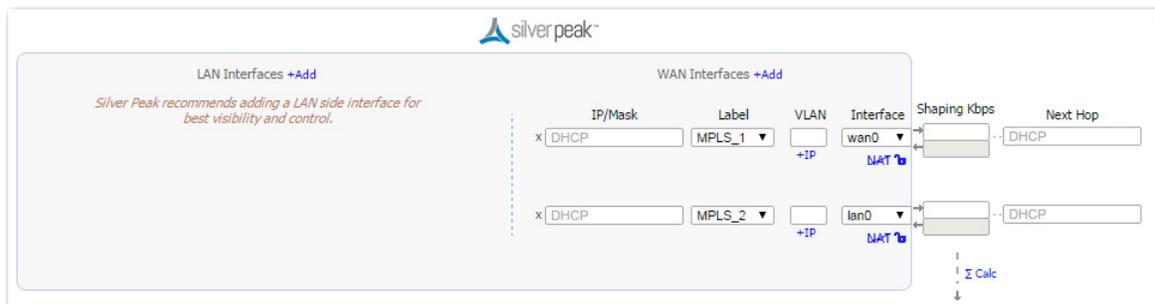
■ #3 - Single WAN Interface Sharing LAN and WAN traffic



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

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

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



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

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

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

■ Considerations for Router Mode Deployments

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

- In the rare case when you need to send inbound WAN traffic to a router other than the WAN next-hop router, use LAN-side routes.

■ Examining the Need for Traffic Redirection

Whenever you place an appliance out-of-path, you must redirect traffic from the client to the appliance.

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

- **PBR** (Policy-Based Routing) — configured on the router. No other special configuration required on the appliance. This is also known as FBR (Filter-Based Forwarding).
If you want to deploy two Silver Peaks at the site, for redundancy or load balancing, then you also need to use **VRRP** (Virtual Router Redundancy Protocol).
- **WCCP** (Web Cache Communication Protocol) — configured on both the router and the Silver Peak appliance. You can also use WCCP for redundancy and load balancing.
- **Host routing** — the server/end station has a default or subnet-based static route that points to the Silver Peak appliance as its next hop. Host routing is the preferred method when a virtual appliance is using a single interface, mgmt0, for datapath traffic (also known as Server Mode).

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



High availability — as configured with **VRRP** and **WCCP** — are covered separately, and in depth, in the *Silver Peak Network Deployment Guide*.

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

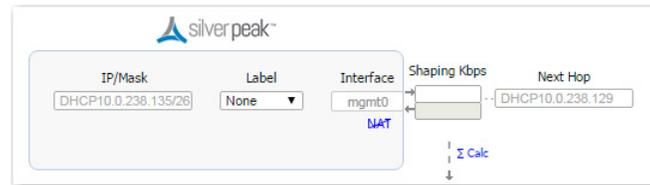
- If you use **subnet sharing** (which relies on advertising local subnets between Silver Peak appliances) or route policies (which specify destination IP addresses), then you only need LAN-side redirection.
- If, instead, you rely on **TCP-based** or **IP-based** auto-optimization (which relies on initial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

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

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

Server Mode

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



How You Can Adjust the Basic Deployments

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



Bonding

- When using an NX 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 **bwano**. This increases throughput on a very high-end appliance and/or provides interface-level redundancy.
- For bonding on a virtual appliance, you would need configure the host instead of the appliance. For example, on a VMware ESXi host, you would configure NIC teaming to get the equivalent of etherchannel bonding.
- Whether you use a physical or a virtual appliance, etherchannel must also be configured on the directly connected switch/router. Refer to the Silver Peak user documentation.

For more information, see [“Configuring Gigabit Etherchannel Bonding” on page 10](#).

Use Fiber Ports

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

Propagate Link Down

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

4-port single bridge

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



Note Changing the deployment mode requires a reboot.

Configuring Gigabit Etherchannel Bonding

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

When you configure bonding, the following is true:

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

◆ To configure etherchannel bonding

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

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

```

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

show etherchannel
show interface port-channel <1>

```

Adding Data Interfaces

- You can create additional data-plane Layer 3 interfaces, to use as tunnel endpoints.
- To add a new interface, click **+IP Address**.
- Add a **VLAN tag** if the appliance is installed on a VLAN trunk and an untagged VLAN is unavailable.

Modifying Interface Configurations

Use the **Configuration > Interfaces** page if you want to change interface parameters such as.

- whether an interface is **admin up** or **down**
- **mgmt1** IP address
- whether or not an IP address is static, or dynamically assigned with DHCP
- speed and duplex
- MTU (Mean Transmission Unit) size
- MAC address

The screenshot shows the 'Interfaces' configuration page with a search bar and a table of network interfaces. The table has the following data:

Name	Admin	Status	IP Address/Mask	DHCP	Speed/Duplex	MTU	MAC	LAN Interface	WAN Interface	Hardened Inter...
mgmt0	up	up	10.0.238.135/26	<input checked="" type="checkbox"/>	auto/auto	1500	00:0C:29:0A:71:96	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
mgmt1	down	down	169.254.0.1/16	<input type="checkbox"/>	auto/auto	1500	Unassigned	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
wan0	up	up		<input type="checkbox"/>	auto/auto	1500	00:0C:29:0A:71:AA	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
lan0	up	up		<input type="checkbox"/>	auto/auto	1500	00:0C:29:0A:71:B4	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
wan1	down	down		<input type="checkbox"/>	auto/auto	1500	Unassigned	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
lan1	down	down		<input type="checkbox"/>	auto/auto	1500	Unassigned	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bv0	up	up	10.1.10.10/26	<input type="checkbox"/>		1500	00:0C:29:0A:71:AA	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Buttons: Apply, Cancel

 **WARNING** DHCP (Dynamic Host Configuration Protocol) can dynamically assign a new IP address to the appliance. **This may result in traffic loss because previously configured tunnel endpoints would now be incorrect.** If you elect to use DHCP, allocate the appliance’s IP address manually in the DHCP server. This prevents the possibility of lost traffic due to the DHCP server dynamically changing the IP address.

Overall, Silver Peak recommends statically assigning IP addresses.

Configuring Next-Hops

Use the **Configuration > Routes** page to configure **next-hops** for management, LAN, and WAN interfaces.

Routes ?

Management

Add new route

Show 10 Search

Subnet	Next-hop IP	Interface	
10.0.238.128/26	0.0.0.0	mgmt0	
0.0.0.0/0	10.0.238.129	mgmt0	X

Showing 1 to 2 of 2 entries First Previous 1 Next Last

WAN

Search

Next-hop IP	Interface
10.1.10.1	bvi0

LAN

Add new route Inter-VLAN Routing

Show 10 Search

Subnet	Next-hop IP	Interface	Metric
No Data Available			

Apply Cancel

Management Next-Hops

- Management routes specify the **default gateways** and local IP subnets for the management interfaces.
- In a Dual-Homed Router Mode configuration, you may need to add a static management route for flow redirection between appliances paired for redundancy at the same site.
- The management routes table shows the configured static routes and any dynamically created routes. If you use **DHCP**, then the appliance automatically creates appropriate dynamic routes. A user cannot delete or add dynamic routes.

WAN Next-Hops

- WAN next-hops provide next-hop addresses for optimized traffic.
- In an in-line deployment (bridge mode), the **wan0** interface displays as **bvi0**, for *bridge virtual interface*.
- When two WAN next-hops are configured Active/Active in 4-port bridge mode:
 - **lan0** ingress traffic is routed to the **wan0** next-hop.
 - **lan1** ingress traffic is routed to the **wan1** next-hop.
- When two WAN next-hops are configured Active/Active in Dual-Homed Router Mode:
 - **wan0** ingress traffic is routed to the **wan0** next-hop.
 - **lan0** ingress traffic is routed to the **lan0** next-hop.

LAN Next-Hops

- LAN routes provide next-hop addresses for traffic going to LAN-side networks that are not directly connected to an in-line (bridge mode) appliance.
- You can create redundant (backup) LAN routes by specifying another next-hop with a larger metric value.

For example, to specify 1.1.1.2 as a backup next-hop for 1.1.1.1, the table would contain:

- default 1.1.1.1 10
 - default 1.1.1.2 20
- Selecting **Inter-VLAN Routing** enables the appliance to route packets over another VLAN when the originally specified VLAN is unavailable.

Taking a Quick Look at the System Page

Odds are, you won't need to make any changes to the **Configuration > System** page. The two typical features to consider are the “**auto-tunnel**” and **subnet sharing**.

System Information for laine-vxa ?

General

Model VX-1000 206002001000 Rev 46839

Serial 000C291953A1

Contact

Location

Optimization

IP Id auto optimization

TCP auto optimization

Automatically establish tunnels

Flows and tunnel failure fail-stick ▾

Subnet Sharing

Use shared subnet information

Automatically include local subnets

Metric for local subnets

Network Memory

Encrypt data on disk

Configured Media Type ram and disk ▾

Media Type ram and disk

Exception Handling

Deciding whether to build tunnels automatically

Before deciding whether or not to use the **auto-tunnel feature**, consider the following:

- If you want the auto-tunnel feature to automatically build tunnels for you, it must be enabled for each appliance involved.
- This feature is useful when setting up a basic Proof of Concept.
- It is **not** recommended if you:
 - Have multiple IP addresses per appliance
 - Would be creating an excess of unnecessary tunnels, based on having a large volume of appliances
 - Need to configure parallel tunnels
 - Want a non-standard or more complex tunnel configuration — for example, configuring for **IPsec** or for **FEC** (Forward Error Correction).



Note Using this feature requires that any necessary outbound and inbound redirection is already configured. For more information, see “*Examining the Need for Traffic Redirection*” on page 8.

Manually Adding Local Subnets

Subnet sharing is a method for automatically routing a flow into the appropriate tunnel for optimization. Because peer appliances can advertise and share their subnet information, it reduces the need to create explicit route map entries to optimize traffic.

Subnets [?](#)

Use shared subnet information

Automatically include local subnets

Metric for automatically added subnets

All Configured Learned

Add new subnet

Show Search

Subnet/Mask	Metric	Is Local	Advertise to Peers	Exclude	Type	SaaS Application Name	Learned from Peer	Comment
2.2.2.0/26	50	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		10.1.11.10	
10.1.10.0/26	10	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Auto (added by system)			
10.1.11.0/26	10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		10.1.11.10	

Showing 1 to 3 of 3 entries

Apply Cancel

First Previous 1 Next Last

How is subnet sharing implemented?

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

The following are system-level choices:

- **Use shared subnet information**, which enables the feature on the appliance. If deselected, the subnet table is not used/available for auto-optimization.
- **Automatically include local subnets**, which adds the local subnet(s) for the appliance's interfaces to the subnet table.

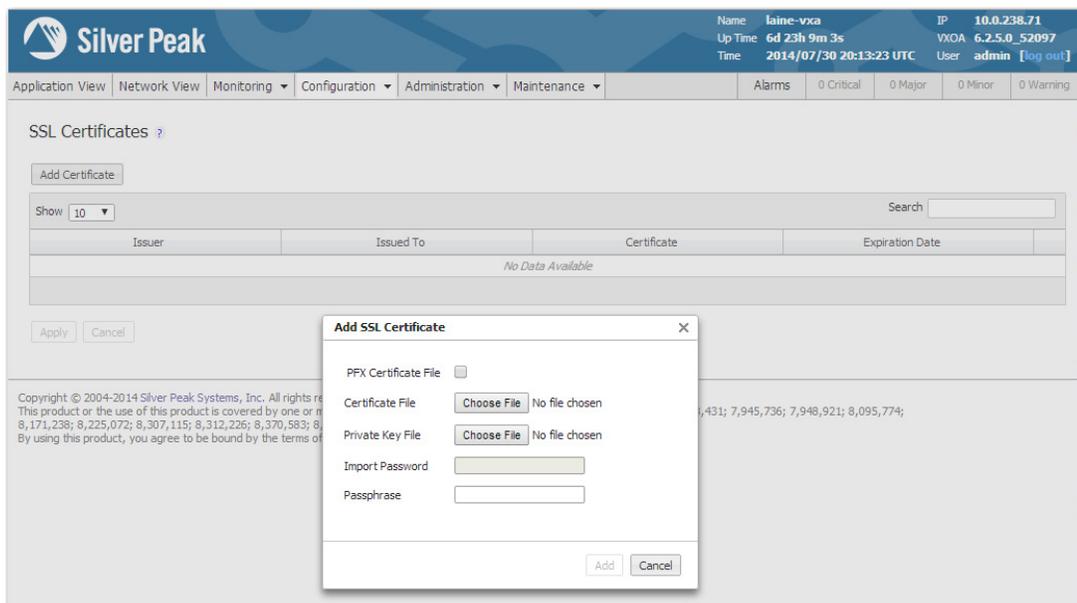
If deselected, the system doesn't create entries for the appliance's local subnets. If these subnets aren't listed, they can't be shared with peer appliances for auto-optimization.
- **Metric for automatically added subnets** = 10. This value can be between 0 and 100 and is the metric assigned to subnets of interfaces on this appliance.

If you're using subnet sharing, and you want to include local subnets that aren't tunnel endpoints, then you need to add them manually.

Adding SSL Certificates and Keys for Deduplication

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

- Silver Peak decrypts SSL data using the configured certificates and keys, optimizes the data, and transmits data over an IPsec tunnel. The peer Silver Peak appliance uses configured SSL certificates to re-encrypt data before transmitting.
- Peers that exchange and optimize SSL traffic must use the same certificate and key.
- Use this page to directly load the certificate and key into this appliance.
 - You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate.
 - The default is PEM when PFX Certificate File is deselected.
 - If the key file has an encrypted key, enter the passphrase needed to decrypt it.
- Silver Peak supports X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
- Silver Peak appliances support:
 - **Protocol versions:** SSLv3, SSLv3.3, TLS1.0, TLS1.1, TLS1.2
 - **Key exchanges:** RSA, DHE, ECDHE
 - **Authentication:** RSA
 - **Cipher algorithms:** RC4, 3DES, AES128, AES256, AES128-GCM, AES256-GCM
 - **Message Digests:** MD5, SHA, SHA256, SHA284



◆ **Before installing the certificates, you must do the following:**

- 1 Configure the tunnels bilaterally for IPsec mode.
To do so, access the **Configuration > Tunnels** page, select the tunnel, and for **Mode**, select **ipsec**.
- 2 Verify that TCP acceleration and SSL acceleration are enabled.
To do so, access the **Configuration > Optimization Policy** page, and review the **Set Actions**.



Configuring the System and Network

This chapter describes the tabs for configuring network and system-level parameters.

In This Chapter

- **System Information** See page 18.
- **Deployment** See page 20.
- **Interfaces** See page 29.
- **Routes** See page 30.
- **Tunnels** See page 32.
- **Shaper** See page 37.
- **Subnets** See page 41.
- **SSL Certificates** See page 43.
- **SSL CA Certificates** See page 45.
- **VRRP** See page 46.
- **WCCP** See page 48.
- **System Limits** See page 51.

System Information

Configuration > [System & Networking] System

Use the **System** page to configure appliance-level features and characteristics.

System Information for Seattle-EC

General

Model: VX-3000 205002006000 Rev 44261

Serial: 000C29475AFE

Contact:

Location:

Optimization

IP Id auto optimization:

TCP auto optimization:

Automatically establish tunnels:

Flows and tunnel failure:

Subnet Sharing

Use shared subnet information:

Automatically include local subnets:

Metric for local subnets:

Network Memory

Encrypt data on disk:

Configured Media Type:

Media Type: ram and disk

Excess Flow Handling

Excess flow policy:

Excess flow DSCP markings:

Miscellaneous

SSL optimization for non-IPSec tunnels:

Bridge Loop Test:

Enable SaaS optimization:

Enable IGMP snooping:

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).
- **Automatically establish tunnels** reduces configuration overhead by removing the need to manually create tunnels.

Subnet Sharing

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

Network Memory

- **Encrypt data on disk** enables encryption of all the cached data on the disks. Disabling this option is not recommended.
- **Configured Media Type** is either ram and disk (VX) or ram only (VRX). Can change for special circumstances, if recommended by Silver Peak.
- **Media Type** displays the actual media being used.

Excess Flow Handling

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

Miscellaneous

- **SSL optimization for non-IPSec tunnels** specifies if the appliance should perform SSL optimization when the outbound tunnel for SSL packets is not encrypted (for example, a GRE or UDP tunnel). To enable Network Memory for encrypted SSL-based applications, you must provision server certificates via the Silver Peak Unity Orchestrator. This activity can apply to the entire distributed network of Silver Peak appliances, or just to a specified group of appliances.
- **Bridge Loop Test** is only valid for virtual appliances. When enabled, the appliance can detect bridge loops. If it does detect a loop, the appliance stops forwarding traffic and raises an alarm. Appliance alarms include recommended actions.
- **Enable SaaS optimization** is a global flag that enables the SaaS optimization feature if you have registered for that cloud-based product on the License & Registration page.
- **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.

Deployment

Configuration > [System & Networking] Deployment

Information for this tab is portioned as follows:

- **Mapping Labels to Interfaces** See page 20.
- **LAN-side Configuration: DHCP** See page 20.
- **WAN-side Configuration** See page 21.
- **Basic Deployments** See page 22.
- **How You Can Adjust the Basic Deployments** See page 27.
- **Adding Data Interfaces** See page 28.
- **Definitions** See page 28.

Mapping Labels to Interfaces



Note In Orchestrator, you can create labels and map them to interfaces. Although you cannot create labels in the appliance, you can apply them if they exist.

- 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*.

LAN-side Configuration: DHCP

- By default, *each* LAN IP acts as a **DHCP Server** when the appliance is in (the default) Router mode.
- The other choices are **No DHCP** and having the appliance act as a **DHCP Relay**.

DHCP Server Definitions

- **DHCP Pool Subnet/Mask** is the full range of IP addresses that you make available for your network.
- **Subnet Mask** is a mask that specifies the default number of IP addresses reserved for any subnet. For example, entering **24** reserves 256 IP addresses.
- **Start Offset** specifies how many addresses not to allocate at the beginning of the subnet's range. For example, entering **10** means that the first ten IP addresses in the subnet aren't available.
- **End Offset** specifies how many IP addresses are not available at the end of the subnet's range.
- **Default lease** and **Maximum lease** specify, in hours, how long an interface can keep a DHCP-assigned IP address.
- **DNS server(s)** specifies the associated Domain Name System server(s).
- **NTP server(s)** specifies the associated Network Time Protocol server(s).
- **NetBIOS name server(s)** is used for Windows (SMB) type sharing and messaging. It resolves the names when you are mapping a drive or connecting to a printer.

- The **NetBIOS node type** of a networked computer relates to how it resolves NetBIOS names to IP addresses. There are four node types:
 - **B-node** = 0x01 Broadcast
 - **P-node** = 0x02 Peer (WINS only)
 - **M-node** = 0x04 Mixed (broadcast, then WINS)
 - **H-node** = 0x08 Hybrid (WINS, then broadcast)

DHCP Relay Definitions

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

WAN-side Configuration

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

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

NAT: If the appliance is behind a NAT-ed interface, select **NAT** (without the strikethrough). When using NAT, use in-line Router mode to ensure that addressing works properly. That means you configure paired single or dual WAN and LAN interfaces on the appliance.

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

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

EdgeConnect Licensing: Only visible on EC appliances

- By default, every EC has a max system bandwidth of 200 Mbps. For more bandwidth, you can purchase **Plus**, and then select it here for this profile.
- If you've purchased a reserve of **Boost** for your network, you can allocate a portion of it here.
- To view your settings for **Plus** and **Boost**, view the **Administration > License & Registration** tab.

Basic Deployments

This page 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.

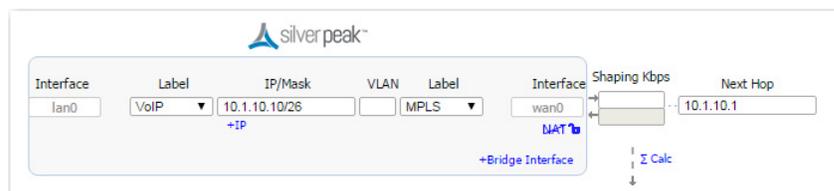
 **Note** In Orchestrator, you can create labels and map them to interfaces. Although you cannot create labels in the appliance, you can apply them if they exist.

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

Bridge Mode

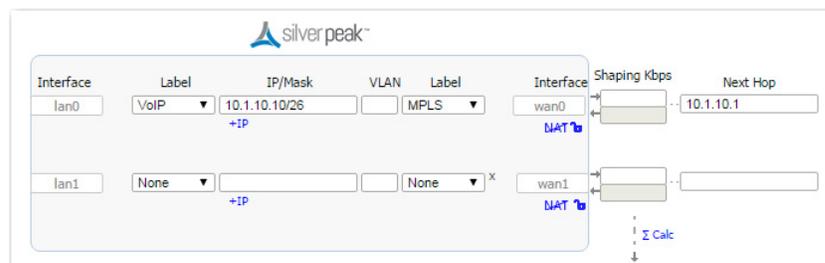
■ Single WAN-side Router

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



■ Dual WAN-side Routers

This is the most common 4-port bridge configuration.



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

■ Considerations for Bridge Mode Deployments

- Do you have a physical appliance or a virtual appliance?

A virtual appliance has no fail-to-wire, so you would need a redundant network path to maintain connectivity if the appliance fails.

- If your LAN destination is behind a router or L3 switch, you need to add a LAN-side route (a LAN next-hop).
- If the appliance is on a VLAN trunk, then you need to configure VLANs on the Silver Peak so that the appliance can tag traffic with the appropriate VLAN tag.

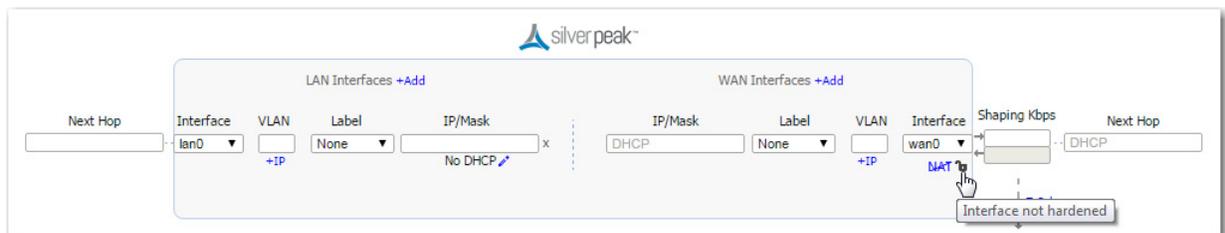
Router Mode

There are four options to consider:

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

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

■ #1 - Single LAN Interface & Single WAN Interface

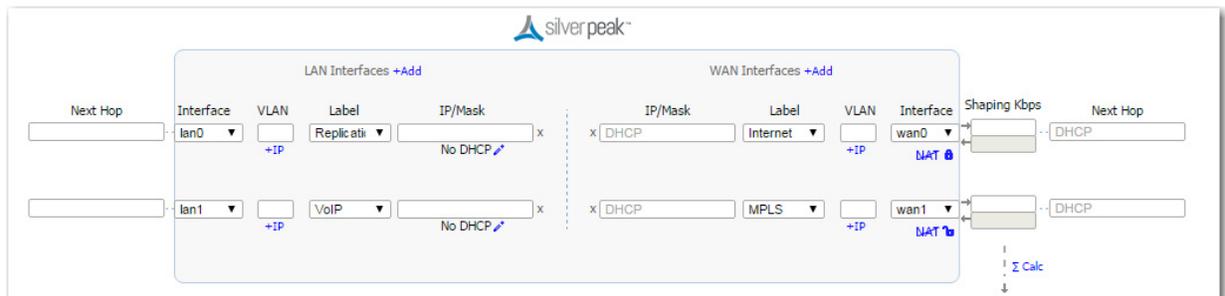


For this deployment, you have two options:

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

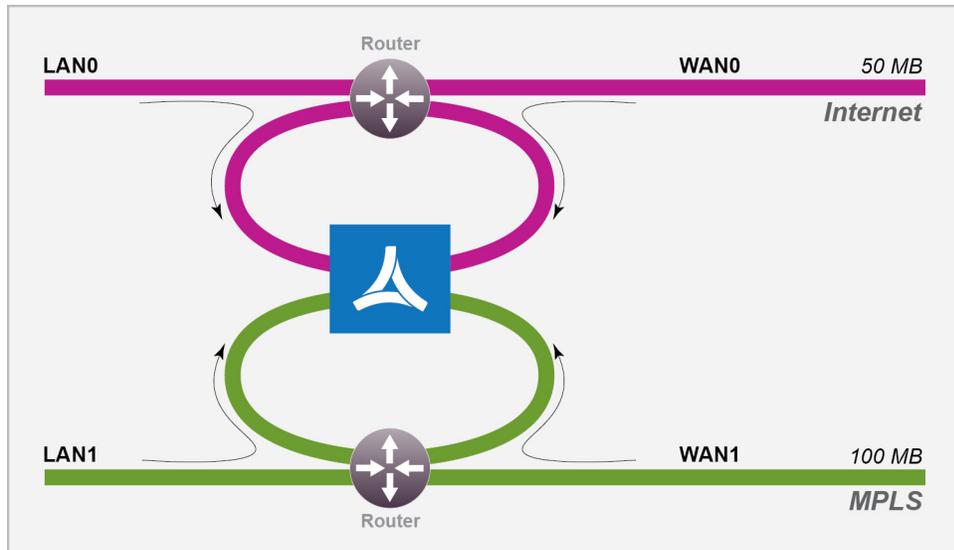
■ #2 - Dual LAN Interfaces & Dual WAN Interfaces



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

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

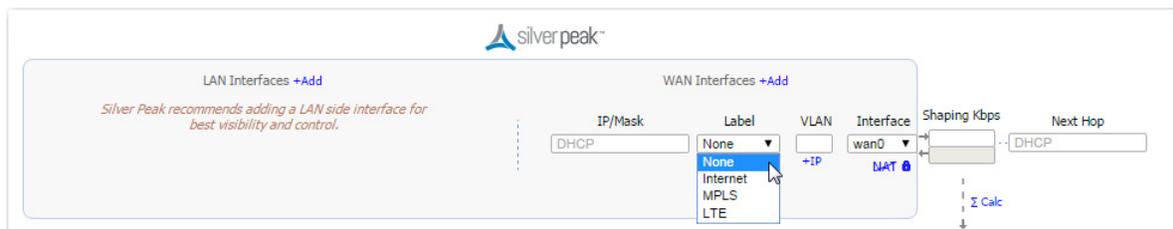
Out-of-path dual LAN and dual WAN interfaces



For this deployment, you have two options:

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

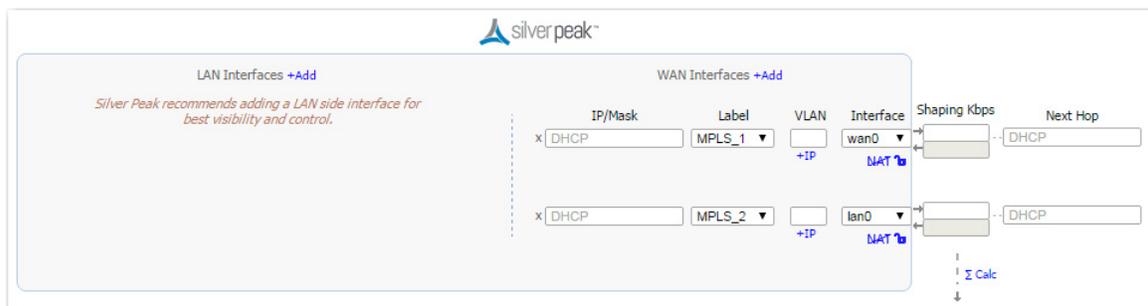
■ #3 - Single WAN Interface Sharing LAN and WAN traffic



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

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

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



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

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

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

■ Considerations for Router Mode Deployments

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

■ Examining the Need for Traffic Redirection

Whenever you place an appliance out-of-path, you must redirect traffic from the client to the appliance.

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

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

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

- **WCCP** (Web Cache Communication Protocol) — configured on both the router and the Silver Peak appliance. You can also use WCCP for redundancy and load balancing.

- **Host routing** — the server/end station has a default or subnet-based static route that points to the Silver Peak appliance as its next hop. Host routing is the preferred method when a virtual appliance is using a single interface, **mgmt0**, for datapath traffic (also known as Server Mode).
To ensure end-to-end connectivity in case of appliance failure, consider using VRRP between the appliance and a router, or the appliance and another redundant Silver Peak.

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

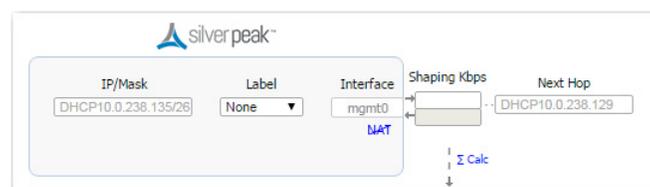
- If you use **subnet sharing** (which relies on advertising local subnets between Silver Peak appliances) or **route policies** (which specify destination IP addresses), then you only need LAN-side redirection.
- If, instead, you rely on **TCP-based** or **IP-based** auto-optimization (which relies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

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

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

Server Mode

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



How You Can Adjust the Basic Deployments

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



Bonding

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

For more information, see *“Configuring Gigabit Etherchannel Bonding” on page 27*.

Use Fiber Ports

Choose this when you want to enable fiber ports (10Gbps/1Gbps) on a physical appliance that also has 1Gbps copper ports.

Propagate Link Down

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

4-port single bridge

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



Note Changing the deployment mode requires a reboot.

Configuring Gigabit Etherchannel Bonding

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

When you configure bonding, the following is true:

- **lan0** plus **lan1** bond to form **blan0**, which uses the **lan0** IP address.
- **wan0** plus **wan1** bond to form **bwan0**, which uses the **wan0** IP address.
- The appliances use flow-based load balancing across the links.
- This configuration provides failover in case one link goes down.

- You can view the statistics on the **Monitoring - Interfaces** page. If you're using bonding, you'll see statistics for **blan0** and **bwan0**, as well as for the interfaces that comprise them (**lan0**, **lan1**, **wan0**, and **wan1**).
 - If a WCCP or VRRP deployment already exists, then you must reconfigure the deployment on the bonding interface. In other words, if you previously configured on **wan0**, then after bonding you must reconfigure on **bwan0**.
 - Rollback to non-bonding mode returns the intact, non-bonded configuration.
 - Enabling/disabling bonding requires an appliance reboot.
- ◆ **To configure etherchannel bonding**

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

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

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

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

Adding Data Interfaces

- You can create additional data-plane Layer 3 interfaces, to use as tunnel endpoints.
- To add a new interface, click **+IP**.
- Add a **VLAN tag** if the appliance is installed on a VLAN trunk and an untagged VLAN is unavailable.

Definitions

LAN Side Next-hop(s): Provide next-hop address(es) for LAN-side networks that are not directly connected to an in-line (bridge mode) appliance. Redundant (backup) LAN Next-hop(s) can be created by the second (lan1) next-hop.

Interfaces

Configuration > [System & Networking] Interfaces

Use this page to **review** and **edit** the appliance interfaces.

The screenshot shows the 'Interfaces' configuration page with a table listing 7 rows of interface information. The table has columns for Name, Status, IP Address/Mask, Public IP, Speed/Duplex, MTU, and MAC. Below the table are 'Apply' and 'Cancel' buttons.

Name ▲	Status	IP Address/Mask	Public IP	Speed/Duplex	MTU	MAC
mgmt0	up	10.0.238.71/26	10.0.238.71	auto/auto	1500	00:0C:29:19:53:A1
mgmt1	down	169.254.0.1/16		auto/auto	1500	00:0C:29:19:53:AB
wan0	up			auto/auto	1500	00:0C:29:19:53:B5
lan0	up	0.0.0.0/0		auto/auto	1500	00:0C:29:19:53:BF
wan1	down			auto/auto	1500	Unassigned
lan1	down			auto/auto	1500	Unassigned
bvi0	up	10.1.153.20/24			1500	00:0C:29:19:53:B5

- 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.
- A **Hardened Interface** provides security specifically on WAN-side interface(s):
 - For traffic inbound from the WAN, the appliance accepts **only** IPSec tunnel packets. It drops all other pass-through traffic and tunnel packets.
 - For traffic outbound from the WAN, the appliance allows IPSec tunnel packets and outgoing pass-through traffic. It drops all other tunnel packets.



WARNING DHCP (Dynamic Host Configuration Protocol) can dynamically assign a new IP address to the appliance. **This may result in traffic loss because previously configured tunnel endpoints would now be incorrect.** If you elect to use DHCP, allocate the appliance's IP address manually in the DHCP server. This prevents the possibility of lost traffic due to the DHCP server dynamically changing the IP address.

Overall, Silver Peak recommends statically assigning IP addresses.

Terminology

blan: Bonded **lan** interfaces (as in **lan0** + **lan1**).

bvi0: Bridge Virtual Interface. When the appliance is deployed in-line (Bridge mode), it's the routed interface that represents the bridging of **wan0** and **lan0**.

bwan: Bonded **wan** interfaces (as in **wan0** + **wan1**).

tlan: 10-Gbps fiber **lan** interface.

twan: 10-Gbps fiber **wan** interface.

Routes

Configuration > [System & Networking] Routes

Use this tab to configure **next-hops** for management, LAN, and WAN interfaces.

Routes ?

Management

Add new route

Show 10 Search

Subnet	Next-hop IP	Interface	
10.0.238.128/26	0.0.0.0	mgmt0	
0.0.0.0/0	10.0.238.129	mgmt0	X

Showing 1 to 2 of 2 entries First Previous 1 Next Last

WAN

Add new route Search

Next-hop IP	Interface
10.1.10.1	bvi0

LAN

Add new route Inter-VLAN Routing

Show 10 Search

Subnet	Next-hop IP	Interface	Metric
No Data Available			

Apply Cancel

Management

- Management routes specify the **default gateways** and local IP subnets for the management interfaces.
- In a Dual-Homed Router Mode configuration, you may need to add a static management route for flow redirection between appliances paired for redundancy at the same site.
- The management routes table shows the configured static routes and any dynamically created routes. If you use **DHCP**, then the appliance automatically creates appropriate dynamic routes. A user cannot delete or add dynamic routes.

WAN

- WAN next-hops provide next-hop addresses for optimized traffic.
- In an in-line deployment (bridge mode), the **wan0** interface displays as **bvi0**, for bridge virtual interface.
- When two WAN next-hops are configured Active/Active in 4-port bridge mode:
 - lan0** ingress traffic is routed to the **wan0** next-hop.
 - lan1** ingress traffic is routed to the **wan1** next-hop.
- When two WAN next-hops are configured Active/Active in Dual-Homed Router Mode:
 - wan0** ingress traffic is routed to the **wan0** next-hop.
 - lan0** ingress traffic is routed to the **lan0** next-hop.

LAN

- LAN routes provide next-hop addresses for traffic going to LAN-side networks that are not directly connected to an in-line (bridge mode) appliance.
- You can create redundant (backup) LAN routes by specifying another next-hop with a larger metric value.

For example, to specify 1.1.1.2 as a backup next-hop for 1.1.1.1, the table would contain:

- default 1.1.1.1 10
 - default 1.1.1.2 20
- Selecting **Inter-VLAN Routing** enables the appliance to route packets over another VLAN when the originally specified VLAN is unavailable.
 - If you're in **dual bridge mode**, then you can manually choose an interface (**bvi0** or **bvi1**) for the route. Otherwise, the appliance chooses.

Tunnels

Configuration > [System & Networking] Tunnels

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

Tunnels Monitoring

Automatically establish tunnels

Use shared subnet information

Refresh
6 mins ago

Add Tunnel Rediscover MTU Default Tunnel Settings

1 Rows, 1 Selected Search

Name	Status	Admin State	Auto MTU	MTU (700..9000)	Local IP	Remote IP	Auto Max BW	Max BW Kbps (1..3600)	Mode	Uptime	Advanced Options	Alert Options
auto_tun_10.1.153.20_to_10...	up - active	up	<input checked="" type="checkbox"/>	1500	10.1.153.20	10.1.154.20	<input checked="" type="checkbox"/>		Auto udp	24d 23h 39m 37s		

Apply Cancel

- To create a tunnel, click **Add Tunnel** and edit within the new row.
- To change the default values for future tunnels, click **Default Tunnel Settings** and edit the values.
- If you've lost network connectivity, one action to try is **Rediscover MTU**.
- You can configure *tunnel-specific* **Threshold Crossing Alerts (TCAs)** in the **Alert Options**.

Tunnel Alert Options

Name	Rising				Falling			
	Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
latency	1000 ms	850 ms	1	<input checked="" type="checkbox"/>	0 ms	0 ms	1	<input type="checkbox"/>
loss post-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
loss pre-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
OOP post-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
OOP pre-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
reduction	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
utilization	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>

Reset to Defaults OK Cancel

TCAs are pre-emptive, user-configurable thresholds that declare a Major alarm when crossed. These would be exceptions to global TCAs configured on the **Configuration > Threshold Crossing Alerts** page.

- You cannot edit a Local IP or Remote IP on an existing tunnel.

Options

- **Automatically establish tunnels** reduces configuration overhead by removing the need to manually create tunnels.
- **Used shared subnet information** enables the appliance to use information from its subnet table, which it builds from entries added automatically by the system or manually by a user. When two appliances are connected by a tunnel, they exchange this information ("learn" it) and use it to route traffic to each other. If this option is deselected, the subnet table is not used for auto-optimization.

Definitions (alphabetically)

- **Admin State** allows you to admin **Up** (or admin **Down**) a tunnel.
- **Auto Discover MTU Enabled** allows the tunnel MTU to be discovered automatically. When selected, this overrides the MTU setting.
- **Auto Max BW Enabled** allows the appliances to negotiate the maximum tunnel bandwidth based upon the lower of the two system bandwidths of the two appliances.
- **Local IP** is a local address on the appliance.
- **Max BW Kbps** is the maximum bandwidth for this tunnel, in kilobits per second. This must be equal to or less than the upstream bandwidth of your WAN connection.
- **Min BW Kbps** is the minimum bandwidth for this tunnel, in kilobits per second.
- **Mode** indicates whether the tunnel protocol is **udp**, **gre**, or **ipsec**. The default is **udp**.
- **MTU (700..9000) Bytes.** (Maximum Transmission Unit). is the maximum tunnel packet size including its payload and Layer-3 header. By default, MTU is automatically discovered because **Auto Discover MTU** is enabled. When setting this value manually, set it to the largest value that won't result in tunnel packets being fragmented by networking equipment in the WAN.
- **Name** is a unique string identifying this tunnel.
- **Remote IP** is the IP address for the remote appliance.
- **Status** indications are as follows:
 - **Down** = The tunnel is down. This can be because the tunnel administrative setting is down, or the tunnel can't communicate with the appliance at the other end. Possible causes are:
 - Lack of end-to-end connectivity / routability (test with iperf)
 - Intermediate firewall is dropping the packets (open the firewall)
 - Intermediate QoS policy (be packets are being starved. Change control packet DSCP marking) Mismatched tunnel mode (udp / gre / ipsec)
 - IPsec is misconfigured: (1) enabled on one side (see show int tunnel configured), or (2) mismatched pre-shared key
 - **Down - In progress** = The tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel.
 - **Down - Misconfigured** = The two appliances are configured with the same System ID. (see show system)
 - **Up - Active** = The tunnel is up and active. Traffic destined for this tunnel will be forwarded to the remote appliance.
 - **Up - Active - Idle** = The tunnel is up and active but hasn't had recent activity in the past five minutes, and has slowed the rate of issuing keep-alive packets.
 - **Up - Reduced Functionality** = The tunnel is up and active, but the two endpoint appliances are running mismatched software releases that give no performance benefit.
 - **UNKNOWN** = The tunnel status is unknown. This can be because the appliance is unable to retrieve the current tunnel status. Try again later.
- **Uptime** is how long since the tunnel came up.

Advanced Tunnel Options

You may use these options as needed.

General

- **IPSec Pre-shared Key** = a shared, secret string of Unicode characters that is used for authentication of an IPSec connection between two parties. If you select Default, the appliance makes the key; if you select Custom (recommended), the user specifies the key.
- **IPSec Anti-replay window** = IP security (IPsec) authentication provides anti-replay 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. Increase this value for networks with a lot of jitter (out-of-order packets).
- **UDP destination port** = Tunnel traffic will be transmitted in a UDP protocol packet using this destination port address. Only valid when the tunnel mode is set to UDP.
- **UDP flows** = the number of flows over which to spread tunnel traffic.

Packet

- **Coalescing Enabled** = whether or not to coalesce smaller packets into larger packets. Default = ON. Packet coalescing is particularly beneficial for web applications, VoIP, and interactive applications, like Citrix.
- **Coalescing Wait (ms)** = determines how long the appliance should hold packets while attempting to coalesce smaller packets into larger packets. Default = 0.

- **Reorder Wait (0..500 ms)** = the maximum time the appliance holds an out-of-order packet when attempting to reorder. The 100ms default value should be adequate for most situations. FEC may introduce out-of-order packets if the reorder wait time is not set high enough.
- **FEC** (Forward Error Correction) reconstructs lost packets (as reported by the remote appliance). The options are **disable**, **enable**, and **auto**.
 - When set to **enable**, FEC reconstructs lost tunnel packets at the destination appliance. FEC achieves this by injecting redundant (called parity) packets in the tunnel traffic. The specified FEC ratio determines the number of parity packets relative to data packets (for example, at 1:5 ratio, a parity packet is added for every 5 data packets).
 - When set to **auto**, it adjusts dynamically based on network conditions, with the upper limit being capped by the **FEC Ratio** value you choose.
- **FEC Ratio** is the ratio of parity packets relative to data packets (for example, at 1:5 ratio, a parity packet is added for every 5 data packets). The selectable values include **disable**, **auto**, **1:2**, **1:5**, **1:10**, and **1:20**. A FEC Ratio of 1:2 is very aggressive and should only be utilized with great care in networks with extremely high loss (10% or greater).

Tunnel Health

- **Retry Count** = the number of failed keep-alive messages that are allowed before the appliance brings the tunnel down. Keep-alive packets are sent once per second. Default = 30.
- **DSCP** = the DSCP value for the tunnel control packets.

FastFail Thresholds

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

- **Fastfail Enabled** – This option is triggered when a tunnel's keepalive signal doesn't receive a reply. The options are **disable**, **enable**, and **continuous**. If the disqualified tunnel subsequently receives a keepalive reply, its recovery is instantaneous.
 - If set to **disable**, keepalives are sent every second, and 30 seconds elapse before failover. In that time, all transmitted data is lost.
 - If set to **enable**, keepalives are sent every second, and a missed reply increases the rate at which keepalives are sent from 1 per second to 10 per second. Failover occurs after 1 second.
 - When set to **continuous**, keepalives are continuously sent at 10 per second. Therefore, failover occurs after one tenth of a second.
- Thresholds for **Latency**, **Loss**, or **Jitter** are checked once every second.
 - Receiving 3 successive measurements in a row that exceed the threshold puts the tunnel into a brownout situation and flows will attempt to fail over to another tunnel within the next 100mS.
 - Receiving 3 successive measurements in a row that drop below the threshold will drop the tunnel out of brownout.

Jumbo Frames and MTU Interworking

Silver Peak supports MTUs (Maximum Transmission Units) up to 9000 bytes. Because of pps (packets per second) limits on the LAN side, using 9000-byte MTUs can significantly improve LAN-side throughput for applications such as storage replication.

More importantly, the appliances support *interworking*. You can configure 9000-byte MTUs on storage arrays even if the replication protocol is running over a WAN with standard (1500-byte) MTUs. This is important because not all service providers allow for jumbo frames on the WAN.

Efficient MTU interworking scenarios include the following:

Local Interface MTU (bytes) [Configuration > Interfaces]	Tunnel MTU (bytes) [Configuration > Tunnels]	Remote Interface MTU (bytes) [Configuration > Interfaces]
1500	1500	1500
9000	9000	9000
9000	1500	9000
1500	9000	1500
9000	1500	1500



CAUTION Across all network devices, you must configure all interfaces on the same subnet to have the same MTU.

- For the Interface MTU, you must configure **each pair** of LAN and WAN interfaces on the appliance to have the same MTU value. For example, you'd configure both LAN0 and WAN0 to have a value of 9000 MTU. These are accessible via the **Configuration - Interfaces** page.
- If either end host has an MTU of 9000 and the tunnel MTU is 1500, then you need to disable the **Adjust MSS to Tunnel MTU** feature on both appliances. This prevents appliances from lowering higher MSS values negotiated by end stations to match the lower MTU of the tunnel. To find this feature, go to the **Optimization Policy**, and in the **TCP Accel Details** column, click the icon to open the **Advanced TCP Options**.

Shaper

Configuration > [System & Networking] Shaper

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

Shaper QoS Policies

Inbound **Outbound** Shaper wan Add Shaper Delete Shaper Enable Shaper

WAN Traffic Classes Max Bandwidth: 3600 (1.. 4000) Kbps - *IMPORTANT: Do not exceed WAN link bandwidth* Restore Defaults Bandwidth in: Kbps - Mbps

10 Rows Search

ID	Traffic Name	Priority	Min Bandwidth		Excess Weighting	Max Bandwidth		Max Wait Time (ms)
			%	kbps		%	kbps	
1	default	5	30	1080	100	100	3600	500
2	real-time	1	30	1080	1000	100	3600	100
3	interactive	2	20	720	1000	100	3600	200
4	best-effort	8	20	720	100	100	3600	500
5		5	30	1080	100	100	3600	500
6		5	30	1080	100	100	3600	500
7		5	30	1080	100	100	3600	500
8		5	30	1080	100	100	3600	500
9		5	30	1080	100	100	3600	500
10		5	30	1080	100	100	3600	500

Pass-through Shaped Traffic
Max Bandwidth 3600 (0.. 4000) Kbps

Apply Cancel

Shaper QoS Policies

Inbound Outbound Shaper Add Shaper Delete Shaper

Click Add Shaper button to add inbound shapers

Dynamic Rate Control
Enable Dynamic Rate Control
Inbound Bandwidth Limit (1.. 2000000) Kbps

Pass-through Shaped Traffic
Max Bandwidth 3600 (0.. 4000) Kbps

Apply Cancel

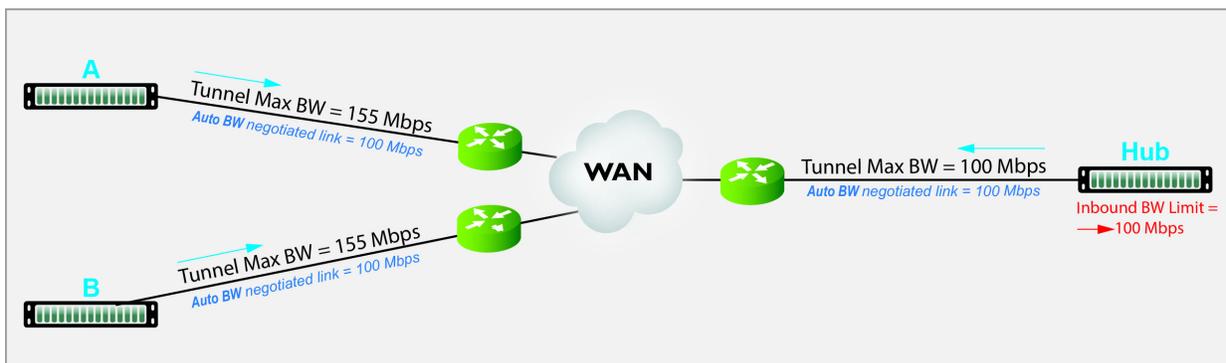
- It shapes traffic by allocating bandwidth as a percentage of the **system bandwidth**.
- The Shaper's parameters are organized into ten traffic classes. Four traffic classes are preconfigured and named --- **real-time**, **interactive**, **default**, and **best effort**.
- The system applies these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

- The default shaper is named, **wan**, and it shapes all outbound traffic on the WAN link.
 - If there are multiple WAN links, they can all use the default shaper, or you can choose to create additional shapers. A WAN link is usually associated with an interface on your appliance.
 - When you click **Add Shaper**, your selected deployment determines what interfaces you can choose. For example, with Dual Home Router Mode, tunnel traffic that uses **wan0** is routed to one WAN link, and **lan0** to another WAN link. So, you can add customized shapers for the **wan0** and **lan0** interfaces.
 - A shaper for a specific interface outranks the generic shaper. For example, if you create a shaper for **wan0**, then the appliance uses that instead of the more generic **wan** shaper.
 - If you have two WAN interfaces and you create a shaper for **wan0**, then the appliance uses the **wan0** shaper for the **wan0** interface, and the generic **wan** shaper for the **wan1** interface.
- You can rename or edit any traffic class.

Dynamic Rate Control

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

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



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

If **Hub** experiences congestion:

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

Definitions

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

If you set **Min Bandwidth** to a value greater than **Max Bandwidth**, then **Max** overrides **Min**.
- **Excess Weighting:** If there is bandwidth left over after satisfying the minimum bandwidth percentages, then the excess is distributed among the traffic classes, in proportion to the weightings specified in the **Excess Weighting** column. Values range from 1 to 10,000.

- **Max Bandwidth:** You can limit the maximum bandwidth that a traffic class uses by specifying a percentage in the **Max Bandwidth** column. The bandwidth usage for the traffic class will never exceed this value.
- **Max Wait Time:** Any packets waiting longer than the specified **Max Wait Time** are dropped.

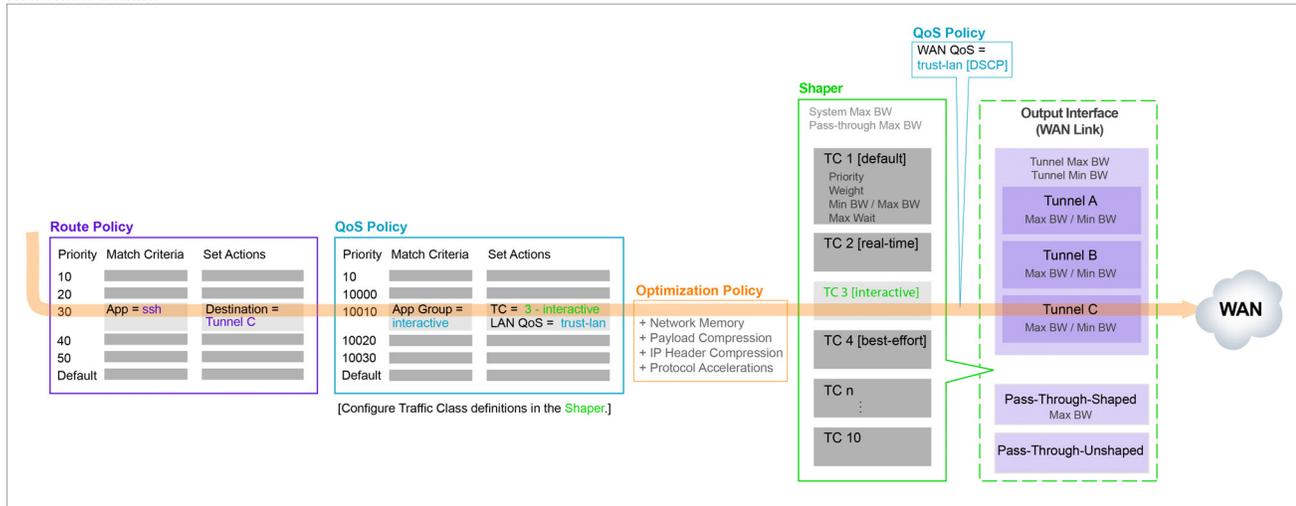
How Flows are Shaped

The following diagram illustrates how flows are shaped when the Route Policy Set Action, **Destination**, is:

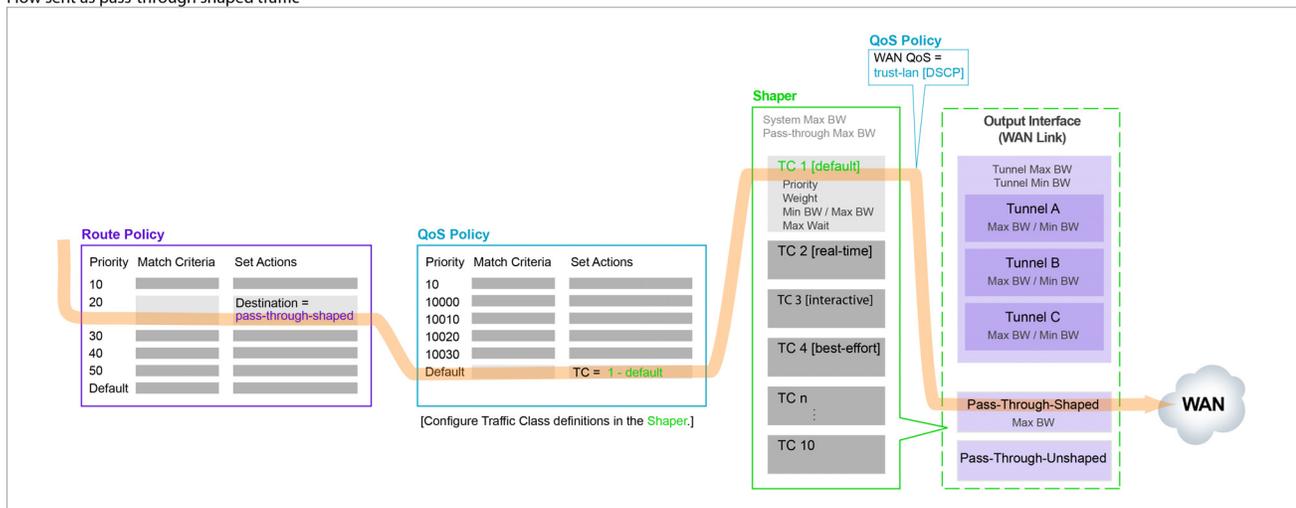
- a specific tunnel
- pass-through shaped
- pass-through unshaped

If the Route Policy's Set Action is *auto-optimized* and the local appliance initiates either TCP-based or IP-based handshaking, then the remote appliance determines which tunnel to use, based on information it receives in the first packets from the local appliance.

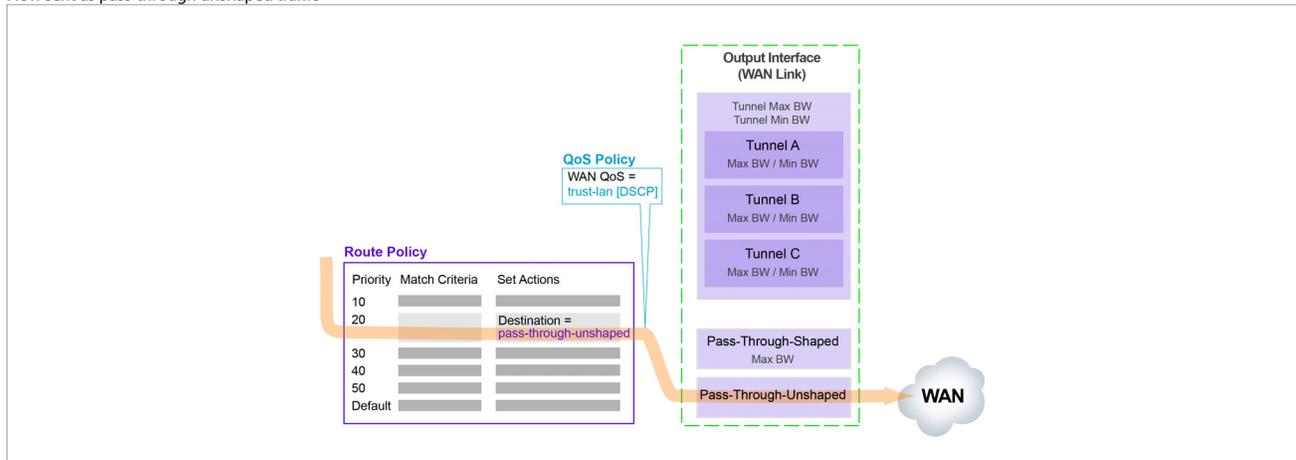
Flow sent to a tunnel



Flow sent as pass-through shaped traffic



Flow sent as pass-through unshaped traffic



Subnets

Configuration > [System & Networking] Subnets

Subnet sharing is a method for automatically routing a flow into the appropriate tunnel for optimization. Because peer appliances can advertise and share their subnet information, it reduces the need to create explicit route map entries to optimize traffic.

Subnets ⓘ

Use shared subnet information

Automatically include local subnets

Metric for automatically added subnets

All Configured Learned

Add new subnet

Show 25 ▾ Search

Subnet/Mask ▾	Metric	Is Local	Advertise to Peers	Exclude	Type	SaaS Application Name	Learned from Peer	Comment
2.2.2.0/26	50	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		10.1.11.10	
10.1.10.0/26	10	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Auto (added by system)			
10.1.11.0/26	10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Learned from peer		10.1.11.10	

Showing 1 to 3 of 3 entries

Apply Cancel

How is subnet sharing implemented?

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

The following are system-level choices:

- **Use shared subnet information**, which enables the feature on the appliance. If deselected, the subnet table is not used/available for auto-optimization.
- **Automatically include local subnets**, which adds the local subnet(s) for the appliance's interfaces to the subnet table.
If deselected, the system doesn't create entries for the appliance's local subnets. If these subnets aren't listed, they can't be shared with peer appliances for auto-optimization.
- **Metric for automatically added subnets** = 10. This value can be between 0 and 100 and is the metric assigned to subnets of interfaces on this appliance.

When would you need to create a Route Policy entry?

Subnet sharing takes care of optimizing IP traffic. Use Route Policy 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, ports, VLAN, DSCP, or ACL (Access Control List)

Subnet table columns

- **Subnet/Mask:** Actual subnet to be shared or learned
- **Metric:** Value must be between 0 and 100. When an appliance finds more than one peer appliance advertising the longest matching subnet (for example, in a high availability deployment), it chooses the peer that advertises the subnet with the lowest metric value - that is, **lower metric values have priority**.
- **Is Local:** Specifies if the subnet is local to this site.

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

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

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

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

- **Type of subnet:**
 - **Auto (added by system)** = automatically added subnets of interfaces on this appliance
 - **Auto (added by saas optimization)** = automatically added subnets from SaaS services
 - **Added by user** = manually added/configured subnets for this appliance
 - **Learned from peer** = subnets added as a result of exchanging information with peer appliances
- **SaaS Application Name:** If the subnet is associated with a SaaS service, the name displays here.
- **Learned from Peer:** Which peer appliance advertised (and share

SSL Certificates

Configuration > [System & Networking] SSL Certificates

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 For SSL decryption of SaaS services, use the **Configuration > SaaS Optimization** page. Because SaaS servers are external to your enterprise network, the appliance creates a *substitute* certificate, which must then be signed by a Certificate Authority (CA).

The screenshot displays the 'SSL Certificates' management interface. At the top, there is an 'Add Certificate' button and a search bar. Below is a table with the following data:

Issuer	Issued To	Certificate	Expiration Date
test1	test1	View c²	Dec 3 00:36:12 2015 GMT
test2	test2	View c²	Dec 3 00:37:15 2015 GMT

Below the table, it indicates 'Showing 1 to 2 of 2 entries' and provides navigation buttons: First, Previous, 1, Next, Last. A modal window titled 'Add SSL Certificate' is overlaid on the page, containing the following fields:

- PFX Certificate File:
- Certificate File: No file chosen
- Private Key File: No file chosen
- Import Password:
- Passphrase:

Buttons for 'Add' and 'Cancel' are located at the bottom right of the modal.

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

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

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

 - Before installing the certificates, you must do the following:
 - Configure the tunnels bilaterally for **IPSec** mode.
To do so, access the **Configuration - Tunnels** page, select the tunnel, and for **Mode**, select **ipsec**.
 - Verify that **TCP acceleration** and **SSL acceleration** are enabled.
To do so, access the **Configuration - Optimization Policy** page, and review the **Set Actions**.



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

SSL CA Certificates

Configuration > [System & Networking] SSL CA Certificates

If the enterprise CA certificate that you use for signing substitute certificates is subordinate to higher level **Certificate Authorities (CA)**, then you must add those CA certificates here.

Those same CA certificates must also be present in the browser. If the browser can't validate up the chain to the root CA, it will warn you that it can't trust the certificate.

SSL CA Certificates ? SaaS Optimization

Add CA Certificate

Show 10 Search

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

Showing 1 to 1 of 1 entries

First Previous 1 Next Last

Apply Cancel

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

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



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

VRRP

Configuration > [System & Networking] VRRP

Use this page to configure the appliance for **Virtual Router Redundancy Protocol (VRRP)**.

The screenshot shows the VRRP configuration interface. At the top, there is a header 'VRRP' with a help icon and a 'Monitoring' link. Below the header is an 'Add VRRP' button. A table displays the current VRRP configuration:

Group ID	Interface	State	Admin	Virtual IP	Advertisement Timer	Priority	Preemption	Authentication String	Description	Details
	wan0	Up			1	128	<input checked="" type="checkbox"/>			

At the bottom of the table are 'Apply' and 'Cancel' buttons.

In an out-of-path deployment, one method for redirecting traffic to the Silver Peak appliance is to configure VRRP on a common virtual interface. The possible scenarios are:

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

For step-by-step procedures for configuring common VRRP deployments, refer to the *Silver Peak Network Deployment Guide*.

CONFIGURATION PARAMETERS (alphabetically)

- **Admin** = The options are **up** (enable) and **down** (disable).
- **Advertisement Timer** = default is 1 second.
- **Group ID** is an identifier assigned to the two peers. Depending on the deployment, the group can consist of an appliance and a router (or L3 switch), or two appliances. The valid range is **1 - 255**.
- **Interface** refers to the interface that VRRP is using for peering.
- **Preemption**. Leave this selected/enabled so that after a failure, the appliance with the highest priority comes back online and again assumes primary responsibility.
- **Priority**. The greater the number, the higher the priority. The appliance with the higher priority is the VRRP Master.
- **State** = There are three options for the VRRP instance:
 - **Backup** = Instance is in VRRP backup state.
 - **Init** = Instance is initializing, it's disabled, or the interface is down.
 - **Master** = Instance is the current VRRP master.
- **Virtual IP**. The IP address of the VRRP instance. VRRP instances may run between two or more appliances, or an appliance and a router.

DETAILS (alphabetically)

- **IP Address Owner** = A Silver Peak appliance cannot use one of its own IP addresses as the VRRP IP, so this will always be **No**.
- **Master IP** = Current VRRP Master's Interface or local IP address.
- **Master State Transitions** = Number of times the VRRP instance went from Master to Backup and vice versa. A high number of transitions indicates a problematic VRRP configuration or environment. If this is the case, check the configuration of all local appliances and routers, and review the log files.
- **State Uptime** = Time elapsed since the VRRP instance entered the state it's in.
- **Virtual MAC address** = MAC Address that the VRRP instance is using. On an NX appliance, this is in 00-00-5E-00-01-**{VRID}** format. On virtual appliances, the VRRP instance uses the interface's assigned MAC Address (for example, the MAC address that the hypervisor assigned to **wan0**).

WCCP

Configuration > [System & Networking] WCCP

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

WCCP

Enable WCCP

Total Entries: 1

< 1 min ago

Group ID	Oper Status	Admin	Router IP	Protocol	Interface	Compatibility	Forwarding Method	Advanced Settings
51	N/A	up	0.0.0.0	tcp	wan0	ios	either	<input type="button" value="X"/>

Web Cache Communications Protocol (WCCP) supports the redirection of any TCP or UDP connections to appliances participating in WCCP Service Groups. The appliance intercepts only those packets that have been redirected to it. The appliance optimizes traffic flows that the Route Policy tunnelizes. The appliance forwards all other traffic as pass-through or pass-through-unshaped, as per the Route Policy.

- For the Service Groups to be active, you must select **Enable WCCP**. Otherwise, the service groups are configured, but not in service.
- The appliance should always be connected to an interface/VLAN that does not have redirection enabled — preferably a separate interface/VLAN would be provided for the appliance.
- If the appliance uses *auto-optimization*, then WCCP redirection must also be applied on the uplinks of the router or L3 switch to the core/WAN.
- Refer to the *Silver Peak Network Deployment Guide* for examples, best practices, and deployment tips.

Definitions (alphabetically)

- **Admin** values are **up** and **down**. The default is **up**.
- **Compatibility Mode**. Select the option appropriate for your router. If a WCCP group is peering with a router running **Nexus OS**, then the appliance must adjust its WCCP protocol packets to be compatible. By default, the appliance is **IOS-compatible**.
- **Forwarding Method**, also known as the *Redirect Method*. Packet redirection is the process of forwarding packets from the router or L3 switch to the appliance. The router or L3 switch intercepts the packet and forwards it to the appliance for optimization. The two methods of redirecting packets are **Generic Route Encapsulation (GRE)** and **L2 redirection**.

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

Service Group Advanced Settings

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

This ensures that a connection will go through the same appliance in both inbound and outbound directions and avoid asymmetry.
 - **custom** provides granular control of the distribution of flows. Contact Silver Peak Technical Support for assistance.

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

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

System Limits

Configuration > [System & Networking] System Limits

You can use this page to **design a template** to create a custom Silver Peak virtual appliance.

System Limits ?

Total memory required **3,954** MB

System Bandwidth 40,000 kbps

Total flows 128,000 flows (978 MB)

CIFS v1 proxy flows OFF (0MB)

SSL proxy flows OFF (0MB)

Network memory (MB) 1,440MB

Permanently lock dynamic system limits settings after reboot



Note This page is specific to service providers subscribing to Silver Peak's CPX services.

This feature requires registration with an **Account Key** on the **License & Registration** page.

The basic steps are as follows:

- 1 Move the sliders to accommodate your planned usage.
Based on your inputs, the value for total memory adjusts to show what the custom appliance will require.
- 2 Select **Permanently lock dynamic system limits settings after reboot**, and click **Apply**.
This finalizes values for the new, custom OVA file.
- 3 Use your hypervisor to export the template, creating a new golden image. You can use this image to create any number of new, identically configured virtual appliances. Each instance requires a License Key.



Configuring Policies

This chapter describes the tabs for configuring appliance policies.

In This Chapter

- **Route Policies** See page 54.
- **QoS Policies** See page 56.
- **Optimization Policies** See page 60.
- **Access Lists** See page 62.
- **Built-in Applications** See page 63.
- **User-Defined Applications** See page 64.
- **Application Groups** See page 66.
- **Flow Redirection** See page 67.
- **SaaS Optimization** See page 68.
- **NAT Policies** See page 71.
- **Threshold Crossing Alerts** See page 74.

Route Policies

Configuration > [Policies] Route Policies

The **Route Policy** specifies where to direct flows.

By default, the Route Policy auto-optimizes all IP traffic, automatically directing flows to the appropriate tunnel. **Auto-optimization** strategies reduce the need to create explicit route map entries for optimization.

Route Policies Route Policies | QoS Policies | Optimization Policies | Access Lists

map1 Active | map1 is active Add Map Delete Map Rename Map

Add Rule 6/400 rules used

Priority	ACL	Protocol	Match Criteria					Set Actions				Comment		
			Source IP/Subnet	Dest IP/Subnet	Application	Source...	DSCP	Interface	Destination	Path	Fallback			
10		igmp	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	pass-through					X
20	offic...	ip	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	auto_tun_10.1.153.20_to_1...		continue			X
300	bett...	ip	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	auto optimized	default	pass-through			X
500	bett...	ip	0.0.0.0/0	0.0.0.0/0	any	0:0	af13	any	auto_tun_10.1.153.20_to_1...		pass-through			X
1000		ip	0.0.0.0/0	0.0.0.0/0	doubletake	0:0	any	any.any	pass-through					X
1010		ip	172.20.0.0/16	0.0.0.0/0	tmbuktu	0:0	any	any.any	drop					X
65535		ip	any	any	any	0:0	any	any	auto optimized	default	pass-through			

Apply Cancel

The three strategies that Silver Peak uses are **TCP-based** auto-opt, **IP-based** auto-opt, and **subnet sharing**. By default, all three are enabled.

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

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

Also, where multiple tunnels exist to the remote peer, the appliance may be configured to dynamically select the best path based on one of these criteria:

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

When using **Path** selection, a best practice is to specify the same path for both appliances (for example, **low-latency**).

Priority

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

Source or Destination

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

QoS Policies

Configuration > [Policies] QoS Policies

The QoS Policy determines how flows are queued and marked.

QoS Policies ? [Route Policies](#) | [QoS Policies](#) | [Optimization Policies](#) | [Access Lists](#) | [Shaper](#)

map2 Active | map2 is active

5/200 rules used

2 Rows Search

Priority	Match Criteria								Set Actions			Comment	
	ACL	Protocol	Source IP/Subnet	Dest IP/Subnet	Application	Sources...	DSCP	Interface	Traffic Class	LAN QoS	WAN QoS		
1000	ip	ip	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	4 - best-effort	trust-lan	trust-lan	evening map	<input type="checkbox"/>
65535		ip	any	any	any	0:0	any	any	1 - default	trust-lan	trust-lan		<input type="checkbox"/>

The QoS Policy's SET actions determine two things:

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

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

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

Priority

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

Source or Destination

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

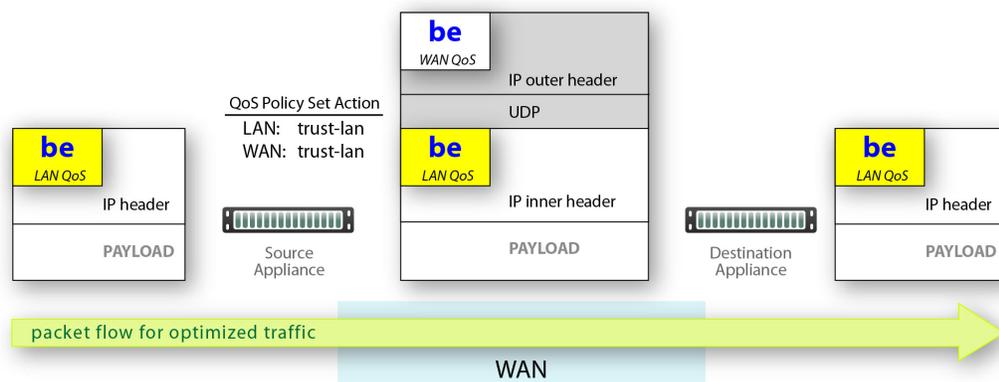
Handling and Marking DSCP Packets

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

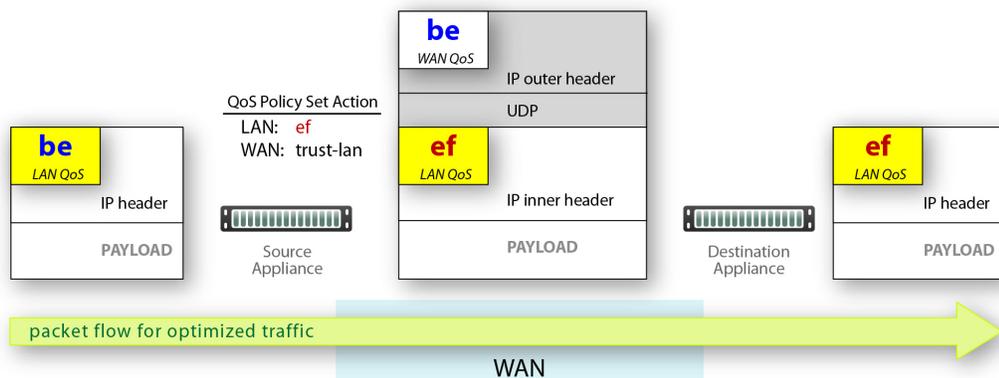
Applying DSCP Markings to Optimized (Tunnelized) Traffic

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

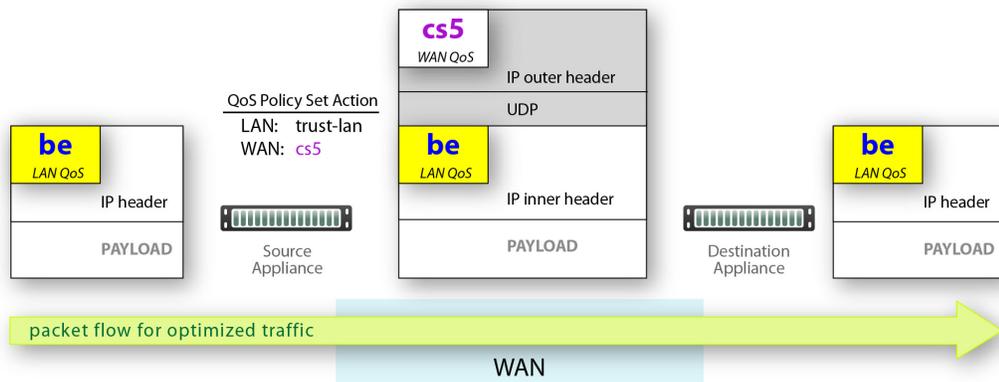
LAN and WAN set to trust-lan



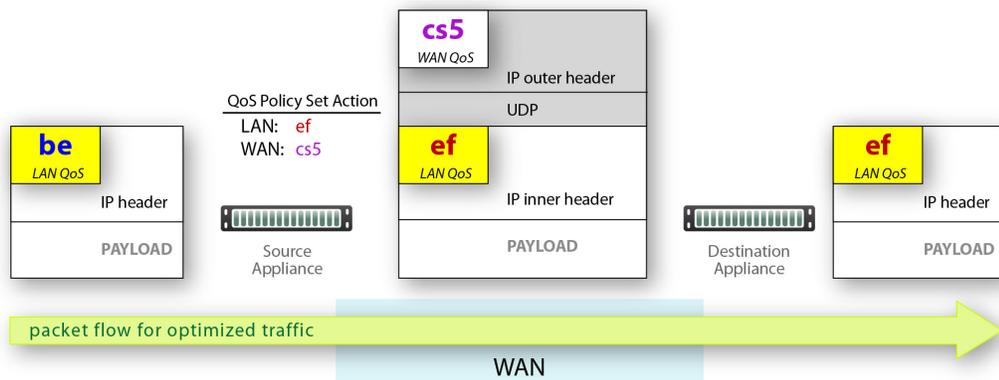
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



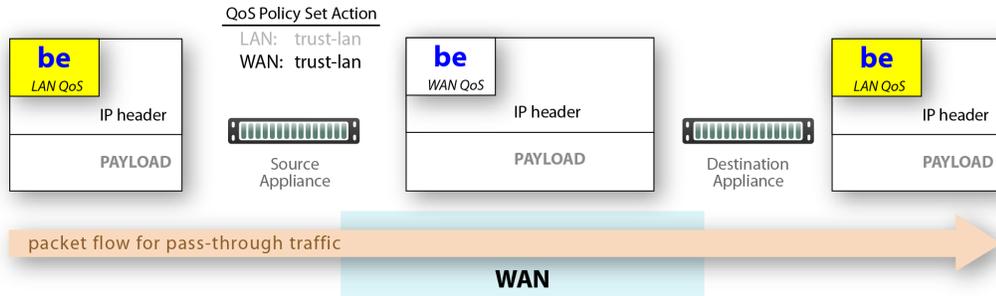
LAN setting changed, WAN setting changed



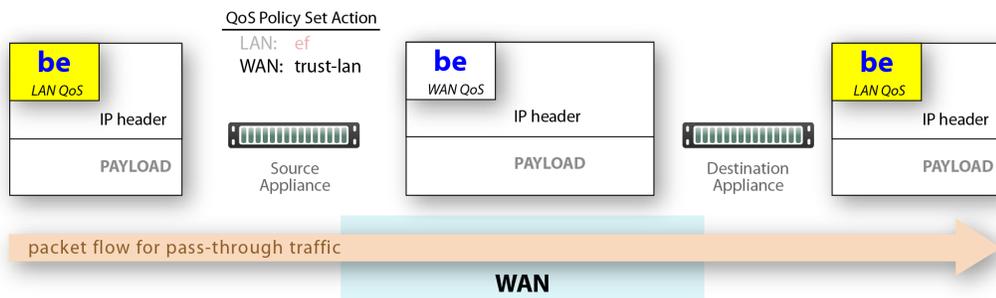
Applying DSCP Markings to Pass-through Traffic

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

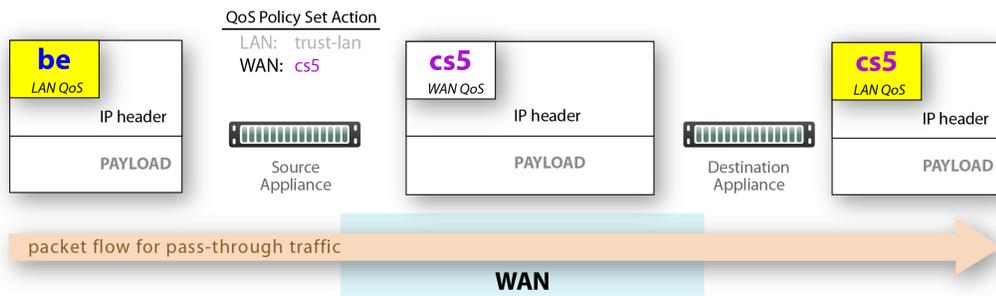
LAN and WAN set to trust-lan



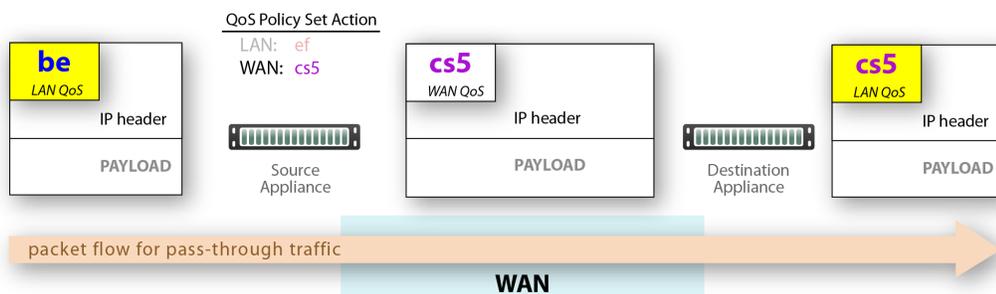
LAN setting changed, WAN is trust-lan



LAN is trust-lan, WAN setting changed



LAN setting changed, WAN setting changed



Optimization Policies

Configuration > [Policies] Optimization Policies

The **Optimization Policy** specifies which optimizations to apply to flows. The table includes the appliance-based defaults for CIFS, SSL, Citrix, and iSCSI.

Optimization Policies ? [Route Policies](#) | [QoS Policies](#) | [Optimization Policies](#) | [Access Lists](#)

map1 Active | map1 is active

7/200 rules used

Priority	Match Criteria								Set Actions						Comment	
	ACL	Protocol	Source IP/Subnet	Dest IP/Subnet	Application	Source...	DSCP	Interface	Network M...	IP Header ...	Payload Co...	TCP Accel	TCP Accel D...	Protocol Accel		
10000		tcp	0.0.0.0/0	0.0.0.0/0	any	0:139	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	cifs	<input checked="" type="checkbox"/>	
10010		tcp	0.0.0.0/0	0.0.0.0/0	any	0:445	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	cifs	<input checked="" type="checkbox"/>	
10020		tcp	0.0.0.0/0	0.0.0.0/0	any	0:443	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	ssl	<input checked="" type="checkbox"/>	
10030		tcp	0.0.0.0/0	0.0.0.0/0	any	0:2598	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	citrix	<input checked="" type="checkbox"/>	
10040		tcp	0.0.0.0/0	0.0.0.0/0	any	0:1494	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	citrix	<input checked="" type="checkbox"/>	
10050		tcp	0.0.0.0/0	0.0.0.0/0	any	0:860	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	iscsi	<input checked="" type="checkbox"/>	
10060		tcp	0.0.0.0/0	0.0.0.0/0	any	0:3260	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	iscsi	<input checked="" type="checkbox"/>	
65535		ip	any	any	any	0:0	any	any	balanced	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	none	<input type="checkbox"/>	

Set Actions Definitions

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

Priority

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

Source or Destination

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

Access Lists

Configuration > [Policies] Access Lists

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

Access Lists ? [Route Policies](#) | [QoS Policies](#) | [Optimization Policies](#) | [Access Lists](#) | [NAT Policies](#)

office-lab

This ACL is used by RouteMap-map1

Add Rule 2 / 100 rules used

2 Rows Search

Priority ▲	Match Criteria							Set Actions		
	Protocol	Source IP/Subnet	Dest IP/Subnet	Application	Source:Dest...	DSCP	Interface	Permit	Comment	
10	l2tp	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any,24	deny		X
20	ip	0.0.0.0/0	0.0.0.0/0	edonkey	0:0	any	any	permit		X

- An **ACL** is a reusable **MATCH** criteria for filtering flows, and is associated with an action, **permit** or **deny**: You can use the same ACL as the **MATCH** condition in more than one policy --- Route, QoS, or Optimization.
- An Access Control List (ACL) consists of one or more ordered access control rules.
- An ACL only becomes active when it's used in a policy.
- **Deny** prevents further processing of the flow by *that ACL, specifically*. The appliance continues to the next entry in the policy.
- **Permit** allows the matching traffic flow to proceed on to the policy entry's associated **SET** action(s). The default is **permit**.
- When creating ACL rules, list **deny** statements first, and prioritize less restrictive rules ahead of more restrictive rules.

Priority

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

Source or Destination

- An IP address can specify a subnet - for example: 10.10.10.0/24.
- To allow **any IP address**, use 0.0.0.0/0.
- Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
- To allow **any port**, use 0.

Built-in Applications

Configuration > [Policies] Applications, Built-in

This page lists the appliance's built-in applications.

Built-in Applications

TCP/IP Ports used by the Silver Peak appliances

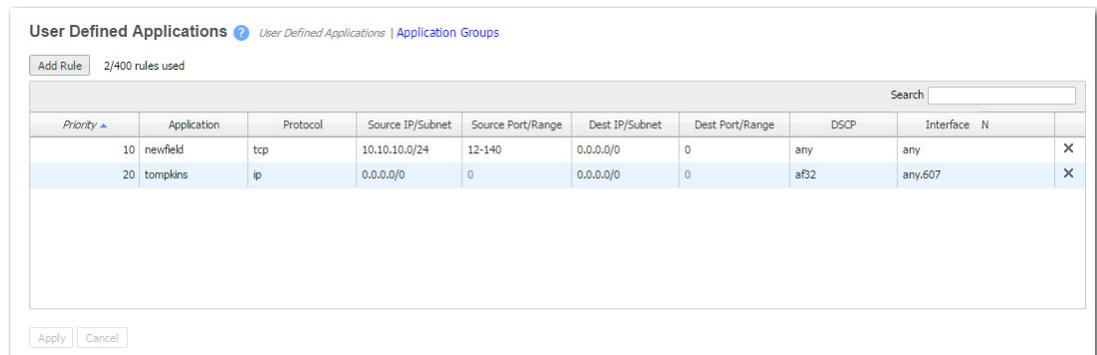
Search

Name ▲	Ports
3par	[TCP] 5781-5783, 5785 [UDP] 5781-5783, 5785
aol	[TCP] 5191-5193
aol_im	[TCP] 4443, 5190
app_assure_replication	[TCP] 8005
app_assure_svr_backup	[TCP] 8001
aspera	[TCP] 33001 [UDP] 33001
avamar	[TCP] 7778, 27000, 28001-28002, 29000
backweb	[UDP] 370
bit_torrent	[TCP] 6881-6999
bluearc	[TCP] 32963
celerra	[TCP] 5085, 8888
centera	[TCP] 3218, 3682 [UDP] 3218
cifs_smb	[TCP] 139, 445
cisco_skinny	[TCP] 2000 [UDP] 2000
citrix-bcast	[UDP] 1604
citrix-cgp	[TCP] 2598
citrix-ica	[TCP] 1494
citrix-ima	[TCP] 2512-2513
commvault	[TCP] 8400-8403 [UDP] 8400-8403

User-Defined Applications

Configuration > [Policies] Applications, User-Defined

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



UDAs are specific to the appliance on which they're defined. Where can you use them?

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

Behavior

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

Priority

- Range = 1 - 50000
- However, if you're using Orchestrator templates, they will overwrite and delete UDAs on the appliance that have priorities over 999.
- By default, adding a rule/application increments the last Priority by 10.

Source or Destination

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

Application Groups

Configuration > [Policies] Application Groups

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

The screenshot displays the 'Application Groups' configuration page. At the top, there is a header 'Application Groups' with a help icon and a breadcrumb 'User Defined Apps | Application Groups'. Below the header is an 'Add Groups' button. A search bar is located on the right. A dropdown menu shows 'Show 10'. The main content is a table with two columns: 'Group Name' and 'Applications'. The table lists five groups with their respective applications and a delete icon (X) in the third column.

Group Name	Applications	
citrix	citrix-cgp,citrix-ica,citrix-ima,citrix-bcast	X
encrypted	nntp,ddm_ssl,ipsec,imap4s,ssh,pop3s,https,smtps,telnet	X
interactive	pcanywhere,ssh,pcop,vnc,telnet,telnet,xwindows	X
real-time	h_323,cisco_skinny,t_120,rtcp,rtsp	X
replication	app_assure_replication,netvault_replication,centera,bluearc,vsphere_replication_ini,celerra,aspera,commv...	X

At the bottom of the table, it says 'Showing 1 to 5 of 5 entries'. Below the table are 'Apply' and 'Cancel' buttons. At the bottom right, there are navigation buttons: 'First', 'Previous', '1', 'Next', and 'Last'.

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

Flow Redirection

Configuration > [Policies] Flow Redirection

Optimizing TCP flows requires that a client request and its server response use the same path through the network. If not, then the network is asymmetric. **Flow redirection** removes the asymmetry locally by merging the traffic of an asymmetric flow into a single appliance.

Flow redirection moves packet traffic between appliances that you assign to a **cluster**:

- A cluster can contain just one appliance (in which no redirection occurs) or several appliances (in which redirection may occur between different pairs).
- All the appliances in a cluster are equal peers.
- You can have up to 32 peers in a cluster.
- The Silver Peak Communication Protocol (SPCP) formalizes peer-to-peer communications in an appliance cluster. SPCP is both a discovery and control protocol. By default, SPCP uses **mgmt1** to communicate between appliances.
- This must be a Layer 2 connection. In other words, you want a switch — not a router — between any two peers.

For each peer appliance in a cluster, flow redirection requires configuration on three pages:

- **Configuration - Interfaces** -- for configuring the **mgmt1** IP address.
- **Configuration - IP Route** -- for configuring the necessary static route(s).
- **Configuration - Flow Redirection** -- for enabling flow redirection, selecting the management interface, and identifying the peers in the cluster.



Note IMPORTANT — When configuring for flow redirection, the **mgmt1** interfaces must be in a separate subnet from the **mgmt0** interfaces.

Definitions

Enable: Enables/disables flow redirection.

Wait time: Allowable timeout between peers when negotiating which appliance should take the flow. This default value rarely requires change -- possibly in environments with longer latencies or when using a datapath interface instead of **mgmt1**.

Interface: The interface on which flow redirection occurs.

PeerIP table: Editable table which lists all Peer IPs and their reachability states.

SaaS Optimization

Configuration > [Policies] SaaS Optimization

Use this page to enable SaaS optimization and accelerate SSL traffic.

For additional compression benefits, set up decryption using a CA (Certificate Authority) certificate.

SaaS Optimization Monitoring

Enable SaaS Optimization

RTT Calculation Interval (1..1440) minutes

Application Name	Optimize	Advertise	RTT Threshold	Domains
GoToTraining	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gototraining.com
GoToWebinar	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	gotoweinar.com, gotoassist.com
Intuit	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	intuit.com
Jobvite	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	careers.jobvite.com, www.jobvite.com, hire.jobvite.com
Lithium	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	lithium.com
LiveOps	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	liveops.com
Marketo	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	marketo.com
NetSuite	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	netsuite.com
Office365	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.officeapps.live.com, *.microsoftonline-p.net, *.microsoftlines...
Parature	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	parature.com
PardotExactTarget	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	pardot.com
PlexSystems	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	plex.com
Salesforce	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	*.na3.force.com, *.salesforce.com
ServiceNow	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	servicenow.com
SuccessFactors	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	successfactors.com
SugarCRM	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	sugarcrm.com
Taleo	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	taleo.com
UltimateSoftwareGroup	<input type="checkbox"/>	<input type="checkbox"/>	10 ms	ultimatesoftware.com

Apply Cancel

SSL Signing Authority for SaaS

Do Not Decrypt SSL Traffic
SSL traffic will be accelerated, but won't be decrypted, so traffic will not be compressed.

Decrypt - Use CA Certificate
For SaaS SSL decryption, the appliance will generate substitute certificates, which must be CA-signed (using the Built-In CA Certificate) or signed by a subordinate CA cert issued from an enterprise CA (Custom CA Certificate). To avoid browser warnings, import the applicable CA Certificate into the Trusted CA Store on client browsers.

Built-In CA Certificate (Issuer: Silver Peak) [View](#)
[Download](#) [Regenerate](#)

Custom CA Certificate [View](#)
Issuer: Silver Peak SSL Proxy
Issued to: Silver Peak SSL Proxy
Expiration: Dec 31 23:59:59 2034 GMT
[Upload and Replace](#) [Download](#) [Delete](#)

[Generate Certificate Signing Request \(CSR\)](#)
Organizations with an enterprise CA can generate a certificate signing request (CSR), submit to a CA administrator, and upload the issued subordinate CA cert.

Setting up SSL Signing Authority for SaaS

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

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

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

- If you **don't** already have a subordinate CA certificate, you can access any appliance's **Configuration > SaaS Optimization** page and generate a Certificate Signing Request (CSR). The workflow would basically follow this pattern:
 1. Click **Generate Certificate Signing Request**, and complete the Certificate Information requested in the dialog box.
 2. Save the CSR and the Private Key.
 3. Submit the CSR to your enterprise CA to obtain a Subordinate CA Certificate.
 4. After approvals are complete and the subordinate CA is in hand, go to the **Configuration > SaaS Optimization** page.
 5. Under **Custom CA Certificate**, click **Upload and Replace** to import the subordinate CA.

What Silver Peak Supports

- When you upload the certificate.
 - You can add either a PFX certificate (generally, for Microsoft servers) or a PEM certificate.
 - The default is PEM when PFX Certificate File is deselected.
- Silver Peak supports:
 - X509 Privacy Enhanced Mail (PEM), Personal Information Exchange (PFX), and RSA key 1024-bit and 2048-bit certificate formats.
 - SAN (Subject Alternative Name) certificates. SAN certificates enable sharing of a single certificate across multiple servers and services.
- Silver Peak appliances support:
 - **Protocol versions:** SSLv3, SSLv3.3, TLS1.0, TLS1.1, TLS1.2
 - **Key exchanges:** RSA, DHE, ECDHE
 - **Authentication:** RSA
 - **Cipher algorithms:** RC4, 3DES, AES128, AES256, AES128-GCM, AES256-GCM
 - **Message Digests:** MD5, SHA, SHA256, SHA284

Optimizing SaaS Services

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. This option is located on the appliance's **Configuration > SaaS Optimization** page.

- If **Advertise** is *selected* for a service (for example, SFDC), the appliance will:
 - Ping active SaaS subnets to determine RTT/metric
 - Add subnet sharing entries locally for subnets within RTT threshold
 - Advertise subnets and their metric (within threshold) via subnet sharing to client-side appliances

- Upon seeing an SFDC flow, generate a substitute certificate for an SFDC SSL domain (one substitute certificate per domain)
- Auto-generate dynamic NAT rules for SFDC (but not for unchecked services)

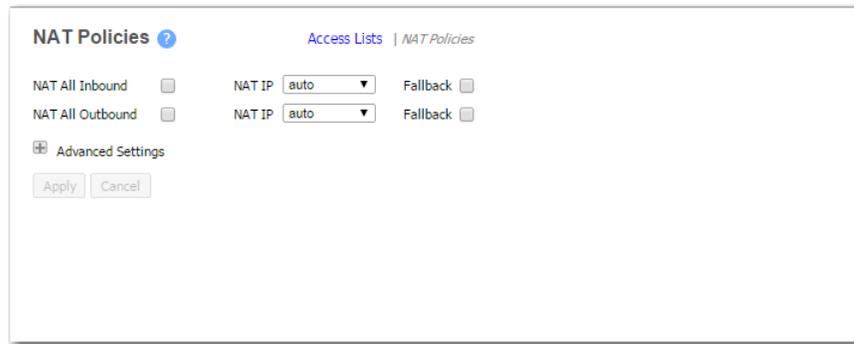
- When **Optimize** is *selected* for a service (for example, SFDC), the appliance will:
 - Ping active SFDC subnets to determine the RTT (metric)
 - Does not advertise metric via subnet sharing (unless **Advertise** is also selected)
 - Receives subnet sharing metric (RTT) from associated appliances
 - Compares its own RTT (local metric) with advertised metric
 - If its own RTT is lower, then the packet is sent pass-through (direct to the SaaS server).
 - If an advertised RTT is lower, then the packet is tunnelized.
 - Generate a substitute certificate for an SFDC SSL domain (one sub cert per domain)
 - No NAT rules created

- When **Optimize** is *not selected* for a service (for example, SFDC), the appliance:
 - Receives subnet sharing advertisements for SFDC but doesn't use them
 - Does no RTT calc pinging
 - Does not participate in SSL
 - Creates no NAT rules
 - Sends all SFDC traffic as pass-through

- For a detailed list of the enabled SaaS applications, click **Monitoring**.

NAT Policies

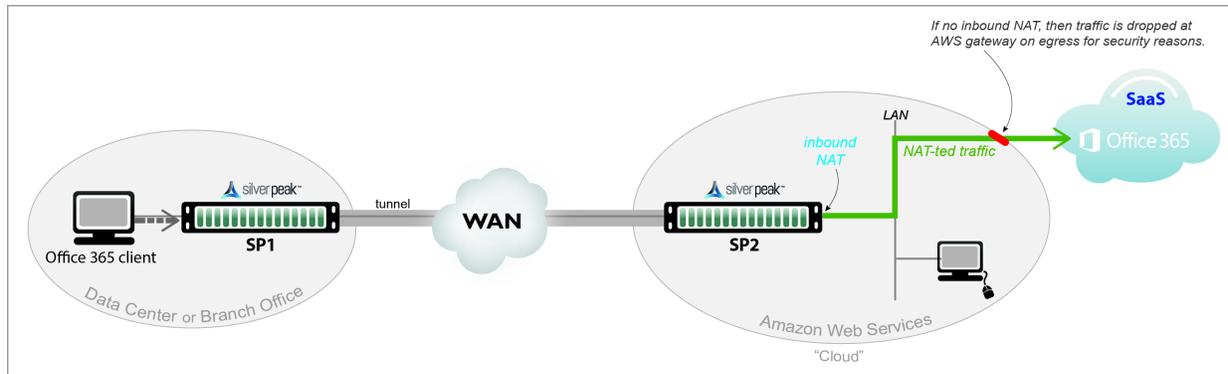
Configuration > [Policies] NAT Policies



Two use cases illustrate the need for NAT:

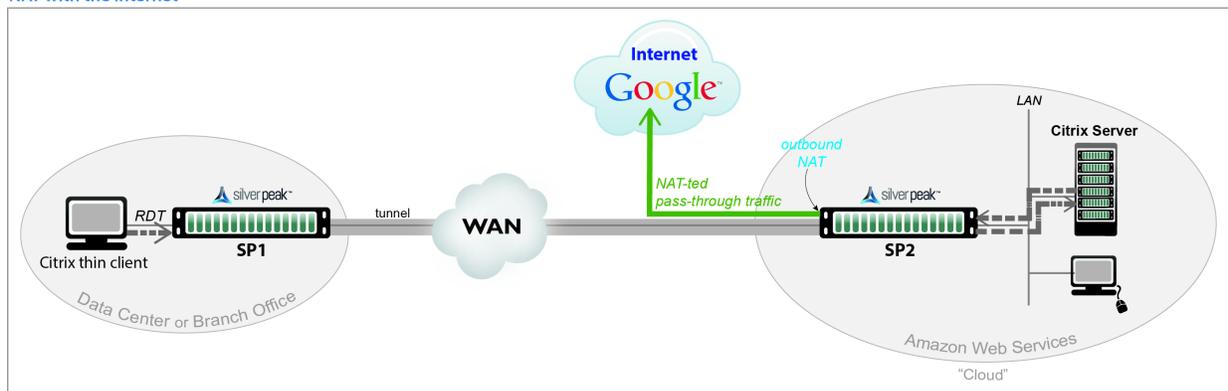
- 1 **Inbound NAT.** The appliance automatically creates a source NAT (Network Address Translation) map when retrieving subnet information from the Silver Peak Cloud portal. This ensures that traffic destined to SaaS servers has a return path to the appliance from which that traffic originated.

NAT with a SaaS Service



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

NAT with the Internet



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

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

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

Advanced Settings

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

The screenshot shows the NAT Policies configuration page. At the top, there are options for 'NAT All Inbound' and 'NAT All Outbound', each with a checkbox and a 'Fallback' checkbox. Below these are 'NAT IP' dropdown menus set to 'auto'. The 'Advanced Settings' link is circled in red. Below the settings, there are buttons for 'Add Map', 'Delete Map', and 'Rename Map', along with a 'Show Dynamic Policies' checkbox. A table below shows the configuration for 'map1' with 'Active' checked. The table has columns for 'Match Criteria' and 'Set Actions'. The table contains one row with the following data:

Priority	Match Criteria						Set Actions					Comment	
	ACL	Protocol	Source IP/Subnet	Dest IP/Subnet	Application	Source:D...	DSCP	Interface	NAT Type	NAT Direction	NAT IP		Fallback
65535		ip	any	any	any	0:0	any	any	no-nat	none	auto	<input type="checkbox"/>	

There are two types of NAT policies:

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

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

- **Source or Destination**
 - An IP address can specify a subnet - for example: 10.10.10.0/24.
 - To allow **any IP address**, use 0.0.0.0/0.
 - Ports are available only for the protocols **tcp**, **udp**, and **tcp/udp**.
 - To allow **any port**, use 0.

- **NAT Type**
 - **no-nat** is the *default*. No IP addresses are changed.
 - **source-nat** changes the source address and the source port in the IP header of a packet.

- **NAT Direction**
 - **inbound** NAT is on the LAN interface.
 - **outbound** NAT is on the WAN interface.
 - **none** -- the only option if the NAT Type is **no-nat**.

- **NAT IP**
 - **auto** -- Select if you want to NAT **all** traffic. The appliance then picks the first available NAT IP/Port.
 - **tunnel** -- Select if you only want to NAT **tunnel** traffic. Applicable only for inbound NAT, as outbound doesn't support NAT on tunnel traffic.
 - **[IP address]** -- Select if you want to make NAT use this IP address during address translation.

- **Fallback** -- If the IP address is full, the appliance uses the next available IP address.

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

Threshold Crossing Alerts

Configuration > [Policies] Threshold Crossing Alerts

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

Threshold Crossing Alerts ?

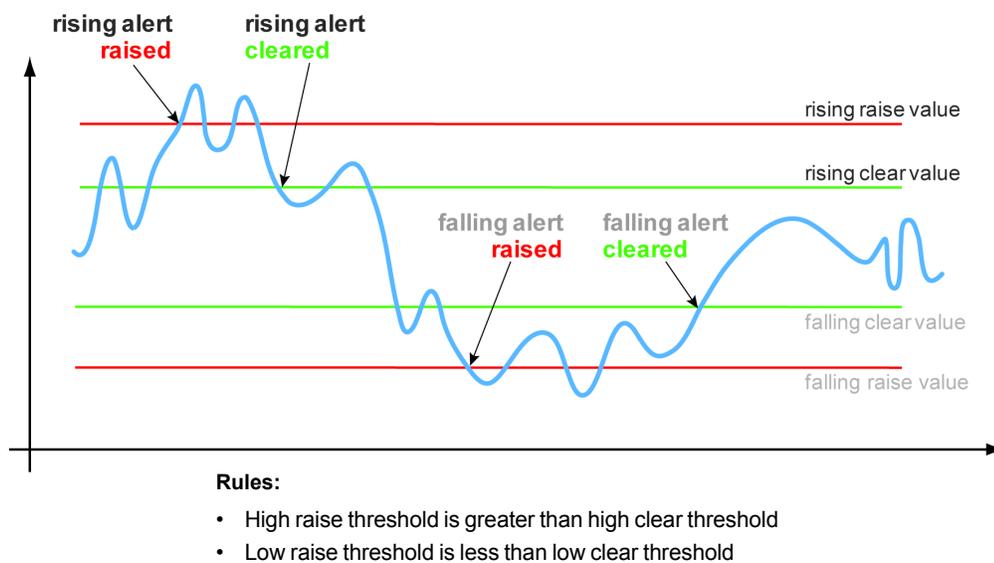
Search

Name ▲	Rising				Falling			
	Raise	Clear	Times to Trigger	Enabled	Raise	Clear	Times to Trigger	Enabled
File-system utilization	90%	85%	1	<input checked="" type="checkbox"/>	75%	75%	1	<input type="checkbox"/>
LAN-side receive throughput	1000000 kbps	1000000 kbps	1	<input type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>
Total number of flows	90%	85%	1	<input checked="" type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Total number of optimized flows	85%	80%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel latency	1000 ms	850 ms	1	<input checked="" type="checkbox"/>	0 ms	0 ms	1	<input type="checkbox"/>
Tunnel loss post-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel loss pre-FEC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel OOP post-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel OOP pre-POC	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel reduction	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
Tunnel utilization	100%	100%	1	<input type="checkbox"/>	0%	0%	1	<input type="checkbox"/>
WAN-side transmit throughput	1000000 kbps	1000000 kbps	1	<input type="checkbox"/>	0 kbps	0 kbps	1	<input type="checkbox"/>

Apply Cancel

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

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



ON by default:

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

OFF by default:

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

Monitoring Traffic

This chapter describes the various tools available for monitoring performance, and reviewing traffic and application statistics.

Alarms are addressed in a separate chapter.

In This Chapter

- **Introduction** See page 78.
- **Application View Tab** See page 79.
- **Network View Tab** See page 80.

Statistics

- **Charts** See page 81.
- **Application Statistics** See page 87.
- **Realtime Charts** See page 89.
- **Flows** See page 90.
- **QoS Statistics** See page 108.
- **Tunnel Statistics** See page 109.
- **Internet Statistics** See page 111.
- **NetFlow Statistics** See page 112.
- **Interface Statistics** See page 113.

Status

- **Bridge Interfaces** See page 115.
- **Route Next Hops** See page 116.
- **VRRP Status** See page 117.
- **Flow Redirection Statistics** See page 119.
- **SaaS Optimization** See page 121.
- **Orchestrator Reachability** See page 122.

Introduction

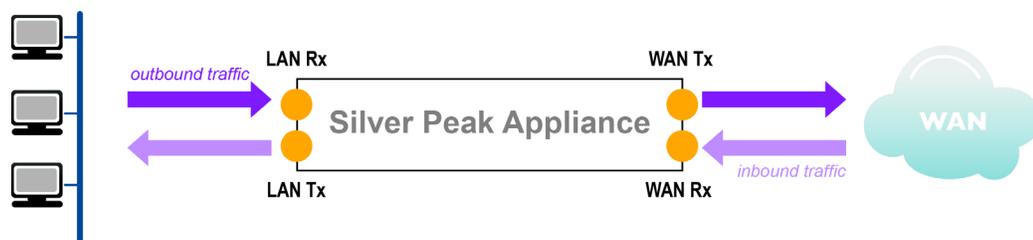
The top-level **Application View** tab and **Network View** tab provide charted summaries of performance. Both tabs display the 10 **Top Flows** — a subset of the **Monitoring** menu's **Flows**.

Additionally, the **Monitoring** menu provides a variety of reports.

Understanding Traffic Direction

In Appliance Manager, statistics and reports either reference the direction of the flow or the point(s) where the data is collected.

When you download raw data, the **.csv** file uses the following abbreviated terms in column headers: **LAN Rx**, **WAN Tx**, **LAN Tx**, and **WAN Rx**.



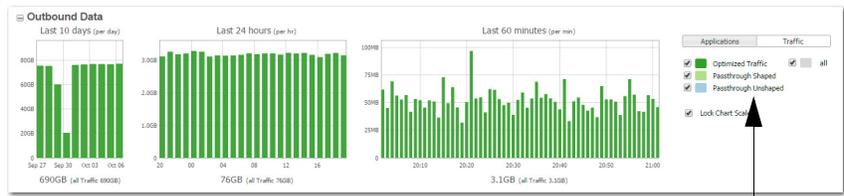
- **LAN-to-WAN** refers to traffic exiting the LAN, destined for the WAN. This flow is also referred to as *outbound traffic*.
- **WAN-to-LAN** refers to traffic coming from the WAN, destined for the LAN. This flow is also referred to as *inbound traffic*.



Tip Here's a helpful mnemonic for remembering the difference:

- **Rx** is "**R**eceive **fR**om", so **LAN Rx** is "receive from LAN"
- **Tx** is "**T**ransmit **T**o", so **LAN Tx** is "transmit to LAN"

Application View Tab



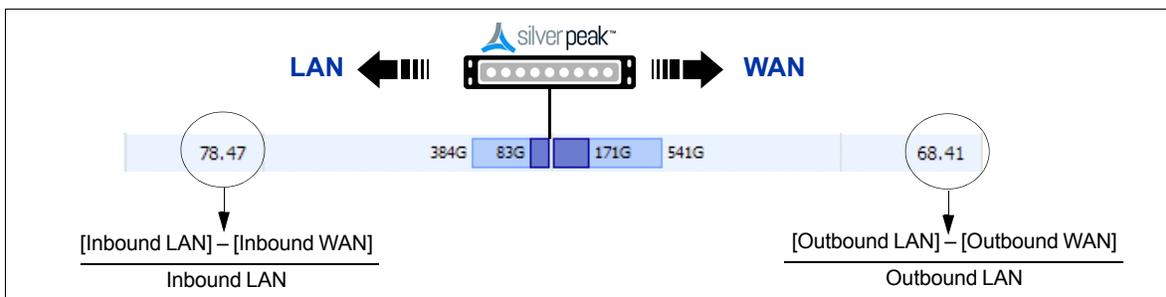
The **Traffic** option profiles the same data, but color-codes it by traffic type.

Top Flows Active 10 Pass-Through 0 Asymmetric 0 Stale 0 | Displayed 10 Matched 10

Application	IP1	Port1	IP2	Port2	Reduction %	Inbound		Outbound		Reduction %	Protocol	Start Time (UTC)	
						Bytes	Bytes	Bytes	Bytes				
gnutella	10.1.153.10	38947	10.1.154.10	6347	69.9	173G	52G	611M	52G	173G	0.0	tcp	01-Oct-15 17:41:42
gnutella	10.1.153.10	6347	10.1.154.10	46500	0.0		564M	52G	173G	69.9	tcp	01-Oct-15 17:42:01	
datadomain	10.1.153.10	44054	10.1.154.10	2051	59.9	130G	52G	501M	52G	130G	0.0	tcp	01-Oct-15 17:41:42
datadomain	10.1.153.10	2051	10.1.154.10	40981	0.0		441M	52G	130G	59.9	tcp	01-Oct-15 17:42:01	
3par	10.1.153.10	48543	10.1.154.10	5781	49.9	104G	52G	436M	52G	104G	0.0	tcp	01-Oct-15 17:41:42
3par	10.1.153.10	5781	10.1.154.10	38819	0.0		363M	52G	104G	49.9	tcp	01-Oct-15 17:42:01	
aol	10.1.153.10	46529	10.1.154.10	5191	29.9	74G	52G	359M	52G	74G	0.0	tcp	01-Oct-15 17:41:42
aol	10.1.153.10	5191	10.1.154.10	59803	0.0		269M	75G	52G	30.0	tcp	01-Oct-15 17:42:01	
unassigned	10.1.153.10	53363	10.1.154.10	6181	20.3	66G	52G	134M	52G	66G	0.0	tcp	01-Oct-15 17:41:42
unassigned	10.1.153.10	6181	10.1.154.10	45548	0.0		288M	66G	52G	20.4	tcp	01-Oct-15 17:42:01	

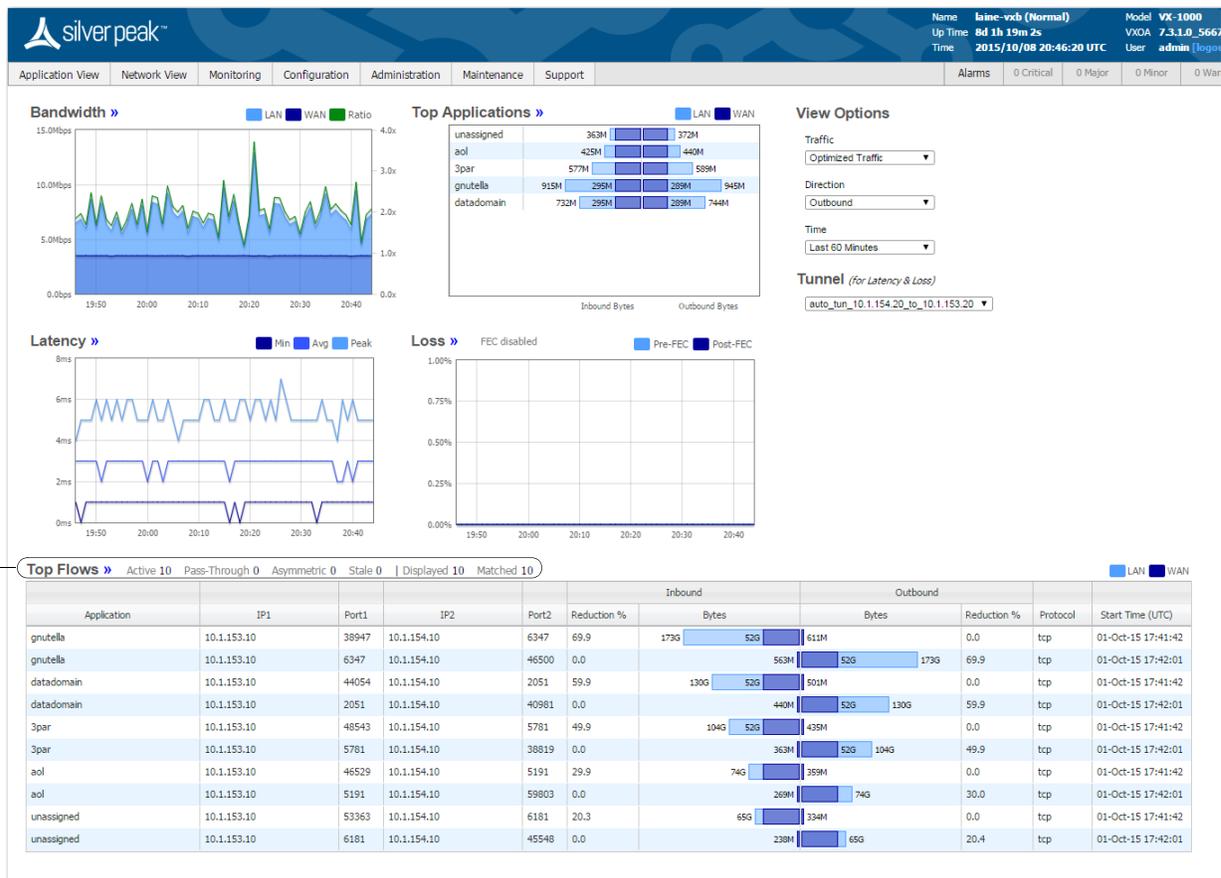
Links to **Current Flows** [in Monitoring]. For more information, see “Flows” on page 90.

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



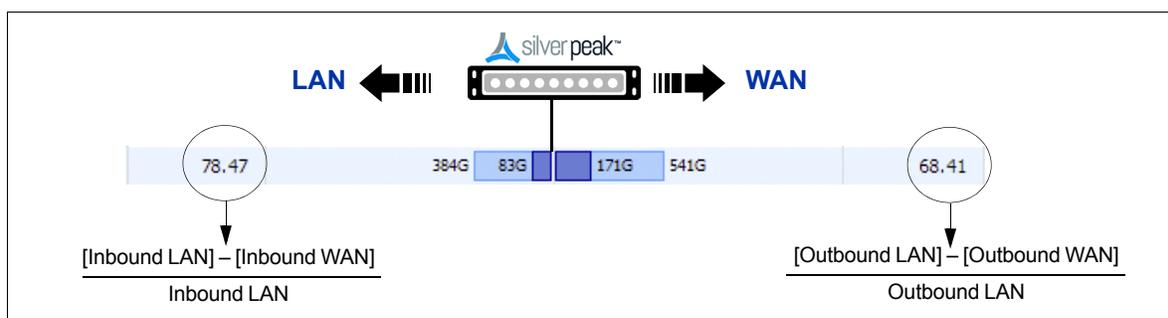
Network View Tab

When you log in, this page opens by default.



Links to **Flows** [in Monitoring]. For more information, see *"Flows"* on page 90.

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



Charts

Charts feature spark lines, as well as selectable (and modifiable) time ranges for any data collected in the last 30 days.

Dynamic charts exist for the following:

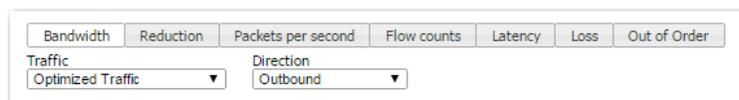
- **Bandwidth** See page 83.
- **Reduction** See page 83.
- **Packets per Second** See page 84.
- **Flow Counts** See page 84.
- **Latency** See page 85.
- **Loss** See page 85.
- **Out-of-Order Packets** See page 86.

Charts consist of filters, a main chart display, and a time selection area.



1

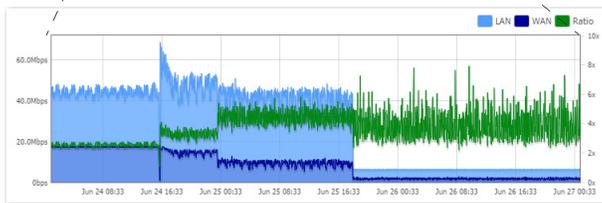
FILTER SELECTION



2 CHART DISPLAY — Legend / X-axis / Y-axis



To zoom in... click, drag, and release. The chart updates to show the new range.



The Y-axis height is calibrated to the maximum Y value shown in the selected range.

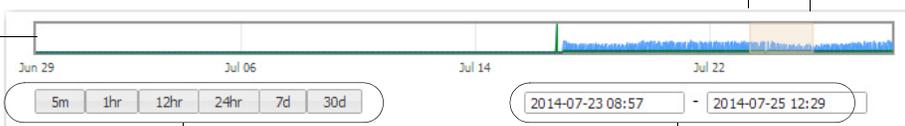
The Y-axis calibration may change if you hide a parameter (LAN, WAN, or Ratio) that has higher values than the remaining parameters.

3 TIME SELECTION

To change the time frame, you can also

- drag the **slider** to change its position
- change the **slider's** size — click and drag an edge

Spark lines showing the activity



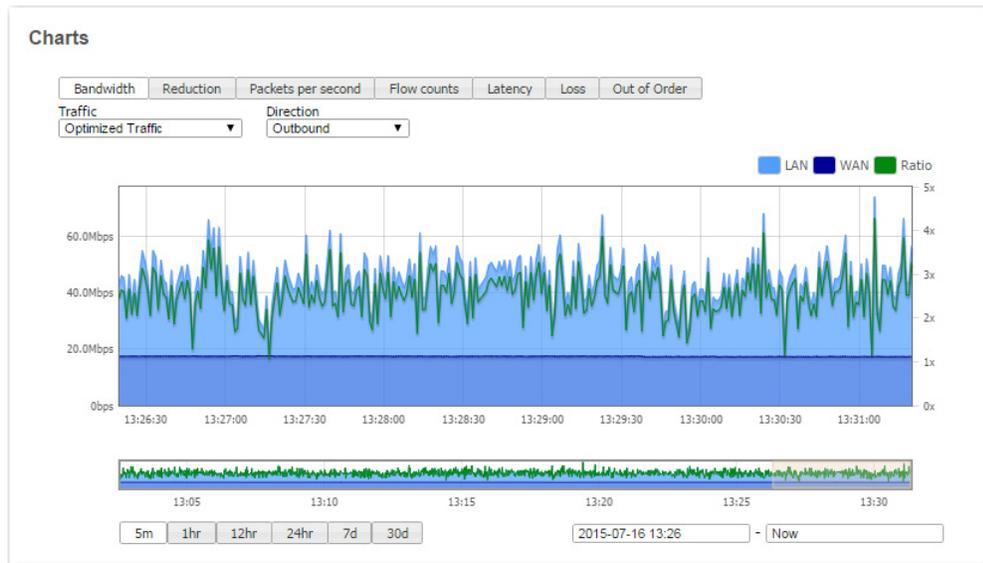
When you select, one endpoint is always NOW.

Displays the slider's range. You can also set it.

Bandwidth

Monitoring > [Statistics] Charts - Bandwidth

The **Bandwidth** chart shows the rate at which data was sent and/or received in each time interval.



Reduction

Monitoring > [Statistics] Charts - Reduction

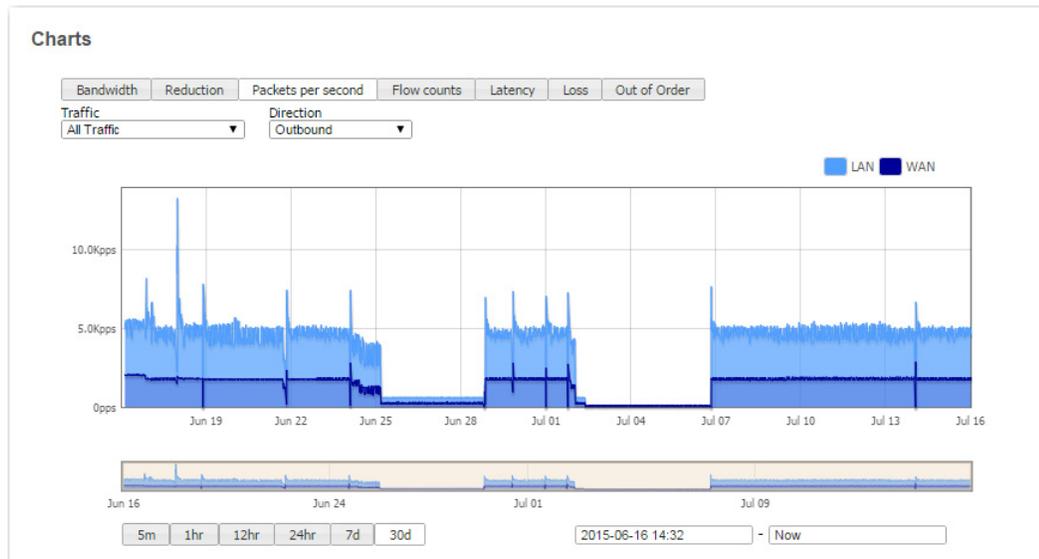
The **Reduction** chart shows the rate at which data was sent and/or received in each time interval.



Packets per Second

Monitoring > [Statistics] Charts - Packets per second

The **Packets per second** chart shows the rate at which data was sent and/or received in each time interval.



Flow Counts

Monitoring > [Statistics] Charts - Flow counts

The **Flow Counts** chart shows the number of flows and differentiates them into TCP (accelerated and unaccelerated) and non-TCP flows.

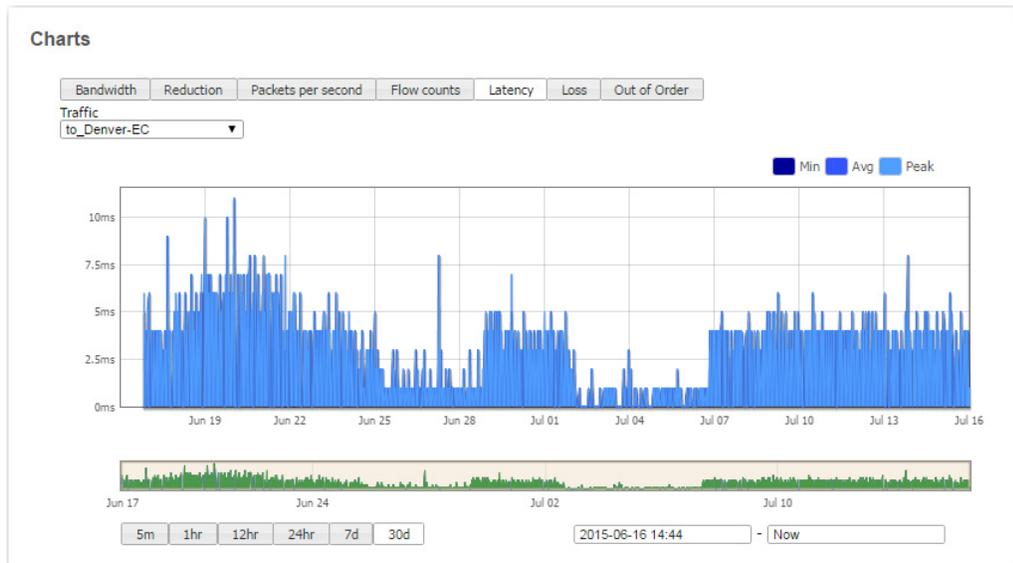


Since CIFS acceleration is a subset of TCP acceleration, that data is included in the accelerated TCP flow data.

Latency

Monitoring > [Statistics] Charts - Latency

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



Loss

Monitoring > [Statistics] Charts - Loss

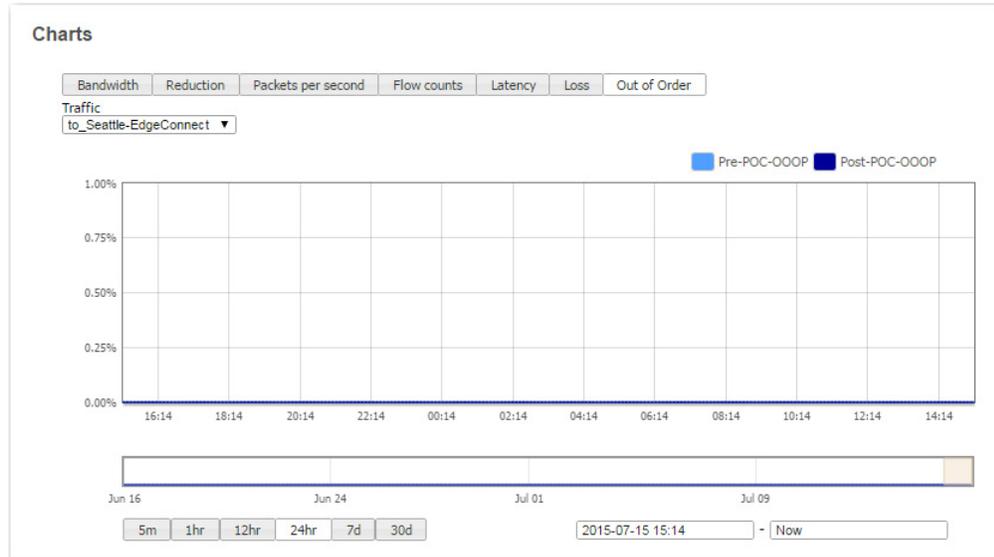
The **Loss** chart shows the percentage of dropped packets in a tunnel, and the effect of **Forward Error Correction (FEC)** on attenuating any loss.



Out-of-Order Packets

Monitoring > [Statistics] Charts - Out of Order

The **Out of Order Packets** chart summarizes, by tunnel, the percentage of packets lost before and after enabling **Packet Order Correction (POC)**.



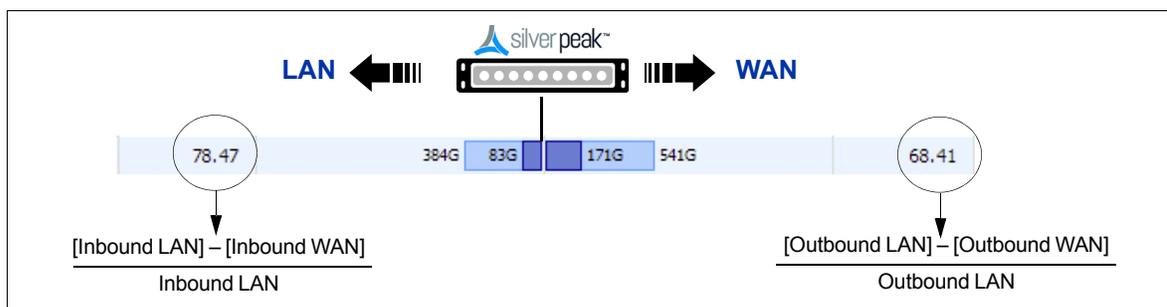
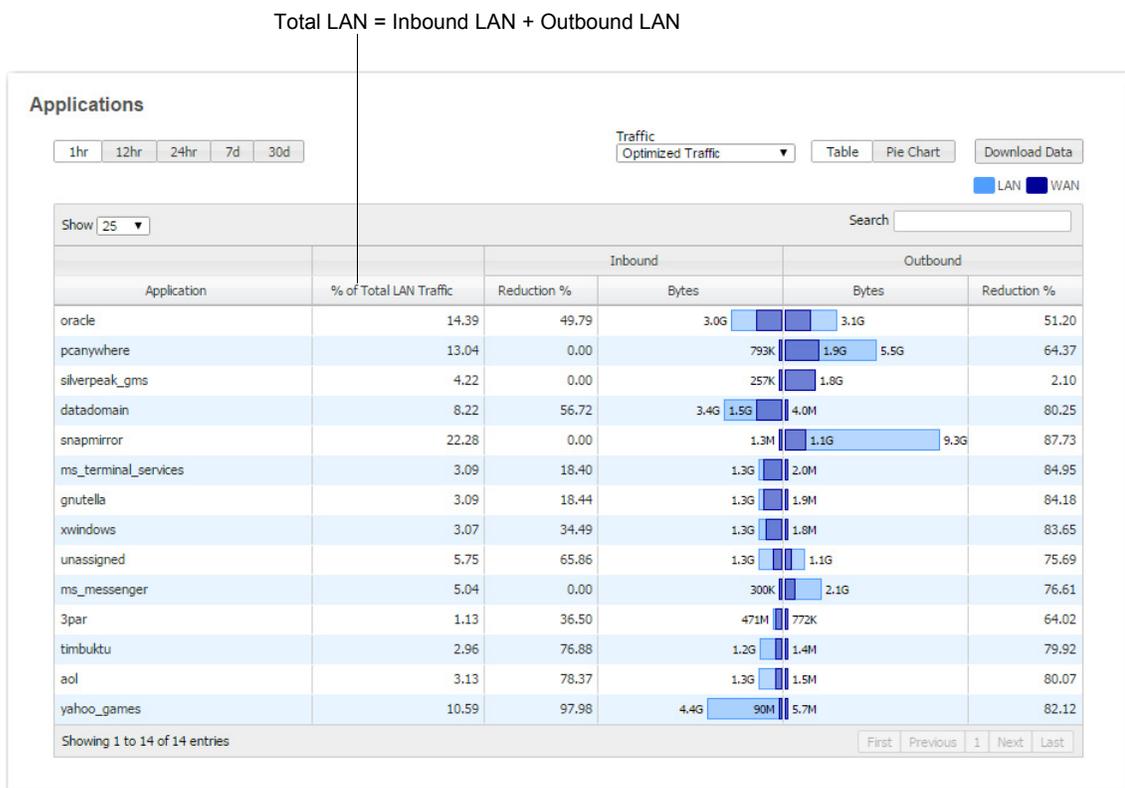
Application Statistics

Monitoring > [Statistics] Applications

The **Applications** page provides table and pie chart views of applications. They show which applications have sent the most bytes:

- What percentage of total LAN traffic does each application comprise?
- What is the data reduction in each direction?
- When comparing outbound and inbound traffic, how are the application distributions different?
- What is the ratio of LAN-to-WAN or WAN-to-LAN traffic for any given application?

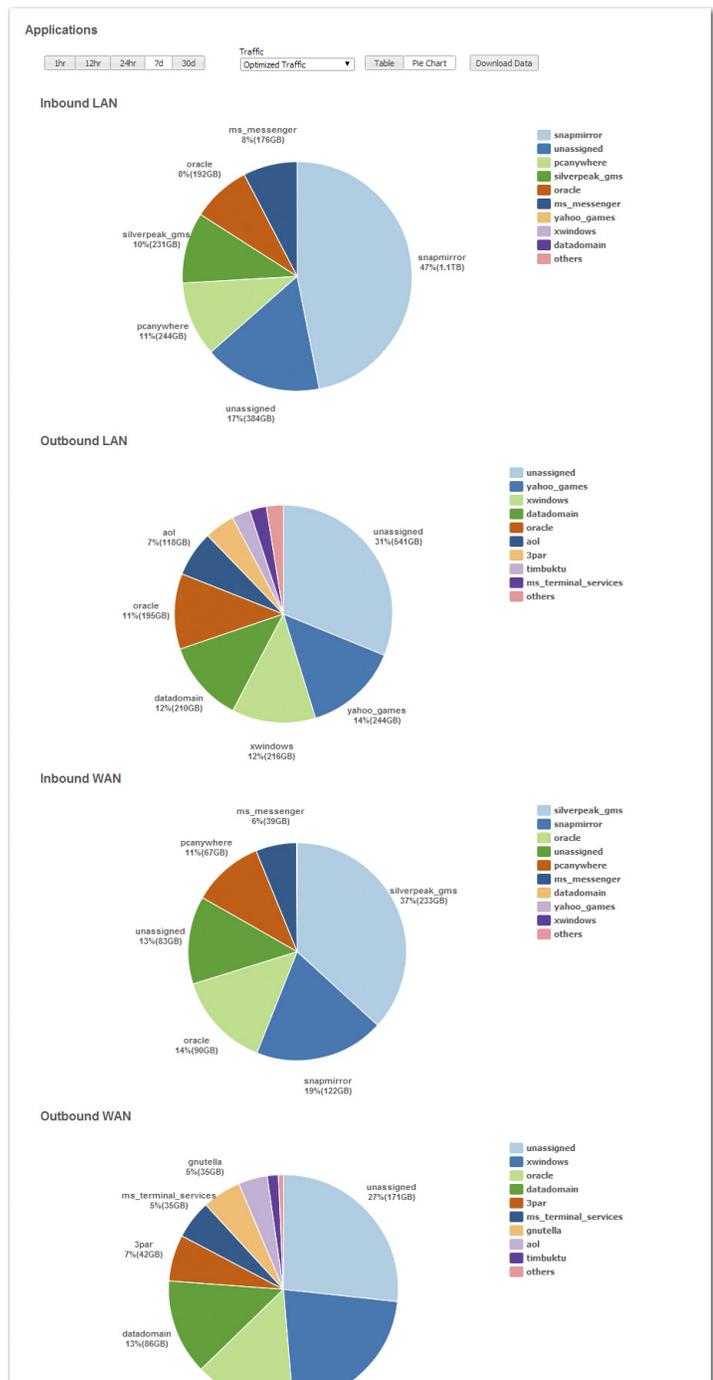
Table View



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

Pie View

The pie view displays the Top 10 applications.



What's the difference between **other** and **unassigned** in the **Applications** stats pie chart?

- The pie chart shows the top ten applications and, possibly, **other**.
- **unassigned** means the sum of traffic (bytes) for which the appliance could not determine the application.
- **other** means the sum of traffic other than the top nine.
- If you don't have more than ten applications, you won't see **other**.

Realtime Charts

Monitoring > [Statistics] Realtime Charts

- For each realtime chart, specify a filter and a metric.

Type of Stats	Filters
Tunnel Stats	Tunnel
Aggregate Tunnel Stats	Traffic Type [Optimized, All]
DSCP Stats	Traffic Type [Optimized, All]; DSCP [1 – 64]
Traffic Class Stats	Traffic Type [Optimized, All]; Traffic Class [1 – 10]
Flow Stats	Traffic Type [Optimized, All]; Flow Type [TCP Acc, TCP Not Acc, Non-TCP]
Application Stats	Traffic Type [Optimized, All]; Application

- You can view multiple realtime charts simultaneously.
- Realtime charts refresh every 3 seconds.
- Although the plotted data doesn't persist when you leave the page (or refresh the browser), the charts do. They begin plotting anew when you return to the page.

Creates the chart

The **Filter** determines the available **Metrics**.

Realtime Charts

Filter

Type of stats: Application Stats

Traffic Type: Optimized

Application: pcanynwhere

Metric to plot

Metric: Inbound LAN PPS

Show in local time:

Plot Delete All Charts

Outbound WAN Rate bps

App: 3par

Traffic: Optimized

Outbound Compression Ratio

Traffic: Optimized

Inbound LAN PPS

App: pcanynwhere

Traffic: Optimized

Click to delete

Flows

Monitoring > [Statistics] Flows

By default, the **Flows** page retrieves a list of active connections. The maximum visible number depends on which browser you use.

- You can also view flows that have ended.
- The page displays a default set of columns, along with individual links to flow details.
- You can display additional columns from a customization list.

This section discusses the following topics:

- **How the Page is Organized** See page 91.
- **Customizing Which Columns Display** See page 93.
- **Flow Details** See page 94.
 - **Error Reasons for TCP Acceleration Failure** See page 98.
 - **Error Reasons for CIFS Acceleration Failure** See page 101.
 - **Error Reasons for SSL Acceleration Failure** See page 102.
 - **Error Reasons for Citrix Acceleration Failure** See page 105.
- **Resetting Flows to Improve Performance** See page 107.

How the Page is Organized

Click to select the filter.
Active filters are highlighted.

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

The screenshot shows the 'Flows' monitoring interface. At the top, there are filter controls for 'Flow Categories' (All, Pass-Through, Asymmetric, Stale Policy), 'Bytes Transferred' (Total, Last 5m), 'Flow Timing' (Active, Active + Ended Last 5min, Started Last 5min, Ended Last 5min, Ended), and 'Flow Characteristics' (Flows to Slow Devices). Below these are input fields for IP1, Port1, IP2, and Port2, along with 'Application' and 'Traffic' dropdowns. A 'Filter' dropdown menu is open, showing options like 'All', 'Optimized Traffic', 'Policy Drop', etc. A table below displays flow data with columns for Application, IP1, Port1, IP2, Port2, Detail, Reduction %, Bytes (Inbound/Outbound), Reduction %, Protocol, Outbound Tunnel, Start Time (UTC), and End Time (UTC). A callout points to the 'Detail' icon in the table row for the 'unassigned' flow.

Detail used by Silver Peak Support for troubleshooting

The following **filters** are available:

Parameter or Action	Definition
Flow Categories	The number after each option specifies how many flows fit the criteria <ul style="list-style-type: none"> All – all flows Pass-through – includes shaped and unshaped traffic that is not being optimized Asymmetric – if a flow's Receive tunnel and the Transmit tunnel are different, the flow is asymmetric. Stale Policy – the policy that was applied to this flow has been modified
Bytes Transferred	Choose from Total or Last 5 minutes .
Flow Timing	Options include Active , Active + Ended Last 5min , Started Last 5min , Ended Last 5min , and Ended .
IP1 (2) / Port1 (2)	The IP address of an endpoint(s) that you want to use as a filter: <ul style="list-style-type: none"> Entering a specific endpoint returns flows that have that endpoint. Entering 0 in any IP address's octet position acts as a wild card for that position. 0 in the Port field is also a wild card. The two IP address (and port) fields are independent of each other. In other words, you can filter on two separate endpoints.

Parameter or Action	Definition (Continued)
Application	Select which standard or user-defined application (or application group) to use as a filter criteria. The default value is All .
Traffic	Select the type of traffic connections you want to retrieve: <ul style="list-style-type: none"> • All – all optimized and pass-through traffic. • Policy Drop – traffic with a Set Action of Drop in the Route Policy • Optimized Traffic – the sum of all optimized traffic. That is, all tunneled traffic. • Pass-through Shaped – all unoptimized, shaped traffic. • Pass-through Unshaped – all unoptimized, unshaped traffic. • [a named Tunnel] – that specific tunnel's optimized traffic.
Protocol	Select from the list. The default value is All .
VLAN Id	Enter only the integer value for the VLAN Id.
Internet Service	For sorting by domain, country, or city.
Max Flows	The upper limit depends on what browser you're using.
Reset Flows	Resetting the flow kills it and restarts it. It is service-affecting. For more information, see "Resetting Flows to Improve Performance" on page 107 .
Reclassify Flows	Reclassifying the flow is not service-affecting. If a policy change makes a flow stale or inconsistent, then reclassifying makes a best-effort attempt to conform the flow to the change. If the flow can't be successfully "diverted" to this new policy, then an Alert asks if you want to Reset.

Customizing Which Columns Display

Following are some customization guidelines:

- The default set of columns includes the following:

Application	Detail	Protocol
IP1	Inbound Reduction %	Outbound Tunnel
PORT1	Inbound Bytes	Start Time (UTC)
IP2	Outbound Bytes	End Time (UTC)
PORT2	Outbound Reduction %	

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

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

- Customizations persist** across sessions and across users. For a given appliance, all users see the same columns.
- When you **Export** the data, **all default and possible custom columns are included** in the .csv file.
- Customize** and **Export** functions are accessible to all users.

◆ To customize the screen display

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

The screenshot shows the 'Flows' monitoring interface. At the top, there are filters for Flow Categories (All, Pass-Through, Asymmetric, Stale Policy), Bytes Transferred (Total, Last 5m), Flow Timing (Active, Active + Ended Last 5min, Started Last 5min, Ended Last 5min, Ended), and Flow Characteristics (Flows to Slow Devices). On the right, there are input fields for IP1, Port1, IP2, Port2, Application, Traffic, Protocol, VLAN Id, Internet Service, and Max Flows. Below these are buttons for 'Reset Flows', 'Reclassify Flows', 'Customize', and 'Export'. The 'Customize' button is circled in red. A dialog box titled 'Customize Current Flows Table' is open, showing a list of columns with checkboxes. The background table shows flow data with columns for Application, IP1, Port1, and Outbound Bytes.

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

Flow Details

Silver Peak Support uses the **Flow Details** for troubleshooting.

The screenshot shows the 'Flows' management interface. At the top, there are filters for Flow Categories (All, Pass-Through, Asymmetric, State Policy), Bytes Transferred (Total, Last 5m), IP1, Port1, IP2, Port2, Application, Traffic, Protocol, VLAN id, Internet Service, and Max Flows. Below the filters is a table of flows. The table has columns for Application, IP1, Port1, IP2, Port2, Detail, Reduction %, Inbound Bytes, Outbound Bytes, Reduction %, Protocol, Outbound Tunnel, Start Time (UTC), and End Time (UTC). A red circle highlights the 'Details' icon in the 'Details' column for the flow with IP1: 10.1.153.10, Port1: 5781, IP2: 10.1.154.10, and Port2: 38819.

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

The screenshot shows the 'Flow details for IP1: 10.1.153.10 Port1: 5781 IP2: 10.1.154.10 and Port2: 38819' dialog box. It has three tabs: General, TCP Info, and NAT Info. The General tab is selected and shows the following information:

Stats		Route	
Outbound Ratio	2.00	Map Name	map1
Inbound Ratio	0.45	Priority in Map	default
Outbound LAN bytes	105,694,649,879	Configured Tx Action	pass-through
Outbound WAN bytes	52,924,154,364	Tx Action	auto_tun_10.1.154.20_to_10.1.153.20
Inbound LAN bytes	368,564,620	Rx Action	auto_tun_10.1.154.20_to_10.1.153.20
Inbound WAN bytes	816,772,366	Tx Reason	Auto-opt
Outbound LAN pkts	74,762,988	Application	3par
Outbound WAN pkts	39,638,631	Protocol	tcp
Inbound LAN pkts	8,840,836	Using Stale Map Entry	No
Inbound WAN pkts	37,002,658	Flow Direction	Outbound
Flow Up Time	7d 5h 43m 57.906s	Flow Redirected From	
Flow ID	2074	Auto-opt Status	Auto Routed
Active	Yes	Auto-opt Transit Node 1	10.1.154.20
TCP Flow Context	2074	Auto-opt Transit Node 2	10.1.153.20
Is Flow Queued For Reset	No	Auto-opt Transit Node 3	0.0.0.0
		Auto-opt Transit Node 4	0.0.0.0
		LAN-side VLAN	None

Optimization		QoS	
Map Name	map1	Map Name	map3
Priority in Map	default	Priority in Map	10030
TCP Acceleration Configured	Yes	Traffic Class	4
TCP Acceleration Status	Yes	LAN DSCP Configured	trust-lan
TCP Acceleration Info		WAN DSCP Configured	trust-lan
TCP Asymmetric	No	Using Stale Map Entry	No
Proxy Remote Acceleration	No		
CIFS Acceleration Configured	No		
CIFS Acceleration Status	No		
CIFS Acceleration Info			
CIFS Server Side	No		
CIFS SMB Signed	No		
SRDF Acceleration Configured	No		
SRDF Acceleration Status	No		
SSL Acceleration Configured	No		
SSL Acceleration Status	No		
SSL Acceleration Reason			
Citrix Acceleration Configured	No		
Citrix Acceleration Status	No		
Citrix Acceleration Reason			
iSCSI Acceleration Configured	No		
iSCSI Acceleration Status	No		
Network Memory	Balanced		
Payload Compression	Yes		
Using Stale Map Entry	No		

At the bottom right of the dialog box, there are 'Refresh' and 'Close' buttons.

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

Field	Definition
Stats Information	
Outbound Ratio	For the outbound traffic, a ratio of the Outbound LAN bytes divided by the Outbound WAN bytes. When this ratio is less than 1.0, it's attributable to a fixed overhead (for WAN transmission) being applied to traffic that either is not compressible or consists of few packets.
Inbound Ratio	For the inbound traffic, a ratio of the Inbound WAN bytes divided by the Inbound LAN bytes.
Outbound LAN bytes	Total number of bytes received from the LAN [outbound traffic]
Outbound WAN bytes	Total number of bytes sent to the WAN [outbound traffic]
Inbound LAN bytes	Total number of bytes sent to the LAN [inbound traffic]
Inbound WAN bytes	Total number of bytes received from the WAN [inbound traffic]
Outbound LAN pkts	Total number of packets received from the LAN [outbound traffic]
Outbound WAN pkts	Total number of packets sent to the WAN [outbound traffic]
Inbound LAN pkts	Total number of packets sent to the LAN [inbound traffic]
Inbound WAN pkts	Total number of packets received from the WAN [inbound traffic]
Flow Up Time	The length of time that there has been a connection between the endpoints.
Flow ID	A unique number that the appliance assigns to the flow.
TCP Flow Context	Silver Peak uses this for debugging purposes.
Is Flow Queued for Reset	Whether the flow is waiting to be reset (after user input) or not.
Route	
Map Name	The name of the Route Policy.
Priority in Map	The number of the entry in the Route Policy that the flow matches.
Configured Tx Action	The SET action configured in the Route Policy's Tunnel field.
Tx Action	How the traffic is actually being transmitted. Usually, this is a tunnel name.
Rx Action	By what path or method the appliance is receiving this flow's traffic.
Tx Reason	Any error associated with packet transmission to the WAN.
Application	Name of the application to which that flow's traffic belongs.
Protocol	The flow's protocol.
Using Stale Map Entry	Whether or not the flow is using a policy entry that has been edited or deleted since the flow began.
Flow Direction	Whether the flow is Inbound or Outbound .
Flow Redirected From	The IP address of the appliance that's redirecting this flow to this appliance.
Auto-opt Status	Whether it matched a specific Route Policy or was Auto Routed.
Auto-opt Transit Node (1, 2, 3, 4)	The IP addresses of the hops between this appliance and the other end of the connection.
LAN-side VLAN	Specifies the VLAN tag (1 – 4095) or None.

Field	Definition (Continued)
Optimization	
Map Name	The name of the Optimization Policy.
Priority in Map	The number of the entry in the Optimization Policy that the flow matches.
TCP Acceleration Configured	Whether or not TCP acceleration is configured in the Optimization Policy.
TCP Acceleration Status	Whether TCP is accelerated [Yes] or not [No].
TCP Acceleration Info	The reason that a TCP flow is not accelerated. For a list of error codes, see “Error Reasons for TCP Acceleration Failure” on page 98.
TCP Asymmetric	When the answer is YES , the Silver Peak appliance is able to intercept connection establishment in only one direction. As a result, this flow is not accelerated. When this happens, it indicates that there is asymmetric routing in the network.
Proxy Remote Acceleration	Which side is accelerating the flow
CIFS Acceleration Configured	Whether or not CIFS acceleration is configured in the Optimization Policy [Yes/No]
CIFS Acceleration Status	Whether CIFS is accelerated [Yes] or not [No].
CIFS Acceleration Info	The reason that a CIFS flow is not accelerated. For a list of error codes, see “Error Reasons for CIFS Acceleration Failure” on page 101.
CIFS Server Side	[Yes/No] If Yes , then this is the server side and the appliance is not accelerating (only the client side accelerates).
CIFS SMB Signed	Specifies whether or not the CIFS traffic is SMB-signed by the server: <ul style="list-style-type: none"> • Yes means it was signed. If that's the case, then the appliance was unable to accelerate any CIFS traffic. • No means it wasn't signed. If that's the case, then server requirements did not preclude CIFS acceleration. • Overridden means that SMB signing is ON and the appliance overrode it.
SRDF Acceleration Configured	Whether or not SRDF acceleration is configured in the Optimization Policy [Yes/No]
SRDF Acceleration Status	Whether SRDF is accelerated [Yes] or not [No].
SSL Acceleration Configured	Whether or not SSL acceleration is configured in the Optimization Policy [Yes/No]
SSL Acceleration Status	If a certificate has been appropriately installed via the Orchestrator, then SSL traffic can be deduplicated. Whether SSL is accelerated [Yes] or not [No].
SSL Acceleration Reason	The reason that an SSL flow is not accelerated. For a list of error codes, see “Error Reasons for SSL Acceleration Failure” on page 102.
Citrix Acceleration Configured	Whether or not Citrix cgp (gateway) or ica protocol acceleration is configured in the Optimization Policy [Yes/No]
Citrix Acceleration Status	Whether Citrix is accelerated [Yes] or not [No].

Field	Definition (Continued)
Citrix Acceleration Reason	The reason that a Citrix flow is not accelerated. For a list of error codes, see “Error Reasons for Citrix Acceleration Failure” on page 105 .
iSCSI Acceleration Configured	Whether or not iSCSI protocol acceleration is configured in the Optimization Policy [Yes/No]
iSCSI Acceleration Status	Whether iSCSI is accelerated [Yes] or not [No].
Network Memory	There are four Network Memory settings: <ul style="list-style-type: none"> • Maximize Reduction — optimizes for maximum data reduction at the potential cost of slightly lower throughput and/or some increase in latency. It is appropriate for bulk data transfers such as file transfers and FTP where bandwidth savings are the primary concern. • Minimize Latency — ensures that no latency is added by Network Memory processing. This may come at the cost of lower data reduction. It is appropriate for extremely latency-sensitive interactive or transactional traffic. It is also appropriate if WAN bandwidth saving is not a primary objective, and instead it is desirable to fully utilize the WAN pipe to increase LAN-side throughput. • Balanced — This is the default setting. It dynamically balances latency and data reduction objectives and is the best choice for most traffic types. • Disabled — No Network Memory is performed.
Payload Compression	Whether or not payload compression is turned on.
Using Stale Map Entry	Whether or not the flow is using a Route Policy entry that has been edited or deleted since the flow began.
QoS Information	
Map Name	The name of the QoS Policy.
Priority in Map	The number of the entry in the QoS Policy that the flow matches.
Traffic Class	The number of the traffic class assigned by the QoS to the flow, based on the MATCH conditions satisfied.
LAN DSCP Configured	The LAN DSCP marking that the QoS policy assigned to the flow, based on the MATCH conditions satisfied.
WAN DSCP Configured	The WAN DSCP marking that the QoS policy assigned to the flow, based on the MATCH conditions satisfied.
Using Stale Map Entry	Whether or not the flow is using a policy entry that has been edited or deleted since the flow began.

Error Reasons for TCP Acceleration Failure

When a flow has an acceleration failure, you can find the reason by clicking the selected flow's **Detail** icon.

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

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

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

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

Error Reasons for CIFS Acceleration Failure

When a flow has an acceleration failure, you can find the reason by clicking the selected flow's **Detail** icon.

Following is a list CIFS reason codes. They use the following format:

- **No [reason]** — The connection is not accelerated, and the “reason string” explains why not.
- **Yes [reason]** — The connection is partially accelerated, and the “reason string” explains why the connection is not fully accelerated.
- **Yes** — The connection is fully accelerated.

Yes/ No	Reason Text	Description
No	CIFS optimization is disabled in the Optimization Policy	CIFS is disabled in the optmap.
No	SMB signing is required by the server	SMB signing is enforced by the server, and this requirement precludes optimization.
No	SMB version 2 is enforced by the client	SMB version 2 protocol is enforced by the client, and this requirement precludes optimization.
No	The flow limit for CIFS optimization has been exceeded	Maximum flow limit reach for CIFS optimized flows.
Yes	Sub-optimal read-write optimization - Non standard server	Sub-optimal read/write optimization due to non-standard server. For example, Windows XP cannot process more than 10 simultaneous outstanding requests.
Yes	Metadata optimization disabled - NTNOTIFY failure	Metadata optimization is disabled due to change notification failure.
Yes	Metadata optimization disabled - OPEN failure	Metadata optimization is disabled because proxy cannot open the root share. To resolve, check the root share permissions.
Yes	Metadata optimization disabled - Unsupported Dialect	Endpoints are using an unsupported CIFS dialect. To resolve, upgrade the CLIFS client/server.
Yes	Metadata optimization disabled - Unsupported Server	Unsupported CIFS server, like UNIX/Samba. To resolve, switch to standard servers like Windows/NetApp..
Yes	Metadata optimization disabled - Unsupported Client	Unsupported CIFS client, like UNIX/smbclient. To resolve, switch to standard clients like Windows/Mac.

Error Reasons for SSL Acceleration Failure

When a flow has an acceleration failure, you can find the reason by clicking the selected flow's **Detail** icon.

Silver Peak supports:

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

Silver Peak appliances support the following:

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



Note To deduplicate SSL (Secure Socket Layer) traffic, appliances must have a valid SSL certificate and key. For information about installing SSL certificates and keys, see *"Adding SSL Certificates and Keys for Deduplication"* on page 16.

Following is a list of the reasons you may receive a failure message for SSL acceleration, along with a possible resolution.

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

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

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

Error Reasons for Citrix Acceleration Failure

When a flow has an acceleration failure, you can find the reason by clicking the selected flow's **Detail** icon.

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

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

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

Resetting Flows to Improve Performance

If you filter for **Asymmetric** flows, it displays flows that don't have the same inbound and outbound paths. To accelerate TCP flows, both directions of the flow must be in the same tunnel:

- TCP connections are not accelerated if they already exist when the tunnel comes up or when the appliance reboots.
- Pass-through connections are not tunnelized (or subsequently accelerated) if they already exist when a new tunnel is added, or when an ACL is added or edited.

Unaccelerated TCP flows can be reset so that the connection end-points have the opportunity to re-establish the flows. When the flows reconnect, the appliance recognizes them as new and accelerates them. The time it takes to reset a flow may vary, depending on the traffic activity.

However, if the network itself is asymmetric, then you'll need to set up flow redirection among peer appliances.



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



Tip For information about configuring the appliance to *automatically* reset TCP flows, see the Advanced TCP Options in the Optimization Policy.

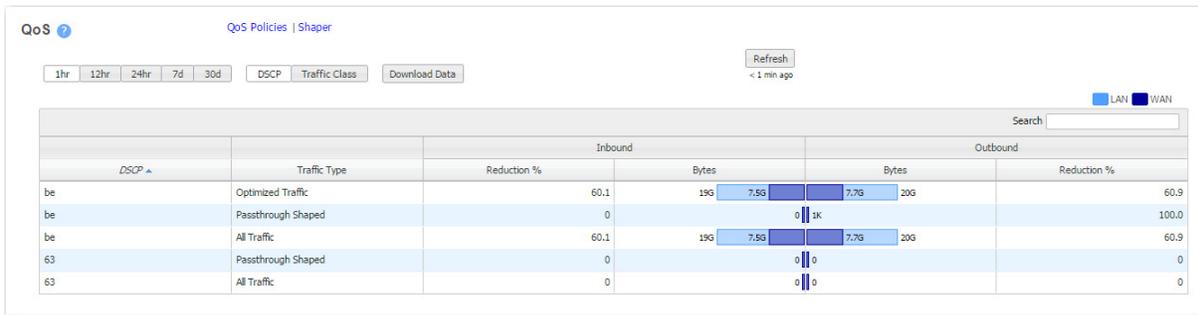
QoS Statistics

Monitoring > [Statistics] QoS

The **QoS** page summarizes optimized traffic on the basis of traffic class and/or WAN DSCP markings.

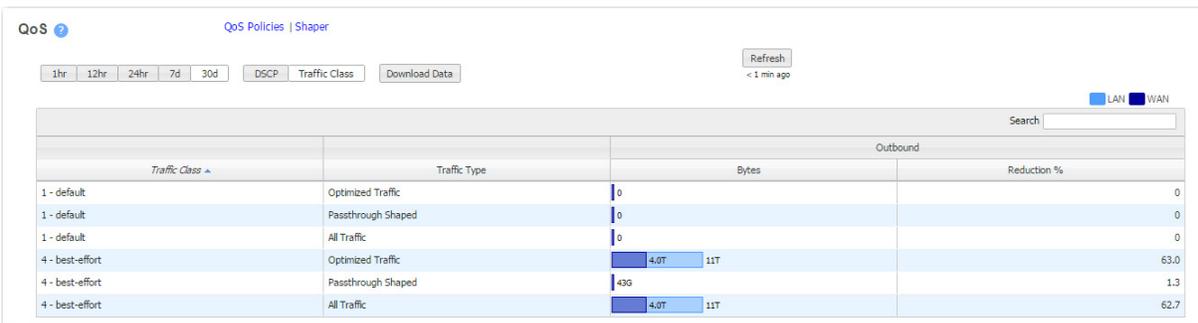
DSCP

The **DSCP Reduction** chart shows how bytes were sent in each DSCP class for the selected time period, as well as how much the data were reduced by the appliance (that is, LAN versus WAN bytes).



Traffic Class

The **Traffic Class Reduction** chart shows bytes sent and reduction for each QoS traffic class.

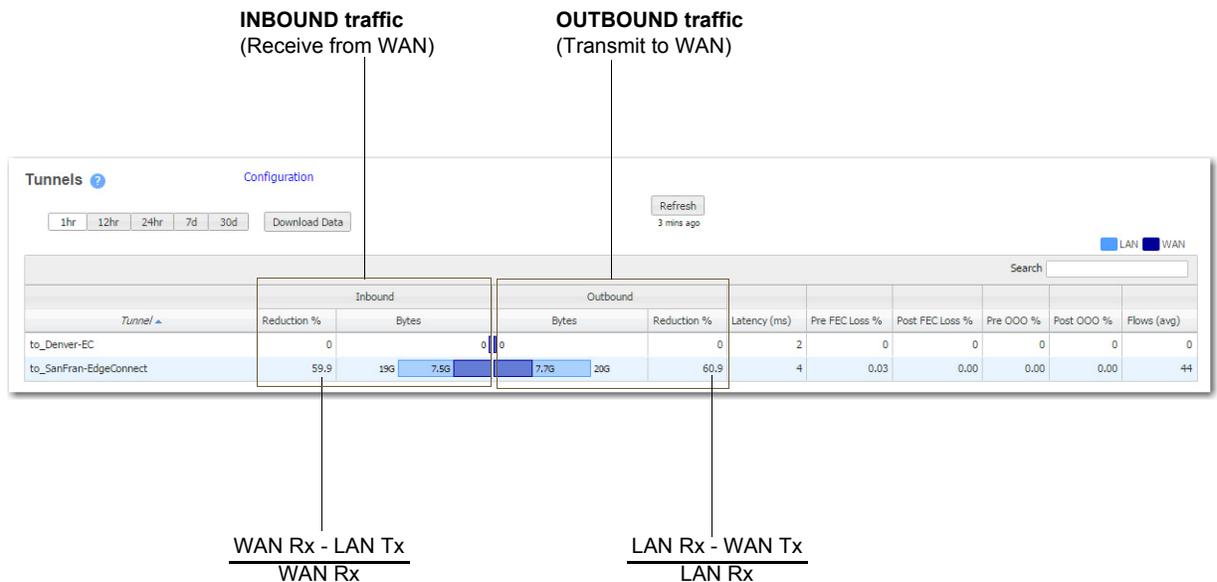


Tunnel Statistics

Monitoring > [Statistics] Tunnels

The **Tunnels** chart shows how bytes were sent on each tunnel for the selected time period, as well as how much the data were reduced by the appliance (that is, LAN versus WAN bytes).

It also shows statistics for latency, loss, out-of-order packets, and average number of flows (sampled every minute).



Section	Definition
Latency (ms)	Duration, in milliseconds, of the round trip latency
Pre FEC Loss %	Percentage of packets lost before enabling Forward Error Correction (FEC).
Post FEC Loss %	Percentage of packets lost after enabling Forward Error Correction (FEC).
Pre OOO %	Percentage of out-of-order packets before enabling Packet Order Correction (POC).
Post OOO %	Percentage of out-of-order packets after enabling Packet Order Correction (POC).
Flows (avg)	Average number of flows during the selected time period

How to Fine-tune Packet Correction

Enabling **Forward Error Correction (FEC)** in the tunnel configuration can sometimes result in the creation of additional **Out-of-Order Packets (OOOP)**.

To view the performance after enabling **FEC**, do either of the following:

- On the **Monitoring > Tunnels** page and review the Stats for **Pre POC Out-of-Order** and **Post POC Out-of-Order**.
- Access the **Monitoring > Charts** page, select the tunnel from **Traffic**, and review its **Out-of-Order** chart.

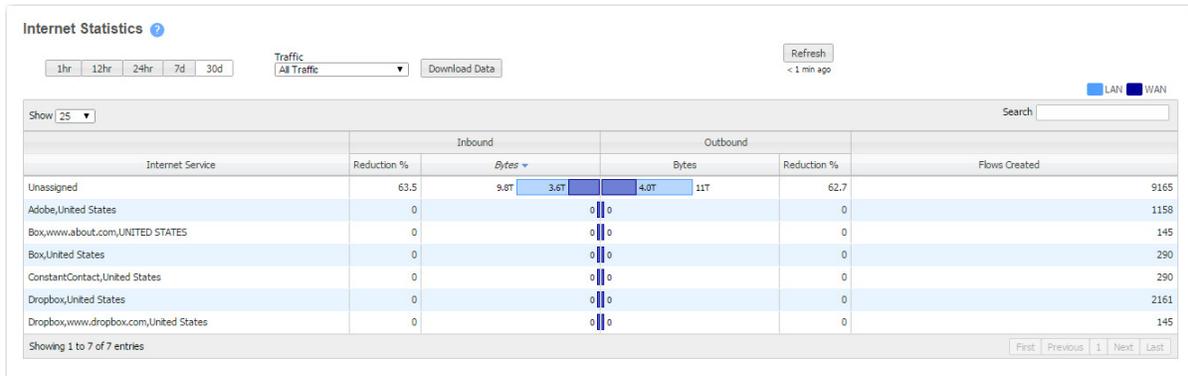
If necessary, adjust as follows:

- 1 If out-of-order packets exist, then you'll need to try another **Reorder Wait** time for the tunnel in question.
- 2 Go to **Configuration > Tunnels**, and in the row of the tunnel in question, click **Advanced Tunnel**.
- 3 At first, set the **Reorder Wait** time to **10ms**, and save the configuration.
- 4 Return to the stats page(s) to see if the out-of-order packets have been eliminated.
- 5 If there are still out-of-order packets, then go back to the tunnel configuration and increase the **Reorder Wait** time.
- 6 Repeat Steps 5 and 6 until there are no more out-of-order packets.

Internet Statistics

Monitoring > [Statistics] Internet

The **Internet Statistics** chart shows which internet services are sending the most data.



NetFlow Statistics

Monitoring > [Statistics] Netflow

The **NetFlow** page summarizes the data exported to NetFlow collectors.

The appliance exports flows against two virtual interfaces -- **sp_lan** and **sp_wan** -- that accumulate the total of LAN-side and WAN-side traffic, regardless of physical interface.

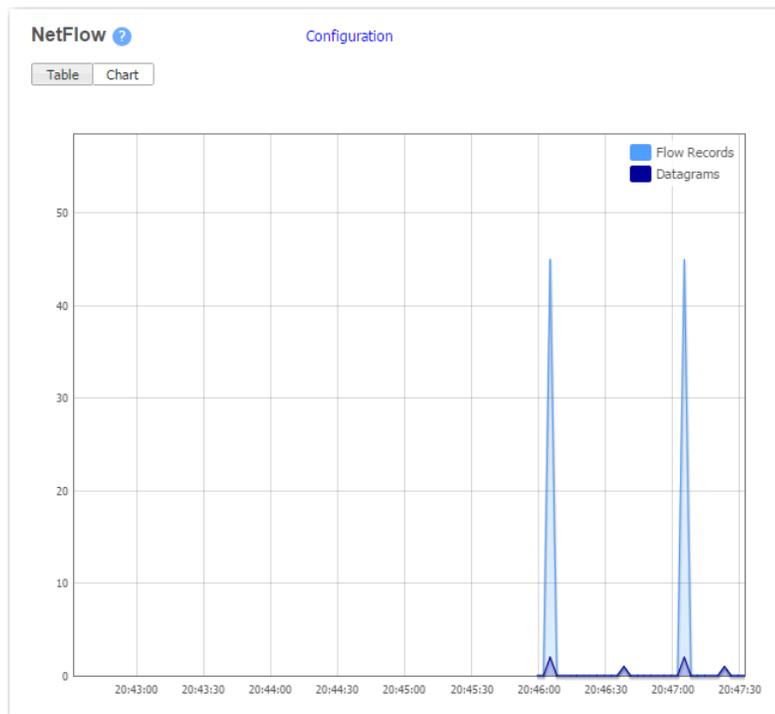
Table



The screenshot shows the NetFlow statistics page in table view. It includes a 'Refresh' button labeled '1 min ago'. The statistics are as follows:

Category	Value
Flow Records Exported:	273067
Datagrams Exported:	12090

Chart



Interface Statistics

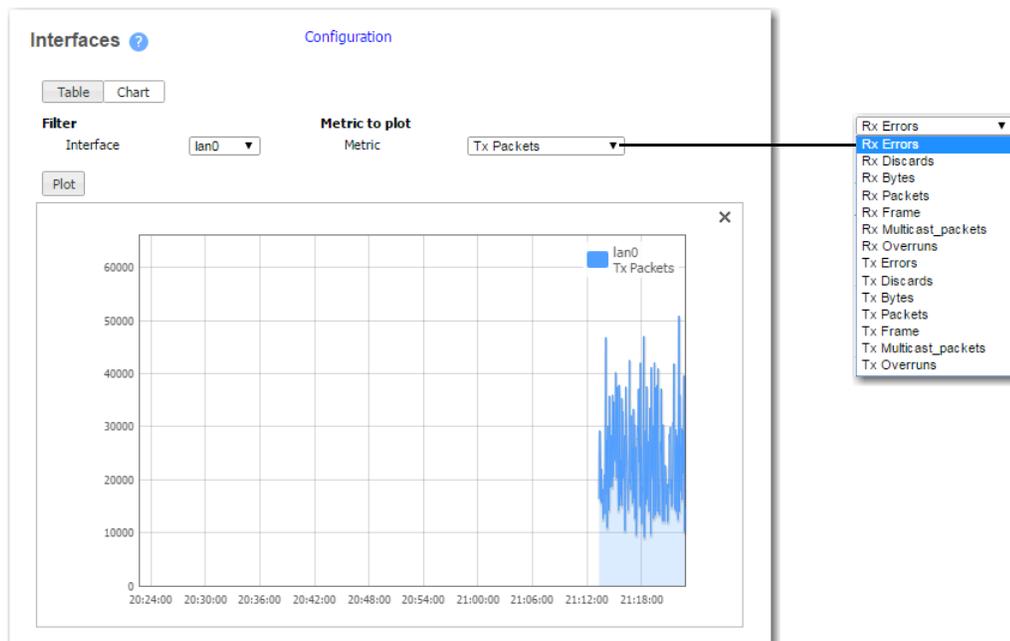
Monitoring > [Statistics] Internet

The **Interfaces** report displays performance data for the LAN, WAN, and management interfaces.

Table

Interface	Bytes		Pkts		Discard Pkts		Error Pkts		Overrun Pkts		Rx MCast Pkts	Tx Carrier Pkts	Rx Frame Pkts	Tx Collision Pkts
	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx	Rx	Tx				
mgmt0	487,197	3,415,771	2,024	3,619	0	0	0	0	0	0	67	0	0	0
eth2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
lo	5,974,607	5,974,607	10,608	10,608	0	0	0	0	0	0	0	0	0	0
bwi0	275,649,995	478,873,990	255,345	380,532	0	0	0	0	0	0	0	0	0	0
lan0	154,970,402	328,842,641	138,036	244,141	0	0	0	0	0	0	10	0	0	0
wan0	124,206,384	149,997,544	117,293	136,422	0	0	0	0	0	0	10	0	0	0
wan1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
mgmt1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
lan1	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Chart



Network Receive Statistics

Field	Definition
Rx Bytes	Number of bytes received inbound from the WAN side
Rx Pkts	Number of packets received inbound from the WAN side, including all packets that were either discarded, contained errors, arrived too quickly for the hardware to receive, or were frame or mcast packets,
Rx Discard Pkts	Number of input packets selected to be discarded even though no errors are found.
Rx Error Pkts	Number of input packets that contained errors.
Rx Overrun Pkts	Number of times the receiver hardware was unable to hand a received packet to a hardware buffer because the rate exceeded the receiver's ability to handle the data.
Rx MCast Pkts	Number of multicast packets received.
Rx Frame Pkts	Number of packets received incorrectly having a CRC error and a non-integer number of octets. On a LAN, this is usually the result of collisions or a malfunctioning Ethernet device.

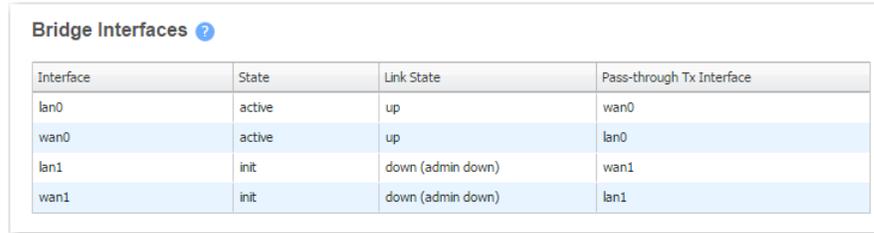
Network Transmit Statistics

Field	Definition
Tx Bytes	Number of bytes transmitted outbound toward the WAN side
Tx Pkts	Number of packets transmitted outbound toward the WAN side, including all packets that were either discarded, contained errors, were overrun, had collisions, or were dropped because the interface detection link is lost.
Tx Discard Pkts	Number of output packets selected to be discarded even though no errors are found.
Tx Error Pkts	Number of outbound packets that could not be transmitted because of errors.
Tx Overrun Pkts	Number of times the transmitter hardware was unable to hand a transmitted packet to a hardware buffer because the rate exceeded the transmitter's ability to handle the data.
Tx Carrier Pkts	Number of packets dropped because the interface detection link is lost.
Tx Collision Pkts	Number of output collisions detected on this interface.

Bridge Interfaces

Monitoring > [Status] Bridge Interfaces

When the appliance is in Bridge mode, this table lists how the data traffic spans the LAN and WAN interfaces.



The screenshot shows a table titled "Bridge Interfaces" with a help icon. The table has four columns: Interface, State, Link State, and Pass-through Tx Interface. The data is as follows:

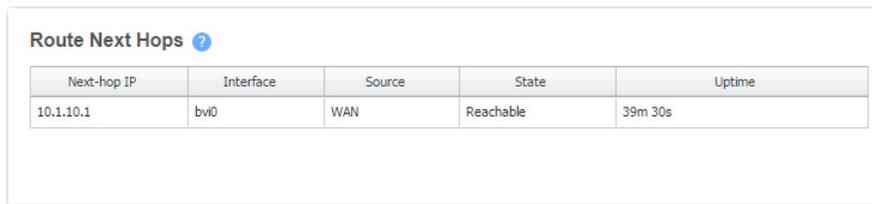
Interface	State	Link State	Pass-through Tx Interface
lan0	active	up	wan0
wan0	active	up	lan0
lan1	init	down (admin down)	wan1
wan1	init	down (admin down)	lan1

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

Route Next Hops

Monitoring > [Status] Route Next Hops

The **Route Next Hops** page displays the state of each management, WAN, and LAN next-hop.



The screenshot shows a web interface titled "Route Next Hops" with a help icon. Below the title is a table with the following data:

Next-hop IP	Interface	Source	State	Uptime
10.1.10.1	bvi0	WAN	Reachable	39m 30s

The page displays the following statistics:

Field	Definition
Next-hop IP	IP address of the router to which the Silver Peak appliance sends datapath (or management) traffic
Interface	Logical port associated with the Next-hop IP
Source	Direction of the next-hop router, relative to the appliance
State	There are four possible states: <ul style="list-style-type: none"> • Initializing • Reachable • Unreachable • Test disabled [when appliance is in Bypass mode]
Uptime	Amount of time that the next-hop router has been reachable

VRRP Status

Monitoring > [Status] VRRP

This report summarizes the appliance's configuration and state when deployed with **Virtual Router Redundancy Protocol (VRRP)**.

Group ID	Interface	State	Admin	Virtual IP	Advertise... Timer	Priority	Preemption	Master IP	Virtual MAC Address	State Uptime	Master State	IP Address Owner
1	wan0	init	Up	10.10.10.1	1	128	<input checked="" type="checkbox"/>	0.0.0.0	00:00:00:00:00:00	0 days 10 secs	0	<input type="checkbox"/>

In an out-of-path deployment, one method for redirecting traffic to the Silver Peak appliance is to configure VRRP on a common virtual interface. The possible scenarios are:

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

DEFINITIONS (alphabetically)

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

- **Virtual MAC address** = MAC Address that the VRRP instance is using. On an NX appliance, this is in 00-00-5E-00-01-**{VRID}** format. On virtual appliances, the VRRP instance uses the interface's assigned MAC Address (for example, the MAC address that the hypervisor assigned to **wan0**).

Flow Redirection Statistics

Monitoring > [Status] Flow Redirection

To accelerate a TCP connection, the local and the remote appliances need to see both directions of that same TCP flow. In a high-availability deployment scenario--where a site has multiple appliances--it's possible for the inbound and outbound directions of a TCP flow to traverse different appliances at that site, thereby preventing acceleration.

Flow redirection ensures that, for a given TCP connection, one of these appliances owns and sees both directions of traffic. At a given site, flow redirection moves traffic among appliances that you assign to a **cluster**:

- A cluster can contain just one appliance (in which no redirection occurs) or several appliances (in which redirection may occur between different pairs).
- All the appliances in a cluster are peers.
- A flow has only **one** owner among peers, and all peers receiving that flow redirect it **to** the owner.

Table View

For each **mgmt1** IP address in the cluster, the **Stats** area summarizes the control packets that keep the connection open to a peer appliance.

These numbers are **cumulative** for all redirected flows, whether they're active or terminated.

The screenshot shows the 'Flow redirection' configuration page. At the top, there is a 'Refresh' button and a timestamp '< 1 min ago'. Below this, there are tabs for 'Table' and 'Chart'. The main content area is divided into two sections: 'Flows' and 'Messages'.

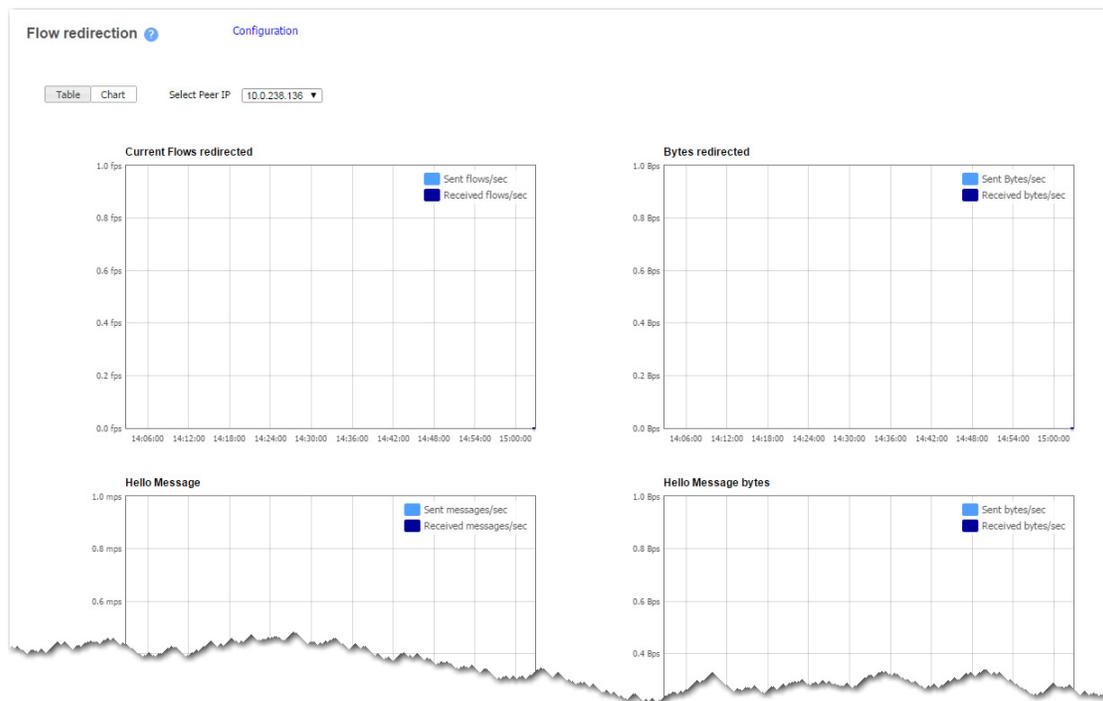
Flows Section:

- Radio buttons for 'Bytes/Pkts/Msgs' (selected) and 'MBytes/KPkts/KMsgs'.
- Table with columns: Peer IP, State, Current Flows, Packets, Bytes (under 'Flows redirected from'); Current Flows, Packets, Bytes (under 'Flows redirected to').
- Row for Peer IP 10.0.238.136, State Unreachable, with all numerical values set to 0.

Messages Section:

- Table with columns: Peer IP, Sent Msgs, Sent bytes, Received Msgs, Received bytes (under 'Hello'); Sent Msgs, Sent bytes, Received Msgs, Received bytes (under 'Redirection').
- Row for Peer IP 10.0.238.136, with all numerical values set to 0.

Chart View



Statistics

The **Flows** stats summarize the number of Packets and Bytes redirected up to the present time.

- The number of flows shown is current, not cumulative.
- This appliance **owns** all the flows redirected **from** a PeerID.
- **State** indicates whether or not a peer is reachable. If **unreachable**, there will be an alarm and recommended actions. Some causes could be:
 - The peer's **mgmt1** may be on a different subnet.
 - The cable between **mgmt1** interfaces may be defective or disconnected.
 - On the **Configuration - Interfaces** page, **mgmt1**'s values for **Admin**, **Status**, **DHCP**, and **Speed/Duplex** may not be the same on this appliance and the unreachable peer.
 - If **Status** is blank, there is no physical connection.

For each PeerIP address in a cluster, the **Messages** stats summarize:

- the control packets that keep open the TCP connection between two peers' **mgmt1** interfaces
- the number of requests to redirect flows.

SaaS Optimization

Monitoring > [Status] SaaS Optimization

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

SaaS Optimization Configuration

Refresh
2 mins ago

Show 25 Search

Application Name	Domain(s)	Subnet	Server IP	Advertised	RTT ▲	RTT Threshold	Ping Method	Ping Port	Location
Box	*.app.box.com, *.box.com, *.bo...	208.184.35.0/27	208.184.35.1	Yes	6 ms	10 ms	TCP	80	White Plains, UNITED STATES
Workday	workday.com	206.169.118.0/24	206.169.118.21	Yes	7 ms	10 ms	TCP	80	San Ramon, United States
Workday	workday.com	162.210.233.64/26	162.210.233.67	Yes	8 ms	10 ms	TCP	443	Sacramento, UNITED STATES
Workday	workday.com	96.46.149.0/24	96.46.149.31	Yes	8 ms	10 ms	TCP	80	Sacramento, United States
Workday	workday.com	173.226.103.0/24	173.226.103.1	Yes	8 ms	10 ms	TCP	80	San Ramon, United States
Workday	workday.com	209.177.168.0/24	209.177.168.4	No	26 ms	10 ms	TCP	80	Pleasanton, United States
Workday	workday.com	209.177.160.0/24	209.177.160.7	No	27 ms	10 ms	TCP	443	Pleasanton, United States
Workday	workday.com	209.177.160.0/20	209.177.160.7	No	28 ms	10 ms	TCP	443	Pleasanton, United States
Workday	workday.com	209.177.162.0/24	209.177.162.4	No	37 ms	10 ms	TCP	443	Pleasanton, United States
Jobvite	careers.jobvite.com, www.jobvi...	104.130.145.184/32	104.130.145.184	No	44 ms	10 ms	TCP	80	San Antonio, United States
Adobe	adobe.com	173.240.103.64/26	173.240.103.65	No	60 ms	10 ms	TCP	80	Indianapolis, United States
Adobe	adobe.com	173.240.102.96/28	173.240.102.97	No	61 ms	10 ms	TCP	80	Indianapolis, United States
Salesforce	*.na3.force.com, *.salesforce.c...	96.43.144.0/20	96.43.144.15	No	61 ms	10 ms	TCP	443	San Francisco, United States
Salesforce	*.na3.force.com, *.salesforce.c...	96.43.144.0/22	96.43.144.15	No	61 ms	10 ms	TCP	443	San Francisco, United States
Salesforce	*.na3.force.com, *.salesforce.c...	136.146.66.0/23	136.146.66.2	No	61 ms	10 ms	TCP	80	San Francisco, United States
Salesforce	*.na3.force.com, *.salesforce.c...	136.146.208.0/23	136.146.208.32	No	61 ms	10 ms	TCP	443	San Francisco, United States
Adobe	adobe.com	173.240.108.112/28	173.240.108.113	No	62 ms	10 ms	TCP	80	Indianapolis, United States
Adobe	adobe.com	173.240.105.0/27	173.240.105.1	No	62 ms	10 ms	TCP	80	Indianapolis, United States
Adobe	adobe.com	173.240.105.32/28	173.240.105.33	No	62 ms	10 ms	TCP	80	Indianapolis, United States
Adobe	adobe.com	173.240.108.160/29	173.240.108.161	No	63 ms	10 ms	TCP	80	Indianapolis, United States
Adobe	adobe.com	173.240.110.128/28	173.240.110.129	No	64 ms	10 ms	TCP	80	Indianapolis, United States
Adobe	adobe.com	173.240.110.160/28	173.240.110.161	No	65 ms	10 ms	TCP	80	Indianapolis, United States
Workday	workday.com	209.177.166.0/24	209.177.166.8	No	66 ms	10 ms	TCP	80	Pleasanton, United States
Salesforce	*.na3.force.com, *.salesforce.c...	136.147.129.0/24	136.147.129.1	No	66 ms	10 ms	TCP	80	Indianapolis, United States
Salesforce	*.na3.force.com, *.salesforce.c...	136.147.133.0/24	136.147.133.11	No	67 ms	10 ms	TCP	443	Indianapolis, United States

Showing 1 to 25 of 40 entries

First Previous 1 2 Next Last

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

Orchestrator Reachability

Monitoring > [Status] Orchestrator Reachability

Use this page to **view**, **add**, and **delete** Orchestrator servers that have access to this appliance.

The screenshot shows the 'Orchestrator Reachability' page. At the top, there are buttons for 'Add', 'Test', and 'Refresh'. Below the table, there are 'Apply' and 'Cancel' buttons. The table has columns for 'Orchestrator IP', 'REST', 'SSH', 'HTTPS', 'Web Socket', and a status column with an 'X' icon. The data rows are as follows:

Orchestrator IP	REST	SSH	HTTPS	Web Socket	
10.0.234.151	Unreachable	Reachable	Unreachable	Unreachable	✘
10.0.236.146	Unreachable	Unreachable	Unreachable	Unreachable	✘
10.0.238.172	Reachable	Reachable	Reachable	Reachable	✘
10.0.239.69	Reachable	Reachable	Reachable	Reachable	✘
172.20.28.47	Unreachable	Unreachable	Unreachable	Unreachable	✘
172.20.28.56	Unreachable	Unreachable	Unreachable	Unreachable	✘
172.23.43.114	Unreachable	Unreachable	Unreachable	Unreachable	✘
172.23.43.197	Reachable	Unreachable	Reachable	Reachable	✘

- The table lists the protocols that the appliance uses to communicate with the Orchestrator.
- Clicking **Test** evaluates reachability to all listed Orchestrator servers.
- The possible states are **Reachable**, **In-Progress**, and **Unreachable**.
 - **Unreachable** indicates a problem in your network. Check your ports, firewalls, and deployment configuration.
 - HTTPS and Web Socket share Port 443.



Monitoring Alarms

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

In This Chapter

- **Understanding Alarms** See page 124.
- **Viewing Current Alarms** See page 135.

Understanding Alarms

This section defines the four alarm severity categories and lists all Silver Peak appliance alarms.

The **Alarms** page lists alarm conditions on the appliance. Each entry represents one current condition that may require human intervention. Because alarms are *conditions*, they may come and go without management involvement.

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

Categories of Alarms

The Appliance Manager categorizes alarms at four preconfigured severity levels: **Critical**, **Major**, **Minor**, and **Warning**.

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

Types of Alarms

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

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

Table 5-1 Silver Peak Appliance Alarms

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

Table 5-1 Silver Peak Appliance Alarms (Continued)

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

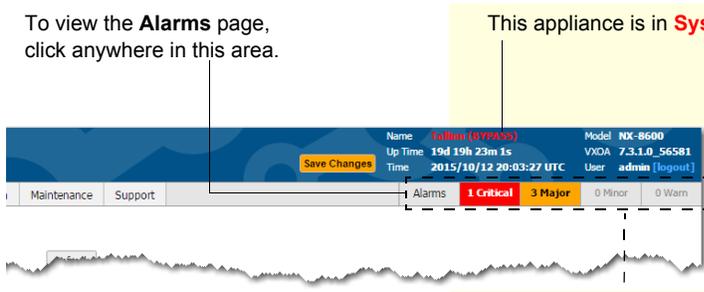
- a. The VX appliances are a family of virtual appliances, comprised of the VX-n000 software, an appropriately paired hypervisor and server, and a valid software license.
- b. If there is either a **LAN Next-Hop Unreachable** or **WAN Next-Hop Unreachable** alarm, resolve the alarm(s) immediately by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

Viewing Current Alarms

Most Silver Peak appliance alarms cannot be cleared by the user. Instead, the appliance generally corrects the alarm condition and clears the alarm by itself.

The alarm summary appears in the banner. You can view current alarms as follows:

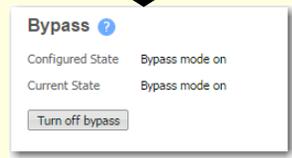
To view the **Alarms** page, click anywhere in this area.



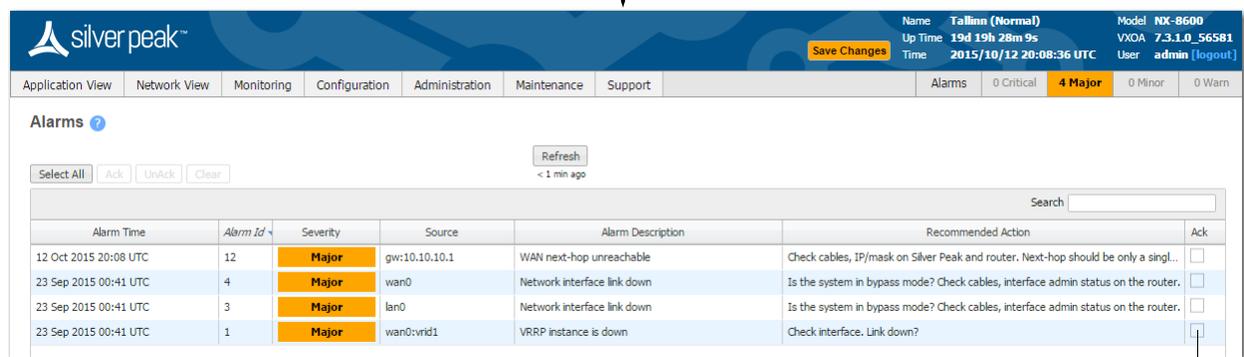
This appliance is in **System Bypass**.

To disable System Bypass:

- Go to the **Maintenance > System Bypass** page
- Click **Turn off bypass**.



...and the **Alarms** page displays.



For acknowledging alarms

The **Alarms** page displays the following information:

Field	Definition/Content
Seq No.	The sequential number of the alarm, based on the time the alarm raised.
Date/Time	The local date and time at the appliance's location, specified by a 24-hour clock.
Type	The type of alarm: <ul style="list-style-type: none"> Tunnel A tunnel-based alarm TC A traffic class-based alarm EQU An equipment-based alarm SW A code- or software-based alarm

Field	Definition/Content (Continued)
Severity	<p>The severity of the alarm, listed here in decreasing order of severity:</p> <ul style="list-style-type: none"> • Critical A critical alarm, such as "Tunnel Down" • Major A major alarm, such as "Disk out of Service" • Minor A minor alarm, such as "Disk Degraded" • Warning A warning, such as "Software Process Restart" • Info For Silver Peak debugging purposes. <p>These are purely related to alarms severities, not event logging levels, even though some of the naming conventions overlap. Events and alarms have different sources. Cleared alarms are at ALERT level in the Support > Log Viewer page.</p>
Source	Refers to the particular subsystem or equipment that is causing the alarm. For example, we can raise the tunnel-based alarm, "Tunnel Down", where the source would refer to a particular tunnel.
Description	A brief description of the alarm.
Recommended Action	Describes what action to take and, when appropriate, provides a link to the page where you need to do it.
Clear	<p>If a checkbox is accessible, a user can clear the alarm.</p> <p>To clear the alarm, click the Clear box, and click Apply. Once cleared, the row is removed and the content is viewable in the read-only page, Alarms - Log Viewer.</p>
Ack	Select Yes to acknowledge the alarm; Select No to remove acknowledgement.



Tip To view resolved alarms, go to the **Support > Log Viewer** page and filter for **Alarms**.

Administration Tasks

This chapter describes menus related to basic appliance administration.

In This Chapter

Basic Settings

- **Date and Time** See page 138.
- **Domain Name Server (DNS)** See page 139.
- **SNMP** See page 140.
- **Netflow** See page 143.
- **License & Registration** See page 144.
- **Management IP / Hostname** See page 145.
- **Silver Peak Cloud Portal** See page 146.

User Management

- **Users** See page 147.
- **Auth/RADIUS/TACACS+** See page 148.
- **Banners** See page 150.
- **Session Management** See page 151.

Date and Time

Administration > [Basic Settings] Date/Time

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

Server IP	Version	
172.20.20.37	3	X

- From the **Time Zone** list, select the appliance's geographical location.
- To manually configure, select **Manual** and enter the **Date** [YYYY/MM/DD] and **Time** [HH:MM:SS] (based on a 24-hour clock).
- To use an NTP server, select **NTP Time Synchronization**.
 1. Click **Add**.
 2. Enter the IP address of the server, and select the version of NTP protocol to use.

Domain Name Server (DNS)

Administration > [Basic Settings] DNS

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



The screenshot shows a configuration window titled "Domain Name Servers" with a help icon. It contains three input fields for "Primary DNS IP address", "Secondary DNS IP address", and "Tertiary DNS IP address". Below these is a section titled "Domain Names" with an "Add" button and a table with a header "Domain name" and a row containing "No Data Available". At the bottom are "Apply" and "Cancel" buttons.

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

SNMP

Administration > [Basic Settings] SNMP

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

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

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

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

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

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

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

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

Loading SNMP MIBs

From Silver Peak's website, you can download the Standard and the Silver Peak proprietary MIBs (Management Information Base) files, for loading into whatever MIBs browser you're using:

- You can choose to install the Standard MIBs, the Silver Peak proprietary MIBs, or both.
- The Standard list and the Silver Peak file list share the same first three files. These are highlighted in **green** below.
- Because there are dependencies, you must load the files in a list in a specific sequence.
- If you choose to load both the Standard and the Silver Peak MIBs, load either list completely and then append the non-common files from the remaining list.

List of Silver Peak MIBs

Load these files in the following order:

- 1 **SNMPv2-SMI.txt**
- 2 **SNMPv2-TC.txt**
- 3 **SNMPv2-CONF.txt**
- 4 SILVERPEAK-SMI.txt
- 5 SILVERPEAK-TC.txt
- 6 SILVERPEAK-PRODUCTS-MIB.txt
- 7 SILVERPEAK-MGMT-MIB.txt

List of Standard SMIBs

Load these files in the following order:

- 1 `SNMPv2-SMI.txt`
- 2 `SNMPv2-TC.txt`
- 3 `SNMPv2-CONF.txt`
- 4 `RFC1155-SMI.txt`
- 5 `RFC1213-MIB.txt`
- 6 `SNMPv2-MIB.txt`
- 7 `SNMP-FRAMEWORK-MIB.txt`
- 8 `SNMP-MPD-MIB.txt`
- 9 `SNMP-TARGET-MIB.txt`
- 10 `SNMP-NOTIFICATION-MIB.txt`
- 11 `SNMP-USER-BASED-SM-MIB.txt`
- 12 `SNMP-VIEW-BASED-ACM-MIB.txt`

Netflow

Administration > [Basic Settings] Netflow

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

The screenshot shows the 'NetFlow Configuration' page. At the top, there is a title 'NetFlow Configuration' with a help icon and a link 'Monitor Statistics'. Below this, there are three settings: 'Flow Exporting Enabled' with a checked checkbox, 'Active Flow Timeout' with a text input field containing '1' and '(1..30) mins', and 'Traffic Type' with four radio buttons: 'Outbound WAN' (checked), 'Inbound WAN', 'Inbound LAN', and 'Outbound LAN'. A section titled 'Collectors' contains an 'Add' button and a table with columns 'Collectors', 'IP Address', and 'Port'. The table is currently empty, with the text 'No Data Available' centered below it. At the bottom of the form are 'Apply' and 'Cancel' buttons.

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

License & Registration

Administration > [Basic Settings] License & Registration

Use this page to enter and apply your appliance's **License Key** (only VX and VRX virtual appliances require licenses). Decoded details display below the input area.

License & Registration ?

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

Current License

Start Date	Tue Dec 3rd 2013
End Date	None
Portal Registration	OK
Model	VX-1000
Serial Number	000C291953A1

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

Account Name	<input type="text" value="Laine"/>
Account Key	<input type="text" value="aSeJqlp4yhczQUG"/>
Appliance Tag	Unassigned
Contact	ltammer@silver-peak.com

Registration

To access the *Silver Peak Unity Cloud Intelligence Service*, you must complete the **Registration** section.

This service is required if you want to enable SaaS optimization.

- 1 After completing registration, go to **Configuration > SaaS Optimization** and select **Enable SaaS Optimization**.
- 2 On the same page, choose whether or not to decrypt the SSL traffic, and if so, which substitute certificate signing method to use.

This service is also required for the EdgeConnect product line.

Management IP / Hostname

Administration > [Basic Settings] Management IP/Hostname

Use this page to configure the **mgmt0** management interface and the appliance hostname.

Management Interface (mgmt0) ?

Appliance Hostname:

DHCP

IP Address/Mask	Next-hop IP
10.0.238.135/26	10.0.238.129

Static

IP Address/Mask	Next-hop IP
<input type="text"/>	<input type="text"/>

- The browser accesses the appliance via the **mgmt0** IP address.
- When deployed in Server mode, a virtual appliance uses **mgmt0** for management *and datapath* traffic.
- If you configure flow redirection between appliances, the **mgmt1** interfaces need to be in a separate subnet from the **mgmt0** interfaces. (Flow redirection is a local way to remove flow asymmetry by redirecting packets so that one appliance owns both directions of traffic.)
- A **hostname** is subject to the following constraints:
 - Maximum number of characters = 24
 - Allowable characters: A-Z, a-z, 0-9, dash (-)

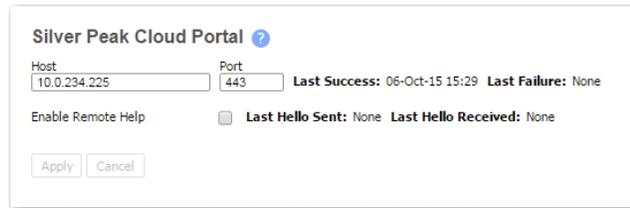
Best Practices

- Assign static IP addresses to management interfaces to preserve their reachability.
- Use different subnets for **mgmt0** and datapath interfaces.

Silver Peak Cloud Portal

Administration > [Basic Settings] Silver Peak Cloud Portal

This page lists the default host and port for the **Silver Peak Cloud portal**.



The screenshot shows a configuration window titled "Silver Peak Cloud Portal" with a help icon. It contains two input fields: "Host" with the value "10.0.234.225" and "Port" with the value "443". To the right of the port field, it displays "Last Success: 06-Oct-15 15:29" and "Last Failure: None". Below these fields is a checkbox labeled "Enable Remote Help" which is currently unchecked. To its right, it shows "Last Hello Sent: None" and "Last Hello Received: None". At the bottom of the window are "Apply" and "Cancel" buttons.

Connectivity to the portal is required for *Silver Peak Unity Cloud Intelligence Services* to work properly. These services include:

- SaaS optimization (and installing any subordinate CA certificate on the appliance)
- Displaying internet statistics about internet services
- CPX licensing
- EdgeConnect licensing

To allow Silver Peak Support to access this appliance's web user interface for one 24-hour period, select **Enable Remote Help for 24h**.

Users

Administration > [User Management] Users

Use this page to **create**, **edit**, and **delete** users.

Active Sessions ?

User Name	Line	From	Login Time	Idle Time
admin	web	10.0.238.172	2015/07/14 03:49:45 PDT	54s
admin	web	10.0.234.151	2015/07/14 03:49:45 PDT	40s
websocket	web	127.0.0.1	2015/07/14 03:50:17 PDT	11s
websocket	web	127.0.0.1	2015/07/14 03:50:53 PDT	45s
admin	web	172.23.43.59	2015/07/14 05:55:25 PDT	1s

User Accounts

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

- The **Active Sessions** table lists who is logged in to the appliance, and from where.
- The **User Accounts** table lists all users known to this appliance, whether or not their accounts are enabled.
- The system user names are **admin** and **monitor**.
 - They **CANNOT** be deleted.
 - You can only disable **monitor**.
 - You can, however, change each one's password.
- A user has either **admin** or **monitor** privileges:
 - **admin** capability allows the user to view and modify.
 - **monitor** capability allows the user to view only.

Auth/RADIUS/TACACS+

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

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

Authentication Order ⓘ

First: Local
Second: None
Third: None

Authorization Information

Map Order: Local Only
Default User: admin

RADIUS Servers

Add

Order	Server IP	Auth Port	Key	Timeout	Retries	Enabled	
1		1812		3	1	<input checked="" type="checkbox"/>	X

TACACS+ Servers

Add

Order	Server IP	Auth Port	Auth Type	Key	Timeout	Retries	Enabled	
No Data Available								

Apply Cancel

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

Authentication and Authorization

To provide authentication and authorization services, Silver Peak appliances:

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

Both RADIUS and TACACS+ are client-server protocols.

Appliance-based User Database

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

RADIUS

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

TACACS+

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

What Silver Peak recommends

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

Banners

Administration > [User Management] Banners

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

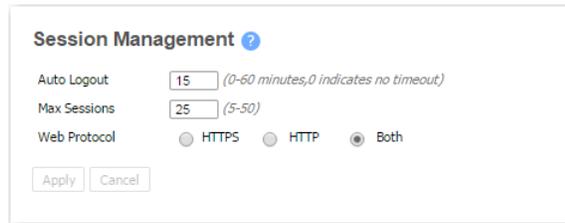
You can configure either, neither, or both.



Session Management

Administration > [User Management] Session Management

Use this page to configure access to the web server.



The screenshot shows a dialog box titled "Session Management" with a help icon. It contains three configuration fields: "Auto Logout" with a value of 15 and a note "(0-60 minutes, 0 indicates no timeout)"; "Max Sessions" with a value of 25 and a note "(5-50)"; and "Web Protocol" with three radio button options: "HTTPS", "HTTP", and "Both", where "Both" is selected. At the bottom are "Apply" and "Cancel" buttons.

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

Maintenance & Support

This chapter describes how to use these menus to perform various system maintenance tasks.

In This Chapter

Software & System Management

- **Viewing System Information** See page 154.
- **Upgrading Appliance Software** See page 155.
- **Backing Up and Restoring Appliance Configuration Files** See page 156.
- **Disk Management** See page 157.

Tools

- **Putting the Appliance in System Bypass Mode** See page 158.
- **Testing Network Connectivity with Ping and Traceroute** See page 159.
- **Capturing Packets for Analysis** See page 163.
- **Erasing Network Memory** See page 164.
- **Rebooting or Shutting Down an Appliance** See page 165.

Support: Technical Assistance

- **How to Contact Support** See page 166.
- **Logging In to the Support Portal** See page 167.
- **Viewing Logs** See page 168.
- **Configuring Logging Parameters** See page 169.
- **Managing Debug Files** See page 171.

Viewing System Information

Maintenance > [Software & System Management] System Information

This page displays information used to identify the appliance.

System Information	
Hostname	laine-vxb
Appliance ID	15046294
Model	VX-1000 206002001000 Rev 46839
System Status	Normal
Uptime	7d 3h 44m 48s
Date/Time	2015/10/07 23:12:06 Etc/UTC
Release	VXOA 7.3.1.0_56670
Serial Number	██████████

Upgrading Appliance Software

Maintenance > [Software & System Management] Software Upgrade

Use this page to manage up to two software images.

The screenshot shows the 'Software Upgrade' page with a help icon. It features three main sections: 'Installed Images', 'Install Image', and 'Install Options'. The 'Installed Images' section has 'Switch Partition' and 'Reboot' buttons above a table with columns for ID, Active, Next Boot, and Version. The 'Install Image' section has 'Local File' and 'URL' buttons above an 'Upload Image' button. The 'Install Options' section has three radio buttons: 'Install and Reboot' (selected), 'Install and set next boot partition', and 'Install Only', followed by an 'Install' button.

ID	Active	Next Boot	Version
1	yes	yes	8.0.2.0_58453
2	no	no	7.3.3.0_57797

- When you switch to the other partition, it only becomes active when you reboot.
- When loading a software image from your PC or a URL, you can:
 - **Install Only.** Loads the image into the inactive partition to store it.
 - **Install and set next boot partition.** Loads the image into the inactive partition. It becomes active at the next reboot.
 - **Install and Reboot.** Loads the image into the inactive partition, and immediately reboots to activate the new image.
- You may need to **downgrade** an appliance to be consistent with the rest of your network if you've purchased an appliance that comes with newer installed software or you've RMA'd an appliance. In that case, refer to the document, [How to Downgrade an Appliance](#).

Backing Up and Restoring Appliance Configuration Files

Maintenance > [Software & System Management] Backup/Restore

Backup creates a copy of the current configuration and saves it on the appliance.

After creating the file, you can download it, upload it, or restore it.

Backup/Restore ?

Backup Current Configuration

Filename

Restore

File Name	Active	Last Modified	
initial_20150527213727	true	2015/07/15 10:08:22	
initial_20150527213727.bak	false	2015/07/14 16:05:28	X
initial_20150416112244	false	2015/05/27 21:39:18	X
initial_20150416112244.bak	false	2015/05/27 14:38:07	X
initial	false	2015/04/16 11:29:12	X
initial.bak	false	2015/04/16 10:58:04	X

Disk Management

Maintenance > [Software & System Management] Disk Management

Physical appliances use RAID arrays with encrypted disks.

Disk failure results in a **critical alarm**.

Disk Management ?

Remove Insert

Slot ID	Pairing Slot ID	Status	Rebuild Status	Size(GB)	Serial Number	Removable
0	1	OK	COMPLETE	465	GTD400P6G2PGWD	yes
1	0	OK	COMPLETE	465	GTD400P6G2PNAD	yes
2	3	OK	COMPLETE	465	GTD400P6G2P0JD	yes
3	2	OK	COMPLETE	465	GTD400P6G2PEGD	yes
4	5	OK	COMPLETE	465	GTD400P6G2P54D	yes
5	4	OK	COMPLETE	465	GTD400P6G2PN1D	yes
6	7	OK	COMPLETE	465	GTD400P6G2PGDD	yes
7	6	OK	COMPLETE	465	GTD400P6G2P3GD	yes
8	9	OK	COMPLETE	465	GTD400P6G2P0ZD	yes
9	8	OK	COMPLETE	465	GTD400P6G2PLGD	yes
10	11	OK	COMPLETE	465	GTD400P6G2PNEED	yes
11	10	OK	COMPLETE	465	GTD400P6G2PN0D	yes
12	13	OK	COMPLETE	465	GTD400P6G2PN6D	yes
13	12	OK	COMPLETE	465	GTD400P6G2PNLD	yes
14	15	OK	COMPLETE	465	GTD400P6G2P7ZD	yes
15	14	OK	COMPLETE	465	GTD400P6G2NZWD	yes

Disk Layout

Follow this procedure when replacing a failed disk:

- 1 Log into your Support portal account, and click **Open a Self Service RMA** for disk replacement.
- 2 Complete the wizard, using the serial number of the appliance (not the disk).
- 3 After you receive the new disk, select the failed disk's row in the table and click **Remove**. This takes the disk off-line.
- 4 Physically remove the old disk from the appliance.
- 5 Physically insert the new disk.
- 6 In the table, select the new disk and click **Insert**. This prompts the software to discover the disk and put it online.

Virtual appliances display read-only data.

Disk Management ?

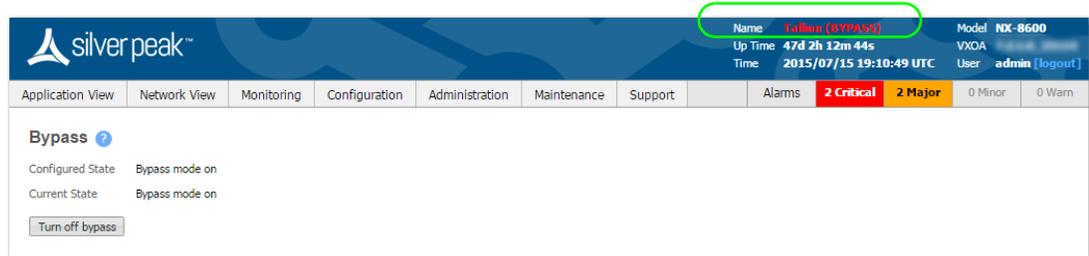
Remove Insert

Slot ID	Status	Rebuild Status	Size(GB)	Serial Number	Removable
0	OK		30		no
1	OK		0		yes

Putting the Appliance in System Bypass Mode

Maintenance > [Tools] System Bypass

In *system bypass mode*, the fail-to-wire (or fail-to-glass) card **DOES NOT** receive or process packets:



- In an in-line deployment (Bridge mode), the **lan** interface is physically connected to the **wan** interface.
- In an out-of-path deployment (Router/Server mode), the appliance is in an open-port state.

Fail-to-wire network interfaces mechanically isolate the appliances from the network in the event of a hardware, software, or power failure. This ensures that all traffic bypasses the failed appliance and maximizes up-time.



Note Virtual appliances don't support *bypass* mode.

Testing Network Connectivity with Ping and Traceroute

Maintenance > [Tools] Ping/Traceroute

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

The screenshot shows the 'Ping/Traceroute' web interface. Under 'Network Connectivity', the 'Ping' tab is selected and the 'IP/Hostname' field contains '10.0.238.136'. The 'Options' field is empty. The 'Output' section displays the following text:

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

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

Using ping

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

The screenshot shows the 'Ping/Traceroute' web interface. Under 'Network Connectivity', the 'Ping' tab is selected and the 'IP/Hostname' field contains '10.10.10.2'. The 'Options' field contains the command '-I 10.10.10.1'. The 'Start' and 'Clear' buttons are visible below the options field.

ping -I <local appliance IP> — sends the **ping** out a datapath interface

- **SYNOPSIS:**

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

The following **ping** options are supported:

Option	Explanation
-A	Adaptive ping. Interpacket interval adapts to round-trip time, so that effectively not more than one (or more, if preload is set) unanswered probes present in the network. Minimal interval is 200 msec if not super-user. On networks with low rtt this mode is essentially equivalent to flood mode.
-b	Allow pinging a broadcast address.
-B	Do not allow ping to change source address of probes. The address is bound to one selected when ping starts.
-c	<i>count</i> : Stop after sending count ECHO_REQUEST packets. With deadline option, ping waits for count ECHO_REPLY packets, until the time- out expires.
-d	Set the SO_DEBUG option on the socket being used. Essentially, this socket option is not used by Linux kernel.
-F	<i>flow label</i> : Allocate and set 20 bit flow label on echo request packets. (Only ping6). If value is zero, kernel allocates random flow label.
-i	<i>interval</i> : Wait interval seconds between sending each packet. The default is to wait for one second between each packet normally, or not to wait in flood mode. Only super-user may set interval to values less 0.2 seconds.
-I	<i>interface address</i> : Set source address to specified interface address. Argument may be numeric IP address or name of device. When pinging IPv6 link-local address this option is required.
-l	<i>preload</i> : If preload is specified, ping sends that many packets not waiting for reply. Only the super-user may select preload more than 3.
-L	Suppress loopback of multicast packets. This flag only applies if the ping destination is a multicast address.
-M	<i>MTU discovery hint</i> : Select Path MTU Discovery strategy. hint may be either do (prohibit fragmentation, even local one), want (do PMTU discovery, fragment locally when packet size is large), or dont (do not set DF flag).
-n	Numeric output only. No attempt will be made to lookup symbolic names for host addresses.
-p	<i>pattern</i> : You may specify up to 16 "pad" bytes to fill out the packet you send. This is useful for diagnosing data-dependent problems in a network. For example, -p ff will cause the sent packet to be filled with all ones.
-Q	<i>tos</i> : Set Quality of Service -related bits in ICMP datagrams. tos can be either decimal or hex number. Traditionally (RFC1349), these have been interpreted as: 0 for reserved (currently being redefined as congestion control), 1-4 for Type of Service and 5-7 for Precedence. Possible settings for Type of Service are: minimal cost: 0x02, reliability: 0x04, throughput: 0x08, low delay: 0x10. Multiple TOS bits should not be set simultaneously. Possible settings for special Precedence range from priority (0x20) to net control (0xe0). You must be root (CAP_NET_ADMIN capability) to use Critical or higher precedence value. You cannot set bit 0x01 (reserved) unless ECN has been enabled in the kernel. In RFC2474, these fields has been redefined as 8-bit Differentiated Services (DS), consisting of: bits 0-1 of separate data (ECN will be used, here), and bits 2-7 of Differentiated Services Codepoint (DSCP).
-q	Quiet output. Nothing is displayed except the summary lines at startup time and when finished.

Option	Explanation (Continued)
-R	Record route. Includes the RECORD_ROUTE option in the ECHO_REQUEST packet and displays the route buffer on returned packets. Note that the IP header is only large enough for nine such routes. Many hosts ignore or discard this option.
-r	Bypass the normal routing tables and send directly to a host on an attached interface. If the host is not on a directly attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it provided the option -I is also used.
-s	<i>packetsize</i> : Specifies the number of data bytes to be sent. The default is 56, which translates into 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.
-S	<i>sndbuf</i> : Set socket sndbuf. If not specified, it is selected to buffer not more than one packet.
-t ttl	Set the IP Time to Live.
-T	<i>timestamp option</i> : Set special IP timestamp options. timestamp option may be either tsonly (only timestamps), tsandaddr (timestamps and addresses) or tsprespec host1 [host2 [host3 [host4]]] (timestamp prespecified hops).
-U	Print full user-to-user latency (the old behavior). Normally ping prints network round trip time, which can be different f.e. due to DNS failures.
-v	Verbose output.
-V	Show version and exit.
-w	<i>deadline</i> : Specify a timeout, in seconds, before ping exits regardless of how many packets have been sent or received. In this case ping does not stop after count packet are sent, it waits either for deadline expire or until count probes are answered or for some error notification from network.

Using traceroute

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

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

The following **traceroute** options are supported:

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

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

Capturing Packets for Analysis

Maintenance > [Tools] Packet Capture

This page allows you to capture packets from the network. System automatically captures packets on all physical interfaces and some virtual debug interfaces. Packets are stored in a pcap file which can later be downloaded for analysis using tools like Wireshark.

Packet Capture ?

Maximum Number of Packets: Estimated size per packet capture file is 14.31MB

IP to capture from:

Port to capture from:

Packet Capture Files

Disk Usage

Usage Percentage: 43% 10,943 MB

Show Search

File Name	File Last Modified Time	File Size	
tcpdumps-dm-vx3000a-2014-05-29-13-59-49.tar.gz	Thu, 29 May 2014 20:59:54 GMT	41.2MB	✕
tcpdumps-DM-VX-A-2015-03-12-10-25-24.tar.gz	Thu, 12 Mar 2015 17:25:28 GMT	40.6MB	✕
tcpdumps-DM-VX-A-2014-11-04-16-42-11.tar.gz	Wed, 05 Nov 2014 00:42:14 GMT	39.1MB	✕
tcpdumps-DM-VX-A-2014-11-04-16-32-01.tar.gz	Wed, 05 Nov 2014 00:32:05 GMT	39.4MB	✕
tcpdumps-DM-VX-A-2014-11-04-16-31-43.tar.gz	Wed, 05 Nov 2014 00:31:48 GMT	39.8MB	✕
tcpdumps-DM-VX-A-2014-11-04-16-31-31.tar.gz	Wed, 05 Nov 2014 00:31:35 GMT	41.4MB	✕

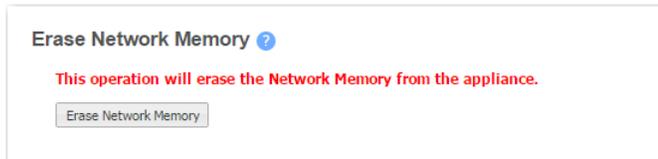
Showing 1 to 6 of 6 entries

- When you click **Start**, packet capture continues until the file reaches the **Maximum Number of Packets** or until you click **Stop**.
- To isolate specific traffic, you can filter by source or destination IP and/or by port.

Erasing Network Memory

Maintenance > [Tools] Erase Network Memory

This command erases the Network Memory cache.
No reboot required.



Rebooting or Shutting Down an Appliance

Maintenance > [Tools] Reboot/Shutdown

The appliance supports three types of reboot:



- **Reboot.** Reboots the appliance gracefully. This is your typical "vanilla" restart.
Use case: You're changing the deployment mode or other configuration parameters that require a reboot.
- **Erase Network Memory and Reboot.** Erases the Network Memory cache and reboots the appliance.
Use case: You need to restart the appliance with an empty Network Memory cache.
- **Shutdown.** Shuts down the appliance and turns the power off. To restart, go to the appliance and physically turn the power on with the Power switch.
Use case:
 - You're decommissioning the appliance.
 - You need to physically move the appliance to another location.
 - You need to recable the appliance for another type of deployment.

Behavior During Reboot

A physical appliance enters into one of the following states:

- *hardware bypass*, if deployed in-line (Bridge mode), or
- *an open-port state*, if deployed out-of-path (Router/Server mode).

Unless a **virtual appliance** is configured for a high availability deployment, all flows are discontinued during reboot.

How to Contact Support

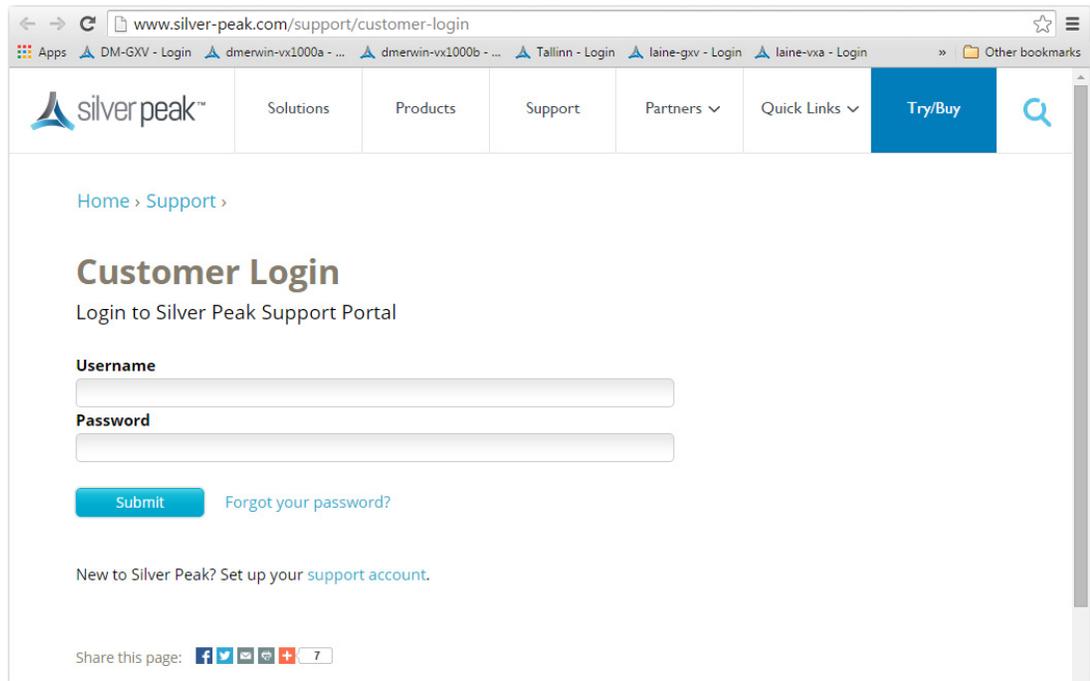
Support > [Technical Assistance] Technical Support

Support	
Technical Support	
Model	VX-3000 205002006000 Rev 44261
Serial Number	00-0C-29-47-5A-FE
Release	VXOA 7.3.0.0_55922
Website	www.silver-peak.com/Support
User Documentation	www.silver-peak.com/Support/user_docs.asp
Email	support@silver-peak.com
Phone	
North America (USA/CAN)	+1 877 210 7325
Global	+1 408 935 1850

Logging In to the Support Portal

Support > [Technical Assistance] Support Portal Log-in

If you have a **maintenance contract** with Silver Peak, then you have a username and password that provides you access to the Silver Peak Support Portal.



The screenshot shows a web browser window displaying the Silver Peak Customer Login page. The browser's address bar shows the URL www.silver-peak.com/support/customer-login. The page features a navigation menu with links for Solutions, Products, Support, Partners, Quick Links, and a prominent blue Try/Buy button. Below the navigation, the breadcrumb trail reads Home > Support >. The main heading is "Customer Login" with the subtext "Login to Silver Peak Support Portal". There are two input fields: "Username" and "Password". A blue "Submit" button is positioned below the password field, and a link for "Forgot your password?" is to its right. At the bottom of the form area, there is a link: "New to Silver Peak? Set up your [support account](#)." Below the form, there is a "Share this page:" section with social media icons for Facebook, Twitter, LinkedIn, and others.

Viewing Logs

Support > [Technical Assistance] Log Viewer

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

Log Viewer ?

Filters

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

Show 100 Search

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

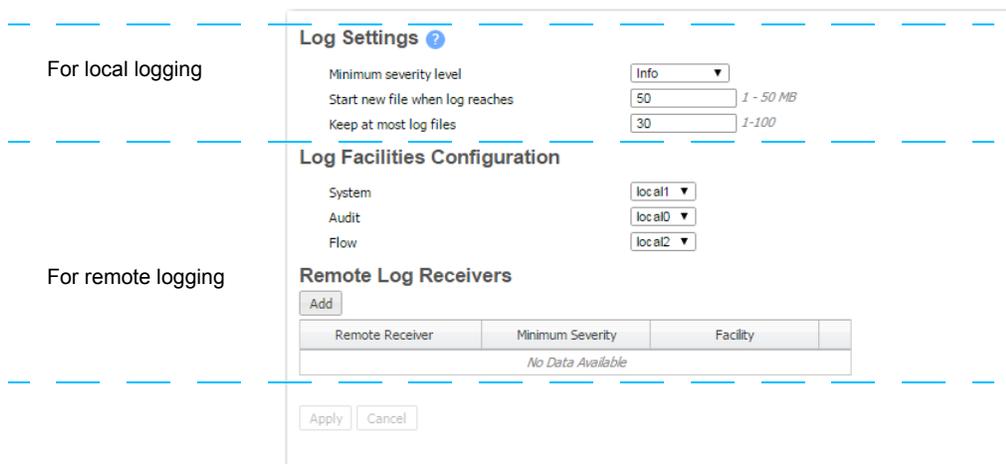
Showing 1 to 100 of 100 entries First Previous 1 Next Last

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

Configuring Logging Parameters

Support > [Technical Assistance] Log Settings

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



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

Minimum Severity Levels

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

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

- The bolded part of the name is what displays in Silver Peak's logs.

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

Configuring Remote Logging

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

Managing Debug Files

Support > [Technical Assistance] Debug Files

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

Silver Peak uses these files for evaluation and debugging.

Debug Files ?

Disk Usage
Usage Percent: 43% 10,984 MB

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

Show: 25 | Search: _____

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

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



APPENDIX A

TCP/IP Ports Used by the Orchestrator and Silver Peak Appliances

Following are lists of ports that are used by the appliances and by the Orchestrator.

In This Appendix

- **Ports used by the Orchestrator** See page 174.
- **Ports used by Silver Peak Appliances** See page 175.

Ports used by the Orchestrator

Application	Client	Server	Comments
SSH	No	Yes	Listens on TCP Port 22
HTTPS	Yes If appliance is added manually, then Orchestrator must be able to talk to the appliance on Port 443.	Yes	Listens on TCP Port 443
HTTP	Yes	No	
FTP	Yes	No	Used for backup
SCP	Yes	No	Used for backup
TACACS+	Yes	No	Used for authentication
RADIUS	Yes	No	Used for authentication
DNS	Yes	No	
NTP	Yes	No	

Ports used by Silver Peak Appliances

Appliance Management Plane

Application	Client	Server	Comments
SSH	No	Yes	Listens on TCP Port 22
HTTPS	Yes Must be able to talk to Orchestrator on Port 443.	Yes	Listens on TCP Port 443
HTTP	Yes	Yes	Listens on TCP Port 80 — can be disabled
FTP	Yes	No	Used for backup
SCP	Yes	No	Used for backup
TACACS+	Yes	No	Used for authentication
RADIUS	Yes	No	Used for authentication
DNS	Yes	No	
NTP	Yes	No	
SNMP	Yes (for Traps)	Yes	Listens on UDP Port 161
Syslog	Yes	No	
Netflow	Yes	No	

Appliance Data Plane

Application	Ports and Protocols	Comments
GRE	IP Protocol 47	Firewalls must allow this protocol for GRE tunnels.
IPsec	Protocol ESP 50 UDP Port 500 UDP Port 4500	Firewalls must allow these for IPsec tunnels.
UDP	Port 4163	Firewalls must allow this port for UDP tunnels.



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

1.877.210.7325
+1.408.935.1850

www.silver-peak.com