



Silver Peak



Silver Peak NX Series Appliances

Operator's Guide

Release 3.3

October 2010

PN 200030-001 Rev I

Silver Peak NX Series Appliance Operator's Guide

Document PN 200030-001 Rev I

Date: October 2010

Copyright © 2005–2010 Silver Peak Systems, Inc. All rights reserved. Information in this document is subject to change at any time. Use of this documentation is restricted as specified in the *End User License Agreement*. No part of this documentation can be reproduced, except as noted in the *End User License Agreement*, in whole or in part, without the written consent of Silver Peak Systems, Inc.

Trademark Notification

Silver Peak Systems™, the Silver Peak logo, Network Memory™, and Silver Peak NX-Series™ are trademarks of Silver Peak Systems, Inc. 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.
4500 Great America Parkway, Suite 100
Santa Clara, CA 95054

1.877.210.7325 (toll-free in USA)
+ 1.408.935.1850

www.silver-peak.com/support

Contents

Preface	i
Who Should Read This Manual?	i
Manual Organization	i
Related Publications	iii
Technical Support	iii
 Chapter 1 Overview	 1
Overview of the Silver Peak NX Series Appliances	2
Benefits	2
Components	3
Features	8
Typical Network Deployments	11
In-Line Deployments	11
Out-of-Path Deployments (Router Mode)	15
 Chapter 2 Installing the Appliance	 21
Before You Begin	22
Summary of Installation Tasks	22
Site Preparation	22
Rack Mounting the Appliance	25
NX-1700	25
NX-2500	27
NX-2600 and/or NX-2610	28
NX-3500	33
NX-3600	34
NX-5500, NX-5504, NX-7500, NX-7504, and NX-8504	41
NX-5600, NX-7600, NX-8600, and NX-9610	45
NX-2700, NX-3700, NX-5700, NX-7700, NX-8700 and NX-9700	51
Connecting the Power and Verifying LEDs	57
NX-1700	57
NX-2500	57
NX-2600 and NX-2610	58
NX-3500	59
NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504	60
NX-5600, NX-7600, NX-8600, or NX-9610	62
NX-2700, NX-3700, NX-5700, NX-7700, NX-8700 and NX-9700	63
Installing the Appliance into the Network	64
Cabling for Configuration Management	68
Running the Initial Configuration Wizard	69
 Chapter 3 The Appliance Manager	 73
Accessing the Appliance Manager	74
Guided Tour of the Appliance Manager	76
The Appliance Manager Home Page	76
Banners	80
Menu Structure	80
Managing Tabular Data	88
Netmask Notations	88
Date and Time Conventions	90
Secure Access Methods	90
Guidelines for Creating Passwords	90

Object Names	90
Saving Your Configuration	91
Definition Help	91
Chapter 4 Configuring Host Settings	93
Overview	94
Configuring Appliance Identity and Max System Bandwidth	94
Selecting a System Deployment	96
Sorting Through the Deployment Options	98
Configuring Gigabit Etherchannel Bonding	99
Setting the Date and Time	100
Configuring Network Parameters	103
Modifying the Physical Interface Parameters	104
Configuring IP Routes	107
Routing Management Traffic	108
Routing LAN-side Traffic to the Next Hop	111
Adding Domain Name Servers	113
Configuring Flow Exports for NetFlow	114
Chapter 5 Creating Tunnels	115
Overview	116
How Policies Affect Tunnel Traffic	116
Tunnel Characteristics	116
Creating a Traffic-Carrying Tunnel	117
Editing a Tunnel	123
Deleting a Tunnel	125
Tunnel Compatibility Mode	126
Chapter 6 Theory of Operations	127
Processing Traffic Flows	128
What Maps and Policies Do	128
Default Behaviors	129
Understanding MATCH Criteria	130
Configuring MATCH Criteria in a Map or Policy	131
Using ACLs to Summarize MATCH Criteria	132
Specifying Applications and Protocols in MATCH Criteria	133
Using ACLs (Access Control Lists)	135
Creating an Access Control List (ACL)	136
Modifying an ACL Rule	139
Removing an ACL Rule	140
Removing an ACL	141
How Policies and ACLs Filter Traffic	145
Managing Applications and Application Groups	147
Built-in Applications	147
Defining Custom Applications	152
Creating and Using Application Groups	156
Chapter 7 Route Policy	161
Introduction	162
How Auto Optimization Works	163
Handshaking for Auto Optimization in Bridge Mode	163
Handshaking for Auto Optimization in Router Mode	164
Where the Route Policy Can Direct Flows	165
Flow directed to a tunnel	165
Flow designated as auto-optimized	166

Flow designated as shaped pass-through traffic	167
Flow designated as unshaped pass-through traffic	168
Flow dropped	169
Continue option used in Tunnel Down Action	170
Route Policy Page Organization	171
Managing the Route Policy	172
Adding an Entry to a Map	172
Editing an Entry	174
Deleting an Entry	175
Adding a New Route Map	175
Deleting a Map	178
Activating a New Policy	178
Chapter 8 Bandwidth Management & QoS Policy	179
Overview	180
How the QoS Policy Affects Flows	180
Flow directed to a tunnel	181
Flow designated as pass-through shaped traffic	182
Flow designated as unshaped pass-through traffic	183
Best Practices for Bandwidth Management	184
Summary of Bandwidth Assessment and Management Tasks	184
Guidelines for Configuring Minimum and Maximum Bandwidth Values	185
Which Appliance Manager Pages to Use	187
Configuring Maximum System Bandwidth	188
2-Port Configurations	188
4-Port Configurations	188
How Tunnel Auto BW Works	189
Configuring Pass-Through Traffic Bandwidths	190
Configuring Traffic Classes	191
Traffic Class Components	192
Handling and Marking Packets	194
Applying DSCP Markings to Optimized Traffic	194
Applying DSCP Markings to Shaped and Unshaped Pass-through Traffic	197
Definitions of DSCP Markings	199
QoS Policy Page Organization	201
Managing the QoS Policy	202
Adding an Entry to a Map	202
Editing an Entry	204
Deleting an Entry	205
Adding a New QoS Map	206
Deleting a Map	208
Activating a New Policy	209
Chapter 9 Optimization Policy	211
Introduction	212
Network Memory	212
Payload Compression	213
TCP Acceleration	213
CIFS Acceleration	213
Making the Best Use of Optimizations	214
When the Appliance Can Apply the Optimization Policy	215
Optimization Policy Page Organization	216
Managing the Optimization Policy	217
Adding an Entry to a Map	217
Editing an Entry	219
Deleting an Entry	220
Adding a New Optimization Map	220

Deleting a Map	223
Activating a New Policy	224
Chapter 10 Using Flow Redirection to Address TCP Asymmetry	225
Introduction	226
Asymmetrical Networks and Flows	226
Removing Asymmetry with Flow Redirection	226
Redirection for WAN-initiated Traffic	227
Redirection for LAN-initiated Traffic	228
Configuring Flow Redirection	229
Example #1: Simple Cluster with Two Physically Connected Peers	230
Flow Reporting	234
Chapter 11 Configuring and Managing VLANs	235
Introduction	236
Configuring a VLAN IP Interface	237
Setting VLAN Tags in Outgoing WAN-side Packets	238
2-Port Bridge	239
Standard 4-Port Bridge	240
Flat 4-Port Bridge	241
Bonded 4-Port Bridge	242
Chapter 12 Reporting Historical Traffic	243
Overview	244
About Viewing Statistics	244
Understanding Traffic Direction	245
Viewing Pie Charts	246
Selecting Time Periods	247
Exporting Table Data	247
Viewing Application Historical Statistics	248
A Sampling of Results	248
What Data Displays	249
Viewing Reduction Statistics	251
A Sampling of Results	251
Viewing Bandwidth Statistics	253
Viewing Flow Counts	255
Viewing Latency	257
Viewing Network Integrity	258
Viewing a Summary of All Historical Reports	261
Chapter 13 Monitoring Realtime Traffic	263
Overview	264
About Viewing Statistics	266
Understanding Traffic Direction	266
Viewing Counters Since Last Reboot	266
Clearing Counters Non-Destructively	267
Viewing Pie Charts	268
Exporting Table Data	269
Viewing Application Realtime Statistics	270
A Sampling of Results	270
What Data Displays	271
Viewing Current Flows	273
Selecting Filters	274
Customizing Which Columns Display	275
Current Flow Details	281

Resetting Flows to Improve Performance	285
Viewing Tunnel QoS Statistics	286
Viewing Tunnel Realtime Statistics	288
LAN/WAN Statistics	290
Flows / Latency / Packet Correction Statistics	291
Viewing Reduction Statistics	293
A Sampling of Results	293
Viewing Bandwidth Statistics	295
Viewing Flow Counts	297
Viewing Latency Statistics	299
Viewing Network Integrity Statistics	300
Viewing Flow Redirection Statistics	302
Viewing NetFlow Statistics	304
Viewing Interface Statistics	305
Viewing Bridge Mode Statistics	307
Two-Port Example	307
Four-Port Example	307
Viewing IP Routes	308
Chapter 14 Administration Tasks	309
Configuring Log Settings	310
Configuring Local Logging	312
Configuring Remote Logging	313
Understanding the Events Log	315
Viewing a Log of All Alarms	316
Viewing the Audit Log	317
Managing Debug Files	318
Types of Debug Files	318
Saving Files to a Remote Server	320
Deleting Log Files	324
Pre-Positioning Data for Enhanced Acceleration Benefits	325
Configuring SNMP	327
Loading SNMP MIBs	327
Configuring SNMP Settings	328
Managing User Accounts	332
Guidelines for Creating Passwords	332
Accessing User Accounts	333
Creating a User Account	334
Modifying a User Account	335
Deleting a User Account	336
Configuring Authentication, RADIUS, and TACACS+	337
Authentication and Authorization	337
Session Idle Time-out	338
Configuring for RADIUS	339
Configuring for TACACS+	346
Configuring Banners	354
Configuring Settings for Web Protocols and Web Users	355
Initial Configuration Wizard	356
Support	358
Chapter 15 System Maintenance	359
Viewing System Information	360
Upgrading the Appliance Manager Software	361
Overview	361
Installing a New Software Image into a Partition	364

Installing the Software Image from the Local Disk	365
Installing the Software Image from a URL	366
Installing the Software Image from an SCP Server	367
Installing the Software Image from an FTP Server	369
Switching to the Other Software Load	371
Managing the Appliance Configuration File	372
Viewing the Appliance Configuration File	372
Saving the Appliance Configuration File	375
Downloading the Appliance Configuration File	380
Testing Network Connectivity	386
Using ping	389
Using traceroute	391
Using tcpdump	393
Erasing Network Memory	399
Restarting the Appliance	400
Chapter 16 Monitoring Alarms	401
Understanding Alarms	402
Categories of Alarms	402
Types of Alarms	403
Viewing Current Alarms	406
Handling Current Alarms	408
Acknowledging Alarms	408
Clearing Alarms	408
Chapter 17 Hardware Maintenance	409
Replacing a Hard Disk Drive	410
Replacing a Power Supply	421
Replacing a Power Supply in the NX-9700, NX-9610, NX-8700, NX-8600, NX-7700, NX-7600, NX-5700, NX-5600, NX-3700, NX-3600, or NX-2700	422
Replacing a Power Supply in the NX-8504, NX-7500, NX-7504, NX-5500, or NX-5504	423
Replacing a Power Supply in the NX-3500	424
Appendix A Specifications, Compliance, and Regulatory Statements	425
Model Specifications	426
Model-specific Specifications	426
Fiber Specifications	431
NX-Series Specifications	431
VX Series Specifications and Requirements	432
Warning Statements	433
NX-9600	433
Compliance Statements	434
FCC Compliance Statement	434
ICES-003 statement	434
Requirements for Rack-Mount Equipment	434
Requirements for Knurled Thumb Screws	435
Cable Pinouts	436
Appendix B Glossary	437
Index	445



Preface

The Silver Peak NX Series appliances enable branch office infrastructure centralization by delivering applications across a WAN with LAN-like performance.

Who Should Read This Manual?

Anyone who wishes to install the NX Series appliances should read this manual. Users should have some background in Windows[®] terminology, Web browser operation, and a knowledge of where to find the TCP/IP and subnet mask information for their system.

Manual Organization

This section outlines the chapters and summarizes their content.

Chapter 1, “Overview,” provides an overview of the Silver Peak NX Series appliances and the Appliance Manager graphical user interface. It also explains the basic concepts and core functionality, along with providing a summary of typical network deployments.

Chapter 2, “Installing the Appliance,” describes the procedures for installing your Silver Peak appliance to prepare for in-line deployment (Bridge mode) and out-of-path deployment (Router mode). It describes the preparations you need to make, how to install the appliance in a rack, how to use a web browser to run the initial configuration wizard, how to add the appliance into the network, and how to verify connectivity.

Chapter 3, “The Appliance Manager,” explains how to access the Appliance Manager through your browser. It also familiarizes you with the task-related and graphical conventions used throughout the interface screens.

Chapter 4, “Configuring Host Settings,” describes how to configure or modify the existing appliance system parameters, including the WAN bandwidth at the far side of the router. Additionally, it describes how to set the date and time, add DNS servers, work with the routing table, modify network interface parameters, configure gigabit ethernetchannel bonding, and set up export to NetFlow collectors.

Chapter 5, “Creating Tunnels,” describes the relationship among tunnels, Access Control Lists (ACLs), and policies, as they relate to directing and processing traffic for acceleration. It describes how to create custom applications that you can use in ACLs. It also prescribes best practices for creating an Up and Active tunnel, and details the procedures.

Chapter 6, “Theory of Operations,” describes how the Silver Peak appliance optimizes traffic by allowing you to define flows with MATCH criteria and direct flows with policy maps. It also describes techniques for streamlining your network management by using Access Control Lists (ACLs), user-defined applications, and application groups.

Chapter 7, “Route Policy,” focuses on the SET actions that are specific to the Route policy. Where applicable, they also provide context relative to the Optimization and QoS policies. It also explains how Auto Optimization works, enabling you to get up and running after only configuring a tunnel.

[Chapter 8, “Bandwidth Management & QoS Policy,”](#) describes the QoS Policy’s SET actions. It also explains how to configure traffic classes for optimized and pass-through traffic, along with providing best practices guidelines for effectively managing bandwidth.

[Chapter 9, “Optimization Policy,”](#) describes how the appliance optimizes tunnelized traffic — improving the performance of applications across the WAN.

[Chapter 10, “Using Flow Redirection to Address TCP Asymmetry,”](#) describes how flow redirection enables Silver Peak appliances to optimize asymmetrically routed flows by redirecting packets between appliances.

[Chapter 11, “Configuring and Managing VLANs,”](#) describes how to configure and manage VLANs when the appliance is in Bridge mode.

[Chapter 12, “Reporting Historical Traffic,”](#) describes how to create reports and view statistics collected over a specific interval for applications, data reduction, bandwidth optimization, flow counts, latency, and network integrity (loss and out-of-order packets).

[Chapter 13, “Monitoring Realtime Traffic,”](#) describes how to view realtime statistics for applications, current flows, QoS, tunnels, data reduction, bandwidth optimization, flow counts, latency, flow redirection, NetFlow, interfaces, and bridge mode. Generally, this includes the last hour’s worth of collected data.

[Chapter 14, “Administration Tasks,”](#) describes administrative tasks such as configuring log settings, viewing event and alarm logs, managing debug files, pre-positioning file server data into Network Memory, configuring SNMP, managing user accounts (as well as their authorization and authentication), configuring settings for web protocols and web users, re-accessing the initial configuration wizard, and contacting Silver Peak Support.

[Chapter 15, “System Maintenance,”](#) describes tasks related to maintaining the hardware, software, and database. This includes tasks such as managing the software images and the configuration files, testing network connectivity, managing the hard disks, erasing Network Memory, and restarting the appliance.

[Chapter 16, “Monitoring Alarms,”](#) describes alarms categories and definitions. It also describes how to view and handle alarm notifications.

[Chapter 17, “Hardware Maintenance,”](#) describes how to rebuild a failed RAID array, replace a power supply, and replace a hard disk drive.

Related Publications

Refer to the following related publications for more information:

	Document	Part Number
Manuals	Silver Peak NX Series Appliances Network Deployment Guide	200059-001
	Silver Peak Command Line Interface Reference Guide	200063-001
	Silver Peak Global Management System User's Guide	200095-001
Mount Instructions	Rack Mount Instructions: 3-RU with rails	200258-001
	Rack Mount Instructions: 1-RU with rails	200259-001
	Rack Mount Instructions: 1-RU without rails	200260-001
	Rack Mount Instructions: 3-RU with rails – NX-9610-8600-7600-5600	200282-001
	Rack Mount Instructions: 2-RU with rails – NX-3600	200371-001
	Rack Mount Instructions: ear mount - NX-1700	200450-001
	Desk / Wall Mount Instructions: NX-1700	200461-001
	Rack Mount Instructions: 2_RU with rails – NX-9700/8700/7700/5700/3700/2700	200486-001
Quick Start Guides	GX-1000 Appliance Quick Start Guide	200080-001
	NX Series Appliances Quick Start Guide	200257-001
	Quick Start Guide – VX Virtual Appliance with VMware ESX/ESXi	200469-001
	Quick Start Guide – GX-V Virtual GMS Server – VMware ESX/ESXi	200471-001
System Requirements	VX Host System Requirements	200468-001
	GX-V Host System Requirements	200476-001
Release Notes	Check www.silver-peak.com/support for the latest version.	

Technical Support

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

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



CHAPTER 1

Overview

This chapter describes the structure, components, and features of the Silver Peak NX Series appliances. It acquaints you with basic concepts and core functionality. It also provides a basic overview of in-line (Bridge mode) and out-of-path (Router mode) deployments.

In This Chapter

- **Overview of the Silver Peak NX Series Appliances** See page 2.
- **Typical Network Deployments** See page 11.

Overview of the Silver Peak NX Series Appliances

Silver Peak's NX Series™ appliances reduce IT costs and enhance enterprise-wide data security and regulatory compliance by enabling the centralization of branch office servers and storage. In addition, they improve the performance and reliability of backup, replication, and recovery across a Wide Area Network (WAN).

The NX Series appliances leverage Local Instance Networking (LIN) to achieve secure, scalable application delivery and significantly improved performance over existing application acceleration approaches. Local Instance Networking effectively localizes information in each office while retaining control where it belongs — centrally. Deployment of the NX Series requires absolutely no client, server, or application reconfiguration.

This section describes features and benefits, along with providing an overview of hardware and software components.

Benefits

♦ Security

The NX Series appliances are the only WAN acceleration appliances with 128-bit encrypted disk drives to protect data stored on the device. IPSec encryption protects data sent between appliances. Hardware acceleration ensures that data security is achieved with little or no impact on application performance.

♦ Resilience

Redundant hardware protects against disk drive and power failures. Additionally, fail-to-wire network interfaces mechanically isolate the appliance in the event of hardware, software, or power failures.

♦ Scalability

Network Memory™ provides a common data store across all locations, preventing storage repetition and ensuring efficient usage of appliance CPU and disk space.

♦ High Availability Deployment

To maximize uptime, you can deploy NX appliances redundantly in 1+1 or N+1 configurations, with failover and load balancing.

♦ Easy to Manage

The intuitive, web-based Appliance Manager Graphical User Interface (GUI) simplifies network monitoring, policy provisioning, and device management. Powerful wizards simplify configuration. The GUI is available via HTTP and HTTPS.

A full-featured CLI is available over the DB-9 console port (RS-232 serial port) or via SSH. For the port pinout, see *“Cable Pinouts” on page 436 in Appendix A*. Also see *“Related Publications” on page iii*.

Larger deployments can easily be managed using Silver Peak's GX-1000 appliance, running the Global Management System (GMS).

♦ Easy Deployment

You can deploy Silver Peak appliances in-line (in Bridge mode) between an Ethernet LAN switch and a WAN router, or out-of-path (in Router mode) using Policy-based-routing redirection, Web Cache Coordination Protocol (WCCP), or Virtual Router Redundancy Protocol (VRRP). Multiple appliances can be clustered for increased scalability. Typical deployment takes less than 30 minutes per appliance.

Components

Silver Peak's NX Series hardware appliances and VX Series virtual appliances enable organizations to centralize branch office server and storage infrastructure on a broad basis, reducing IT costs and enhancing enterprise-wide data security and compliance.

Silver Peak appliances are deployed in each office of a distributed enterprise network and typically sit “behind” the Wide Area Network (WAN) router. The appliances support a variety of different installation modes and robust fallback mechanisms, making them a perfect fit for all enterprise situations.

Hardware Appliances

The NX Series includes the following models, designed to fit seamlessly into any enterprise network and to accommodate a wide range of enterprise office environments:



NX-1700

- A 1-RU appliance that supports 4 Mbps of WAN bandwidth and 500 GB of secure local data storage.
- The **NX-1700** is ideal for branch or remote offices.



NX-2500

- A 1-RU appliance that supports 2 Mbps of WAN bandwidth and 250 GB of secure local data storage.
- The **NX-2500** is ideal for branch or remote offices.



NX-2600

- A 1-RU appliance that supports 4 Mbps of WAN bandwidth and includes 250 GB of secure local data storage.
- The **NX-2600** is ideal for branch or remote offices.



NX-2610

- A 1-RU appliance that supports 8 Mbps of WAN bandwidth and includes 500 GB of secure local data storage.
- The **NX-2610** is ideal for branch or remote offices.

**NX-2700**

- A 2-RU appliance that supports 10 Mbps of WAN bandwidth and 1 TB of secure local data storage.
- The **NX-2700** is ideal for branch or remote offices.

**NX-3500**

- A 2-RU appliance that supports 10 Mbps of WAN bandwidth and 500 GB of secure local data storage.
- The **NX-3500** is ideal for mid-size offices or corporate data centers.

**NX-3600**

- A 2-RU appliance that supports 20 Mbps of WAN bandwidth and 1 TB of secure local data storage.
- 4 LAN/WAN data ports
- The **NX-3600** is for mid-size offices or corporate data centers.

**NX-3700**

- A 2-RU appliance that supports 20 Mbps of WAN bandwidth and 1 TB of secure local data storage.
- The **NX-3700** is ideal for mid-size offices or corporate data centers.

**NX-5500 / NX-5504**

- A 3-RU appliance that supports 50 Mbps of WAN bandwidth and 2 TB of secure local data storage.
- The **NX-5500** brings application acceleration to medium and large offices.

**NX-5600**

- A 3-RU appliance that supports 50 Mbps of WAN bandwidth and 2 TB of secure local data storage.
- The **NX-5600** brings application acceleration to medium and large offices.

**NX-5700**

- A 2-RU appliance that supports 50 Mbps of WAN bandwidth and 4 TB of secure local data storage.
- The **NX-5700** brings application acceleration to medium and large offices.

**NX-7500 / NX-7504**

- A 3-RU appliance that supports 155 Mbps of WAN bandwidth and 2 TB of secure local data storage.
- The **NX-7500** is intended for deployment in larger data centers.

**NX-7600**

- A 3-RU appliance that supports 155 Mbps of WAN bandwidth and 3 TB of secure local data storage.
- The **NX-7600** is intended for deployment in larger data centers.

**NX-7700**

- A 2-RU appliance that supports 155 Mbps of WAN bandwidth and 5 TB of secure local data storage.
- The **NX-7700** is ideal for deployment in larger data centers.

**NX-8504**

- A 3-RU appliance that supports 500 Mbps of WAN bandwidth and 7 TB of secure local data storage.
- The **NX-8500** is intended for deployment in larger data facilities, such as regional hubs, multinational data centers, and disaster recovery locations.

**NX-8600**

- A 3-RU appliance that supports 500 Mbps of WAN bandwidth and 8 TB of secure local data storage.
- The **NX-8600** is intended for deployment in larger data facilities, such as regional hubs, multinational data centers, and disaster recovery locations.

**NX-8700**

- A 2-RU appliance that supports 622 Mbps of WAN bandwidth and 5 TB of secure local data storage, enhanced by 4 x 64GB SSDs.
- The **NX-8700** is intended for deployment in larger data facilities, such as regional hubs, multinational data centers, and disaster recovery locations.

**NX-9610**


- A 3-RU appliance that supports 1 Gbps of WAN bandwidth and 8 TB of secure local data storage.
- The **NX-9610** is intended for deployment in larger data facilities, such as regional hubs, multinational data centers, and disaster recovery locations.

**NX-9700**

- A 2-RU appliance that supports 1 Gbps of WAN bandwidth and 5 TB of secure local data storage, enhanced by 4 x 64GB SSDs.
- The **NX-9700** is intended for deployment in larger data facilities, such as regional hubs, multinational data centers, and disaster recovery locations.

Virtual Appliances

Silver Peak's VX, or Virtual, appliances are software versions of the company's winning NX appliances. They support all of Silver Peak's realtime Network Acceleration, Network Integrity, and Network Memory™ features to overcome common WAN bandwidth, latency, and quality challenges.

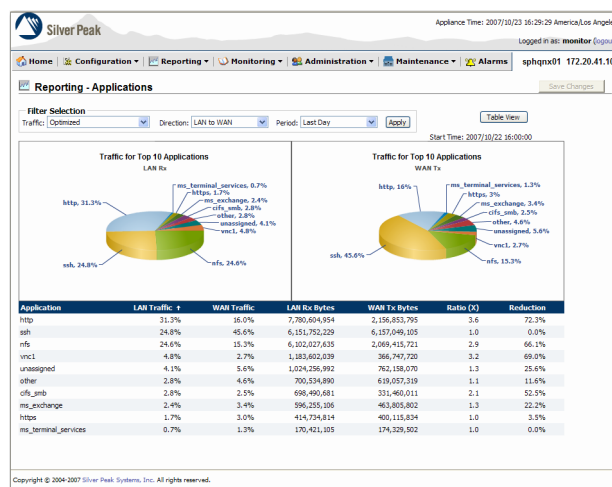


By running on industry-standard appliances, VX appliances leverage all the benefits of server virtualization, which include ease of deployment, reduced hardware costs, mobility, and high availability.

Software Component

The Appliance Manager features intuitive and powerful graphing tools to monitor your network's performance, application load, and to generate key ROI (Return On Investment) metrics.

The web-based Appliance Manager GUI is available via HTTP/HTTPS and features powerful wizards to simplify common appliance configuration tasks.



Features

The Silver Peak solution is fully transparent to clients, servers, networking equipment, and applications. Absolutely no client, server, or application reconfiguration is necessary. Once you deploy the appliances in a network, enterprise-wide benefits are realized immediately, with continued gains over time as content repetition increases.

This section describes features in terms of the following categories:

- **Network Memory™** See page 8.
- **Network Integrity** See page 8.
- **Network Acceleration** See page 8.
- **Management** See page 9.
- **In-Line and Out-of-Path Deployments** See page 10.

Network Memory™

All Silver Peak NX Series appliances are equipped with Network Memory™ technology — the cornerstone of the Silver Peak solution. With Network Memory, each NX Series appliance uses advanced fingerprinting algorithms to examine all incoming and outgoing WAN traffic. Network Memory™ localizes information and transmits only modifications between locations while retaining the control where it belongs — centrally.

Silver Peak appliances support state-of-the-art IP header compression, cross-flow payload compression, packet acceleration, and packet coalescing.

Network Integrity

Silver Peak ensures network integrity by using QoS management, Forward Error Correction, and Packet Order Correction.

Quality of Service Management

QoS management consists of packet classification into application flows, application-to-traffic-class mapping, and queuing and service disciplines. You can configure multiple traffic maps, each defining a mapping of applications to traffic classes. Only one traffic map can be active at any given time. Additionally, the user can also configure the priority of each traffic class for a given tunnel, and for pass-through traffic.

Forward Error Correction & Packet Order Correction

When Adaptive Forward Error Correction (FEC) is enabled, the appliance introduces a parity packet, which helps detect and correct single-packet loss within a stream of packets, reducing the need for retransmissions. Silver Peak dynamically adjusts how often this parity packet is introduced in response to changing link conditions. This maximizes error correction while minimizing overhead.

To avoid retransmissions that occur when packets arrive out of order, Silver Peak NX appliances use Packet Order Correction (POC) to resequence packets on the far end of a WAN link, as needed.

Network Acceleration

Silver Peak mitigates the impacts of latency across the WAN by using various TCP acceleration techniques, like adjustable window sizing and selective acknowledgements, as well as CIFS acceleration techniques, such as read-aheads and write-behinds. These tools help to overcome inherent chattiness that can otherwise hamper application performance across a WAN.

Management

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

System Access

Silver Peak supports four methods:

- **Appliance Manager Graphical User Interface (GUI):** The Silver Peak Appliance can be managed through the Appliance Manager. The Appliance Manager is implemented as a Java applet that can be downloaded directly from the appliance using a Web browser.
- **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.
- **Global Management System (GMS):** 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 NX Series appliances, GMS 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.

Real-Time/Summary Statistics

Interface- or tunnel-related real-time statistics and summaries represent the current running total of a given counter since the last appliance reboot or counter reset. Some real-time statistics, such as round-trip-time (RTT) measurements and number of active traffic flows, represent the current snapshot values over the last hour.

Historical/Interval Statistics

In addition to real-time statistics, the Silver Peak appliance maintains historical/interval statistics for report generation. These statistics represent delta values taken at predetermined intervals.

All historical reports support two levels of granularity:

- **Hourly interval:** values collected over 24 hours for periods beginning on the **Last Day**, **2 Days Ago**, **3 Days Ago**, and **4 Days Ago**
- **Daily interval:** cumulative daily values for the last 7 days (**Last Week**) or last 30 days (**Last Month**)

User Access Management

User Access Management controls and monitors access to the Silver Peak Appliance. User Access Management consists of:

- **User Authentication:** Supports local password protection and centralized authentication using RADIUS and/or TACACS+.
- **Access Privilege Control:** Two levels of privileges are supported – administration and monitoring. *Monitoring* privileges provide the user with read access to the configuration database, statistics, etc. *Administration* privileges include add, change, and delete, as well as monitoring privileges.
- **Access Audit:** This feature provides a mechanism to track access to the system. It also tracks unauthorized attempts to access the system.

Software Image Management

Image management is responsible for the loading and activation of system-bootable images. You can store downloadable images in one of two specified partitions, as well as specify which partition to boot from the next time the appliance restarts.

Configuration Database Backup and Restore

This feature provides the capabilities to upload and download configuration data. Whereas you can store multiple configuration files in the appliance, only one of the configuration files is active at any given time.

Logging/Debugging

All alarms and events are logged to the local disk. Optionally, you can configure Access Manager to send events of a that meet a specified minimum severity level to a remote syslog server.

Fault Management

Fault management is responsible for detection, isolation, and correcting faults in the Silver Peak Appliance. Fault Management provides the following:

- Generating an event to raise or clear an alarm condition
- The ability to acknowledge/un-acknowledge alarms
- The ability for the user to clear an alarm (applies to clearable alarm types only)
- An Active alarm table that holds existing outstanding fault conditions
- A log file that holds all alarms and events generated by the appliance

In-Line and Out-of-Path Deployments

Silver Peak appliances can be installed in the data path (in-line; bridge mode) between an L2/L3 switch and the edge WAN router, with fail-to-wire in case of failure.

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.

Alternatively, Silver Peak appliances can be installed out-of-path (router mode) to the WAN router. In an out-of-path deployment, policy-based routing (PBR), VRRP, or WCCP redirect the traffic to the Silver Peak appliance for processing.

The next sections outlines typical network deployments.

Typical Network Deployments

This section provides an overview of the supported in-line and out-of-path deployments, complete with diagram and summary.



For detailed configuration information, see [Silver Peak NX Series Appliances Network Deployment Guide](#).

For a discussion of Flow Redirection in both Bridge and Router modes, see [“Using Flow Redirection to Address TCP Asymmetry” on page 225](#).

In-Line Deployments

Silver Peak supports these typical in-line deployments:

- 1 **Bridge Mode - Two Ports** See page 11.
- 2 **Bridge Mode - Four Ports** See page 12.

Bridge Mode - Two Ports

This deployment does not require any configuration modifications to the L2 switch or the WAN router. Typically, this is a common deployment in a branch office.

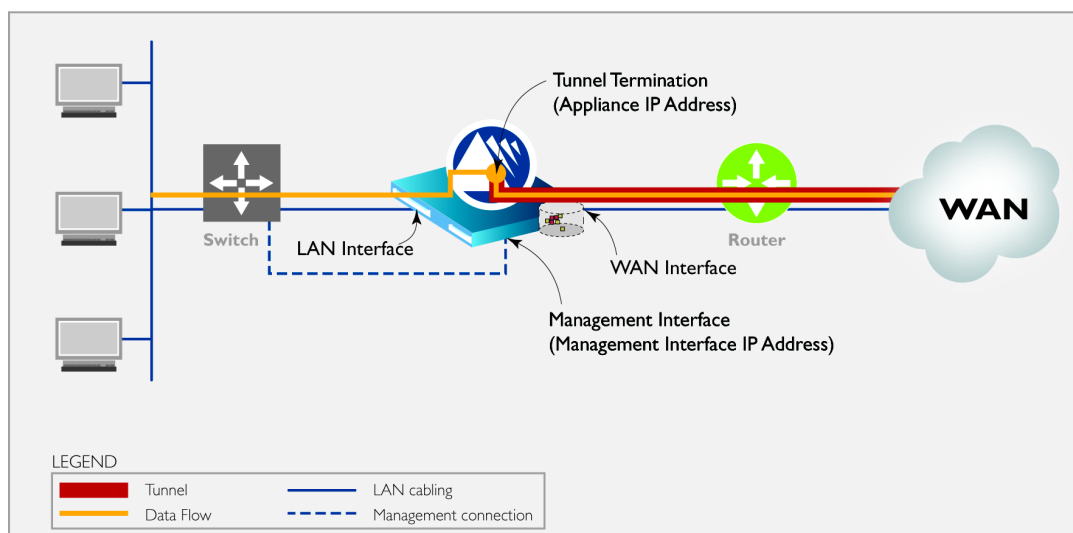


Figure 1-1 In-Line Deployment: Bridge Mode [Bridging with Fail-to-Wire]

Summary

Appliance Placement	<p>Appliance placed in-line between Ethernet LAN switch and WAN router</p> <ul style="list-style-type: none"> • Appliance LAN interface connects to Ethernet LAN switch • Appliance WAN interface connects to WAN router
Fail-Safe Behavior	<p>Fails-to-Wire: The appliance behaves as a crossover cable between the Ethernet LAN switch and the WAN router in any failure scenario (hardware, software, power).</p> <p>IMPORTANT: Ensure that the Ethernet LAN's switch and the WAN router have compatible Ethernet interface physical configuration settings (speed and duplex settings). This is to ensure that traffic flows correctly if the Silver-Peak appliance "Fails-to-wire".</p>
IP Addresses	<p>This deployment model requires two IP addresses (on the same or separate subnets)</p> <ul style="list-style-type: none"> • Silver Peak Appliance IP address (to originate and terminate tunnel) • Silver Peak Management IP Address (for appliance configuration and management)

Bridge Mode - Four Ports

This deployment does not require any configuration modifications to the L2 switch or the WAN router. Typically, this is a common deployment in a branch office.

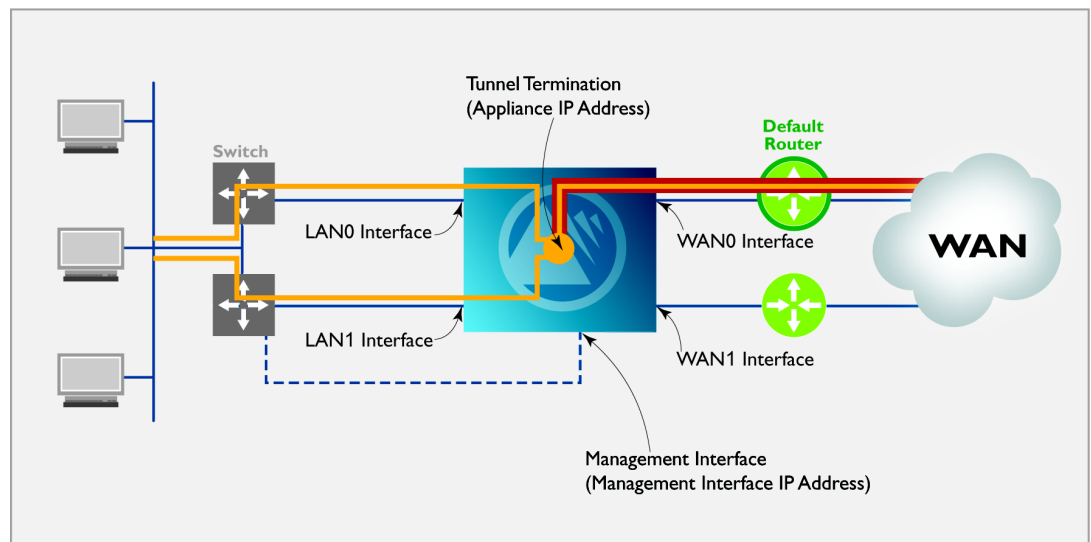


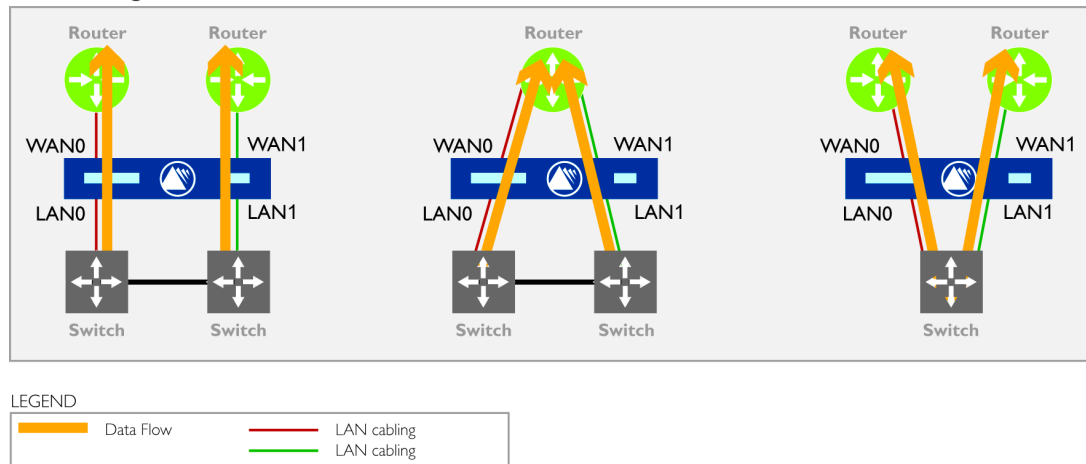
Figure 1-2 In-Line Deployment: Bridge Mode [Bridging with Fail-to-Wire]

NX Series appliances support the four-port configurations for both pass-through and tunneled traffic, as shown next.

Pass-through Traffic Summary

The following 4-port topologies are supported for pass-through traffic:

Pass-through Traffic



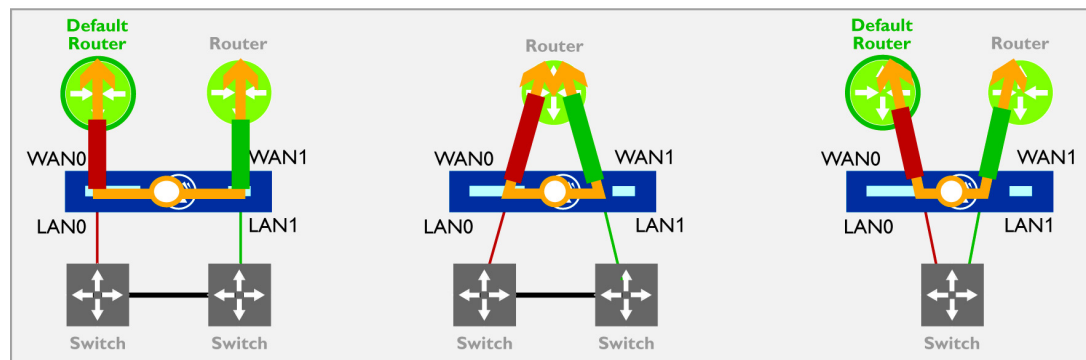
Appliance Placement	<p>Appliance placed in-line between Ethernet LAN switch and WAN router</p> <ul style="list-style-type: none"> • Appliance LAN interface(s) connects to Ethernet LAN switch(es) • Appliance WAN interface(s) connects to WAN router(s)
Fail-Safe Behavior	<p>Fails-to-Wire: The appliance behaves as a crossover cable between the Ethernet LAN switch and the WAN router in any failure scenario (hardware, software, power).</p> <ul style="list-style-type: none"> • LAN0 maps to WAN0 • LAN1 maps to WAN1 • No crossover of traffic in the appliance <p>IMPORTANT: Ensure that the Ethernet LAN's switch and the WAN router have compatible Ethernet interface physical configuration settings (speed and duplex settings). This is to ensure that traffic flows correctly if the Silver-Peak appliance "Fails-to-wire".</p>
IP Addresses	<p>This deployment model requires two IP addresses (on the same or separate subnets)</p> <ul style="list-style-type: none"> • Silver Peak Appliance IP address (to originate and terminate tunnel) • Silver Peak Management IP Address (for appliance configuration and management)
Link Propagation	<p>Default is Enable. When an interface goes down, it forces the interface that's paired with it to fail. For example, if LAN1 goes down, it forces WAN1 to fail.</p>

Tunnelized Traffic Summary

The appliances' WAN next-hops can be configured Active/Active or Active/Backup.

The following 4-port topologies are supported for tunnelized traffic:

Tunnelized Traffic



LEGEND

—	Tunnel	—	LAN cabling
—	Tunnel	—	LAN cabling
—	Device		

Appliance Placement	<p>Appliance placed in-line between Ethernet LAN switch and WAN router</p> <ul style="list-style-type: none"> • Appliance LAN interface(s) connects to Ethernet LAN switch(es) • Appliance WAN interface(s) connects to WAN router(s)
Fail-Safe Behavior	<p>Fails-to-Wire: The appliance behaves as a crossover cable between the Ethernet LAN switch and the WAN router in any failure scenario (hardware, software, power).</p> <p>IMPORTANT: Ensure that the Ethernet LAN's switch and the WAN router have compatible Ethernet interface physical configuration settings (speed and duplex settings). This is to ensure that traffic flows correctly if the Silver-Peak appliance "Fails-to-wire".</p>
IP Addresses	<p>This deployment model requires two IP addresses (on the same or separate subnets)</p> <ul style="list-style-type: none"> • Silver Peak Appliance IP address (to originate and terminate tunnel) • Silver Peak Management IP Address (for appliance configuration and management)
Link Propagation	<p>Default is Enable. When an interface goes down, it forces the interface that's paired with it to fail. For example, if LAN1 goes down, it forces WAN1 to fail.</p>

Out-of-Path Deployments (Router Mode)

Silver Peak supports these typical out-of-path deployments:

- 1 **Out-of-Path with Policy-Based-Routing (PBR Redirection)** See page 16.
- 2 **Out-of-Path with Web Cache Coordination Protocol (WCCP)** See page 17.
- 3 **Out-of-Path with VRRP Peering to WAN Router** See page 18.
- 4 **Out-of-Path with Policy-Based-Routing (PBR) and VRRP Redundant Silver Peak Appliances** See page 19.
- 5 **Out-of-Path with Web Cache Coordination Protocol (WCCP) Redundant Silver Peak Appliances** See page 20.

Out-of-Path with Policy-Based-Routing (PBR Redirection)

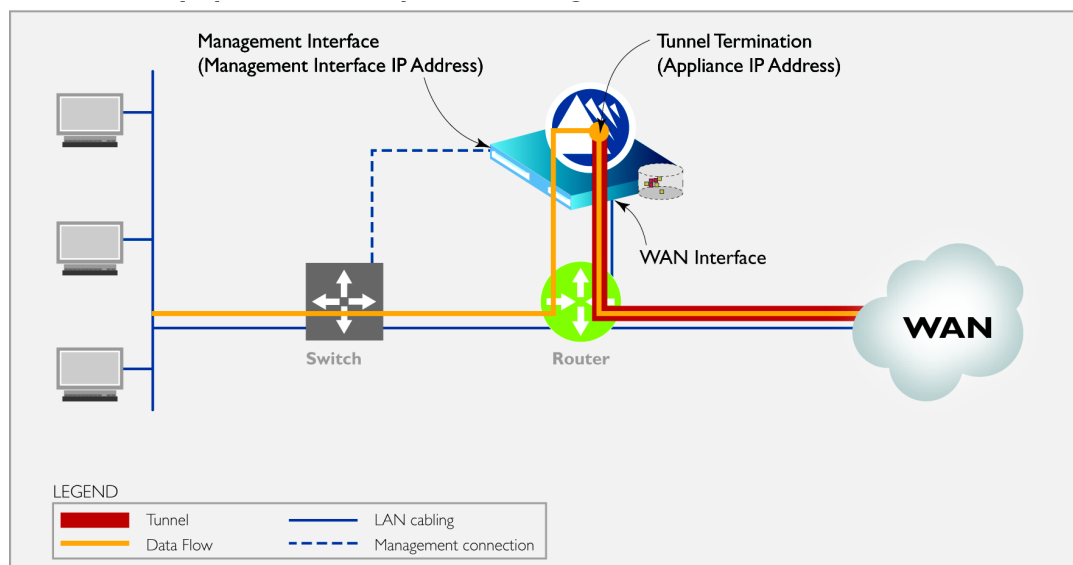


Figure 1-3 Out-of-Path Deployment with Policy-Based Routing (PBR): Router Mode [Spare Router Port Available]

Summary

Appliance Placement	<p>Attached to available router interface:</p> <ul style="list-style-type: none"> • Appliance WAN interface connects to available WAN interface • Do not connect LAN interface
Failure Method	<p>Fails-Open:</p> <ul style="list-style-type: none"> • The appliance behaves as unconnected port in all failure cases (hardware, software, power) • The WAN router sees the link to the appliance go down, Policy-Based-Routing fails, unicast routing forwards traffic normally.
IP Addresses	<p>This deployment model requires two IP addresses (on the same or separate subnets):</p> <ul style="list-style-type: none"> • Silver Peak Appliance IP Address (to originate and terminate tunnel) • Silver Peak Management IP Address (for appliance configuration and management) <p>Configure PBR on WAN router</p> <ul style="list-style-type: none"> • Direct traffic from LAN (subnet/interface) destined for WAN to Silver Peak appliance • Do NOT enable this PBR on the interface to which the Silver Peak appliance connects

Out-of-Path with Web Cache Coordination Protocol (WCCP)

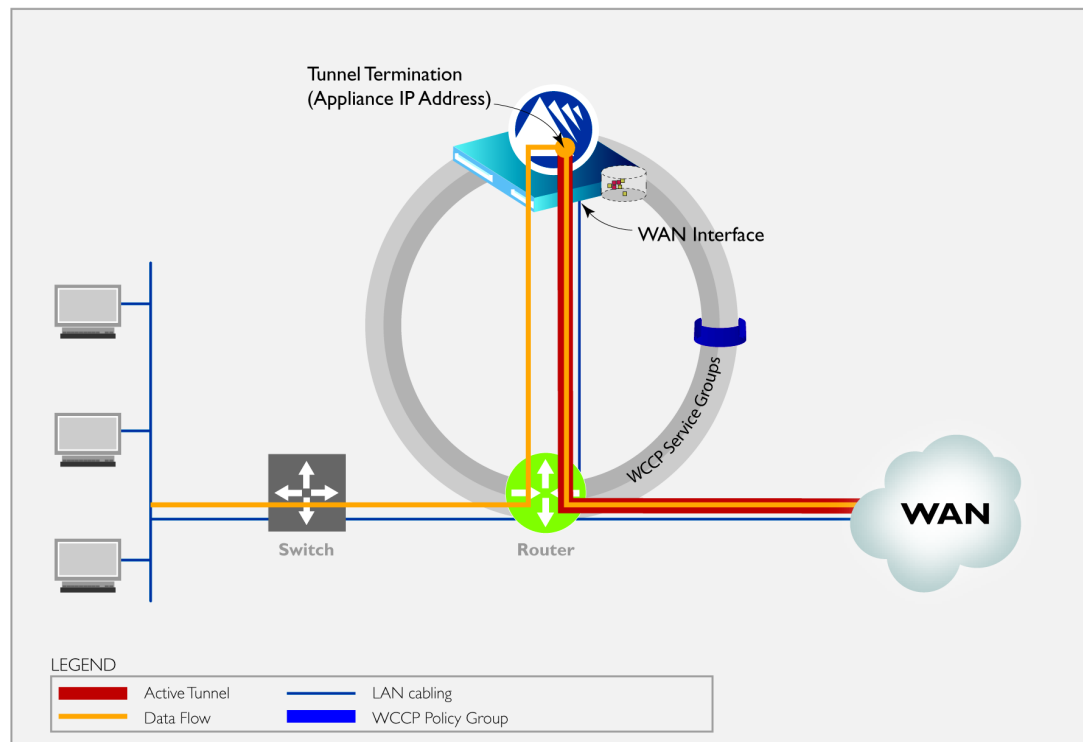


Figure 1-4 Out-of-Path Deployment: Silver Peak Appliance peered with an L3 router using WCCP

Summary

Appliance Placement	<p>Appliance attached in network, reachable by WAN router</p> <ul style="list-style-type: none"> • Appliance WAN interface connects to network • Do not connect LAN interface
Fail-Safe Behavior	<p>WCCP recognizes failed appliance</p> <ul style="list-style-type: none"> • Appliance removed from WCCP Service Groups • WAN router resumes forwarding traffic normally according to its routing tables • Capable of load balancing across multiple NX Series appliances
IP Addresses	<p>This deployment model requires two IP addresses (on the same or separate subnets)</p> <ul style="list-style-type: none"> • Silver Peak Appliance IP Address (to originate and terminate tunnels) • Silver Peak Management IP Address (for appliance configuration and management) <p>Configure WCCP on the Silver Peak appliance and the WAN router</p> <ul style="list-style-type: none"> • Configure two WCCP Service Groups on the Silver Peak appliance (one for TCP and one for UDP) • Configure two WCCP Service Groups on the WAN router (one for TCP and one for UDP)

Out-of-Path with VRRP Peering to WAN Router

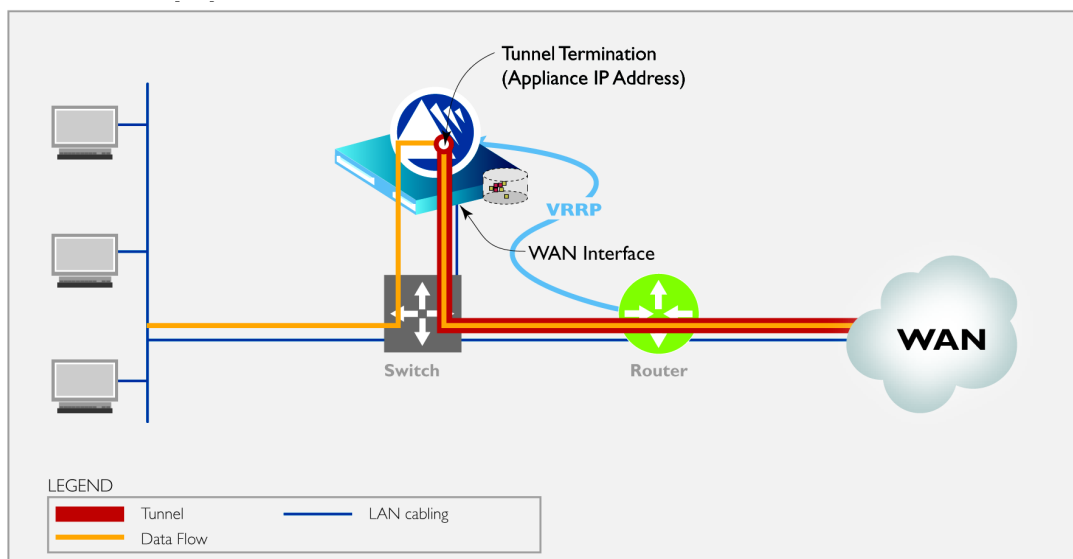


Figure 1-5 Out-of-Path Deployment: Silver Peak Appliance peered with an L3 router using Virtual Router Redundancy Protocol (VRRP)

Summary

Appliance Placement	<p>Appliance shares LAN segment with existing equipment</p> <ul style="list-style-type: none"> • Appliance WAN interface connects to Ethernet LAN switch • Do not connect LAN interface
Failure Method	<p>Fails - Open:</p> <ul style="list-style-type: none"> • The appliance behaves as an unconnected port in all failure cases (hardware, software, power) • WAN router assumes Virtual IP Address and forwards traffic normally
IP Addresses	<p>This deployment model requires three IP addresses:</p> <ul style="list-style-type: none"> • Silver Peak Appliance IP Address (to originate and terminate tunnel) • Silver Peak Management IP Address (for appliance configuration and management) • Virtual IP Address (VIP) shared by Silver Peak appliance and the WAN router <p>The VIP must be the default gateway for the clients and servers on the LAN subnet. NOTE: Typically, this would be the current default gateway, to avoid client reconfigurations.</p> <p>The Silver Peak appliance must share the default gateway VIP with WAN router using VRRP.</p> <ul style="list-style-type: none"> • The Silver Peak appliance must be configured with higher priority and preemption to ensure VRRP reverts to the appliance.

Out-of-Path with Policy-Based-Routing (PBR) and VRRP Redundant Silver Peak Appliances

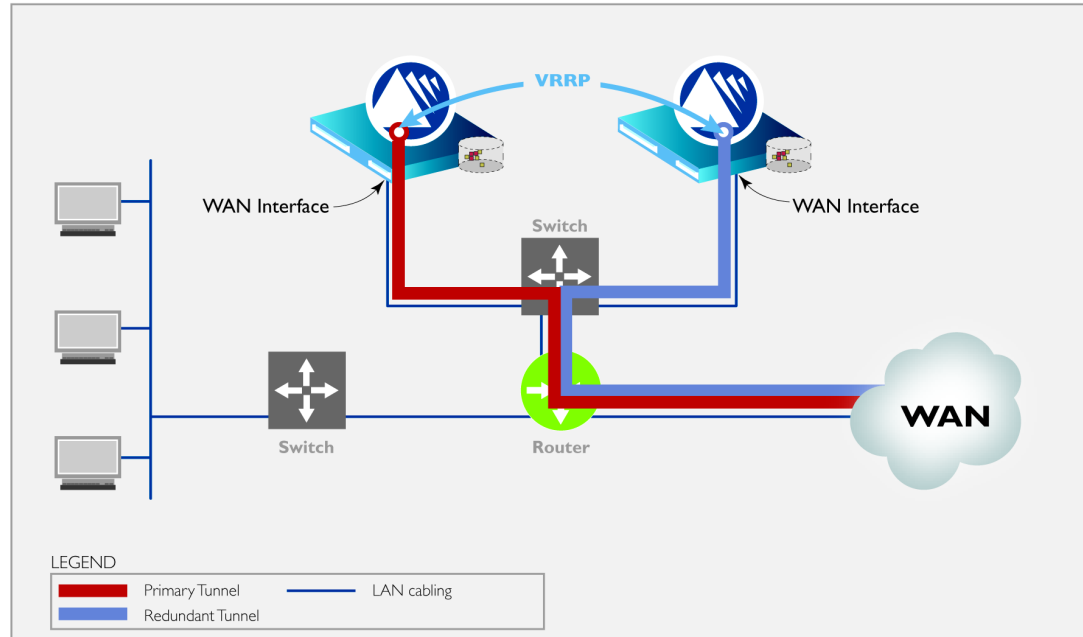


Figure 1-6 Out-of-Path Deployment: Redundant Silver Peak Appliances using Policy-Based-Routing (PBR)

Summary

Appliance Placement

Both appliances are attached to the same available interface via an Ethernet LAN switch:

- Each appliance's WAN interface connects to the Ethernet switch that is connected to the available WAN interface
- Do not connect LAN interface of either appliance

Failure Method

Fails Open:

- The failed appliance behaves as unconnected port in all failure cases (hardware, software, power)
- The backup Silver Peak appliance assumes the Silver Peak Appliance Virtual IP Address. Router forwards traffic to the backup Silver Peak appliance.
- Remote appliances switch to the backup appliance

IP Addresses

This deployment model requires five IP addresses:

- Each appliance needs a Silver Peak Appliance IP Address (to originate and terminate tunnels)
- The two appliances share one Silver Peak Appliance Virtual IP Address for VRRP
- Each appliance needs a Silver Peak Management IP Address (for appliance configuration and management)

Configure PBR on WAN router

- Direct traffic from LAN (subnet/interface) destined for WAN to Silver Peak Appliances' Virtual IP Address
- Do NOT enable this PBR on the interface to which the Silver Peak appliances connect

Out-of-Path with Web Cache Coordination Protocol (WCCP) Redundant Silver Peak Appliances

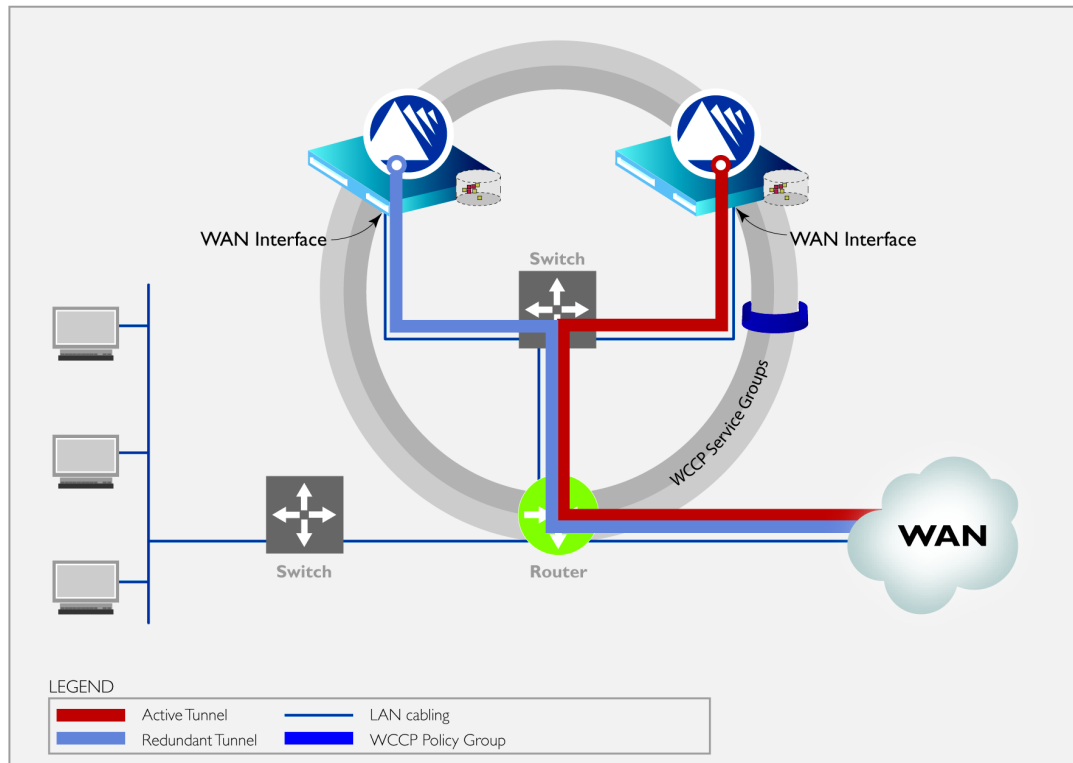


Figure 1-7 Out-of-Path Deployment: Redundant Silver Peak Appliances peered with an L3 router using WCCP

Summary

Appliance Placement	<p>Both appliances are attached in network, reachable by WAN router</p> <ul style="list-style-type: none"> Each appliance's WAN interface connects to network Do not connect LAN interface of either appliance
Fail-Safe Behavior	<p>WCCP recognizes the failed appliance</p> <ul style="list-style-type: none"> Failed appliance is removed from WCCP Service Groups WCCP forwards all traffic to the backup Silver Peak appliance Remote appliances switch to the backup appliance
IP Addresses	<p>This deployment model requires four IP addresses:</p> <ul style="list-style-type: none"> Each appliance needs a Silver Peak Appliance IP Address (to originate and terminate tunnels) Each appliance needs a Silver Peak Management IP Address (for appliance configuration and management) <p>Configure WCCP on the Silver Peak Appliances and the WAN router</p> <ul style="list-style-type: none"> Configure two WCCP Service Groups on each Silver Peak appliance (one for TCP and one for UDP) Configure two WCCP Service Groups on the WAN router (one for TCP and one for UDP)



CHAPTER 2

Installing the Appliance

This chapter describes the procedures for installing your Silver Peak appliance. It describes the preparations you need to make, how to install the appliance in a rack, how to access and run the configuration wizard, and how to verify connectivity.

In This Chapter

- **Before You Begin** See page 22.
- **Rack Mounting the Appliance** See page 25.
- **Connecting the Power and Verifying LEDs** See page 57.
- **Installing the Appliance into the Network** See page 64.
- **Cabling for Configuration Management** See page 68.
- **Running the Initial Configuration Wizard** See page 69.



CAUTION Disconnect AC power before servicing.

Before You Begin

Summary of Installation Tasks

The following table summarizes the tasks, and points you to the appropriate section of this chapter.

	Task	Notes	For detailed instructions, see...
1	Review the specifications and prepare the site	Ensure that the physical environment supports the product requirements. Also, make sure that you have all the equipment and network information needed for initial configuration.	Appendix A “Specifications, Compliance, and Regulatory Statements”. “Site Preparation” on page 22.
2	Rack mount the appliance		“Rack Mounting the Appliance” on page 25.
3	Connect the power	Connect the power cords and check the power LED. Do NOT turn the power on yet.^a	“Connecting the Power and Verifying LEDs” on page 57.
4	Cable the appliance into the network, based on the chosen deployment	Cable the interface ports and verify LED behavior.	“Installing the Appliance into the Network” on page 64.
5	Turn the power on		“Installing the Appliance into the Network” on page 64.
6	Connect the management cables	Instructions detail how to connect to the appliance.	“Cabling for Configuration Management” on page 68.
7	Access and run the GUI-based configuration wizard	Configure the appliance for either Bridge mode or Router mode, based on the selected deployment.	For detailed procedures for each deployment, see Silver Peak NX Series Appliances Network Deployment Guide .

a. Some appliances power up automatically, and therefore are not subject to this step.

Site Preparation

These preparations ensure quick and smooth installation.

- 1 Inspect the package contents and verify them against the packing list.
- 2 Decide on the appliance's location in the network topology.



*For a **brief overview** of network deployments, see “Typical Network Deployments” on page 11 in Chapter 1, “Overview.”*

*For **detailed deployment scenarios**, see the [Silver Peak NX Series Appliances Network Deployment Guide](#).*

3 Survey the physical site.

You should make note of the following for the physical installation of the appliance:

This appliance...	requires...	in a ...
NX-1700	1 Rack Unit (1 RU)	2-post Telco Rack, 4-post Server Rack, rack mount kit for wall or under-desk installation
NX-2500	1 Rack Unit (1 RU)	2-post Telco Rack
NX-2600	1 Rack Unit (1 RU)	4-post Server Rack
NX-2610	1 Rack Unit (1 RU)	4-post Server Rack
NX-2700	2 Rack Units (2 RU)	4-post Server Rack
NX-3500	2 Rack Units (2 RU)	2-post Telco rack or 4-post Server Rack
NX-3600	2 Rack Units (2 RU)	2-post Telco rack or 4-post Server Rack
NX-3700	2 Rack Units (2 RU)	4-post Server Rack
NX-5500	3 Rack Units (3 RU)	4-post Server Rack
NX-5504	3 Rack Units (3 RU)	4-post Server Rack
NX-5700	2 Rack Units (2 RU)	4-post Server Rack
NX-7500	3 Rack Units (3 RU)	4-post Server Rack
NX-7504	3 Rack Units (3 RU)	4-post Server Rack
NX-8504	3 Rack Units (3 RU)	4-post Server Rack
NX-5600	3 Rack Units (3 RU)	4-post Server Rack
NX-7600	3 Rack Units (3 RU)	4-post Server Rack
NX-7700	2 Rack Units (2 RU)	4-post Server Rack
NX-8600	3 Rack Units (3 RU)	4-post Server Rack
NX-8700	2 Rack Units (2 RU)	4-post Server Rack
NX-9610	3 Rack Units (3 RU)	4-post Server Rack
NX-9700	2 Rack Units (2 RU)	4-post Server Rack

- Ensure that sufficient power is available. Supply circuits should be protected by a minimum 15A, maximum 20A circuit breaker.
- On the network devices that will provide the WAN, LAN, and management connections to the appliance, identify the 10/100/1000 Ethernet ports, and determine their speed and duplex settings.
- Verify that the ambient temperature does not exceed 35° C (95° F).
- Make sure you have a standard Phillips screwdriver for installing the brackets and rails.



CAUTION IT (Information Technology) Power System: To ensure safe operation of this equipment, connect only to an AC power source that contains a protective earthing (PE) conductor. IT power systems do not provide adequate grounding and are not recommended.



CAUTION Silver Peak NX Appliances contain embedded hard disk drives. Accordingly, securely rack mount the appliance to stabilize it against excessive vibration.

4 Collect all the data needed for system configuration.

You will need the following information on hand to complete the initial turn up of your Silver Peak appliance. If in doubt, ask your Network Administrator for help.

What the initial configuration wizard requires...	Your information...	Definition
Host name for Appliance		A name you give the appliance that makes it easy for you to remember.
Appliance Mode		Bridge (b) : in-line deployment Router (r) : out-of-path deployment
Administrator password		When you create a user name, ensure that the first character of the name is alphabetical (a-z or A-Z). The remaining characters must include one of the following: <ul style="list-style-type: none"> • alphabetical (upper or lower case) • numerical • dash (-) • underscore (_) • dot (.)
mgmt0 IP Address/Netmask		This is the management interface. It's the name of the actual physical interface. Labeled as such on the appliance.
mgmt0 Next-hop IP Address		The default gateway IP address. Depending on the network topology, this may or may not be the same as the wan0 Next-hop IP Address . If you're unsure, contact your network administrator.
Appliance IP Address / Netmask		The IP address assigned to the Silver Peak Appliance
Second IP / Netmask		For use in a dual-home router configuration.
wan0 Next-hop IP Address		IP address of WAN edge router connected to appliance's WAN interface cable.
LAN Interface speed [Bridge mode only]		The mode (auto, full duplex, half duplex) and the speed (10, 100, 1000 Mbps) of each of the Ethernet interfaces connected to the appliance. By default, the Silver Peak appliance supports auto-sense. If the network device connected to the appliance supports auto-sense, both sides should be configured as auto . Otherwise, both sides need to be configured with the same mode and speed.
LAN Interface duplex mode [Bridge mode only]		
WAN Interface speed		
WAN Interface duplex mode		
MGMT (mgmt0) Interface speed		
MGMT (mgmt0) Interface duplex mode		
LAN Next-hop IP [Bridge mode only; optional]		If your LAN has multiple subnets, you can enter this information for one of the subnets.

The next section describes how to rack mount the appliance.

Rack Mounting the Appliance

This section describes how to rack mount the various appliances:

- **NX-1700** See page 25.
- **NX-2500** See page 27.
- **NX-2600 and/or NX-2610** See page 28.
- **NX-3500** See page 33.
- **NX-3600** See page 34.
- **NX-5500, NX-5504, NX-7500, NX-7504, and NX-8504** See page 41.
- **NX-5600, NX-7600, NX-8600, and NX-9610** See page 45.
- **NX-2700, NX-3700, NX-5700, NX-7700, NX-8700 and NX-9700** See page 51.

NX-1700

The NX-1700 has multiple installation options:

- You can mount the chassis into either a 2-post telco rack or a 4-post server rack.
- With a separate rack mount kit, you can mount the chassis either horizontally under a desk or vertically onto a wall.

The chassis weighs 8.5 lbs. (3.9 kg) and has dimensions of 1.8" H x 17.5" W x 8.2" D (45mm H x 445mm W x 209 mm D).


Installation is easier if one person holds the chassis in position while another attaches the screws.

♦ To mount the NX-1700 into the telco or server rack




Note If your rack requires its own screws and/or nuts, then use those instead of the ones that come with the appliance.

- 1 Attach the cage nuts to the rear of the front rack posts.



- 2 Use the screws to secure the chassis to the outside front of the rack. When using two screws on each side, secure the top and bottom holes for greater stability.



◆ **To mount the individual rack and NX-1700**

- 1 Ensure that the mounting surface (wall or underside of desk) is sturdy enough to support the weight of the NX-1700 and the rack.
- 2 Test fit the rack to the surface to ensure proper fit, and mark the 6 mounting points.



- 3 For wood surfaces only, use the provided self-tapping screws to affix the rack to the mounting surface. If you're using another type of surface, obtain and use the appropriate screws.

After the rack is properly mounted to the surface, you can install the NX-1700.


- 4 Use the supplied square cage nuts to provide the mounting points for the appliance.
- 5 Place the appliance on to the outside face of the rack.
- 6 Secure the appliance to the rack with the supplied cabinet screws.

If only using two screws on each side, secure the top and bottom holes for greater stability.

Exploded view of mounting components



Horizontal positioning under desk



Wall mount positioning



NX-2500

You can mount the chassis in either a 2-post telco rack or a 4-post server rack.


The chassis weighs 14 lbs. (6.4 kg) and has dimensions of 1.7" H x 16.8" W x 14.0" D. Realistically, it requires one person to hold the chassis in position while another attaches the screws.

♦ To mount the NX-2500 into the telco or server rack

Attach one rear rack mounting bracket to each side of the chassis, using the screws provided. If only using two screws per side, make sure that you secure the top and bottom holes on the front ears.

Attach front ears to outside front of telco or server rack.

The appliance ships with four 12-24 x 5/8" screws — 2 for each ear.



NX-2600 and/or NX-2610

Mount the chassis in a 4-post server rack. Because it's rack-specific, no screws are provided for attaching to your particular rack. You'll need eight (8) screws to attach the rails to the server rack and two (2) longer screws to secure the chassis' front ears to the rails' front brackets.

Both chassis have dimensions of 1.7" H x 16.9" W x 22.4" D.

The NX-2600 weighs 22 lbs. (10.0 kg), and the NX-2610 weighs 24.2 lbs. (11.0 kg).




CAUTION Due to the appliance's weight and size, Silver Peak recommends that **two people** mount the chassis in a 4-post server rack, using the chassis rails that ship with the appliance.

Please make sure that the chassis covers and chassis rails are installed on the chassis before you install the chassis into the rack.


Installing the Chassis Rails

Before installing the chassis rail, be sure to remove all external devices and connectors.

- 1 Included in the shipping package are a pair of rail assemblies. In each rail assembly, locate the inner rail and the outer rail.
- 2 Press the locking tab to release the inner rail from its locking position and pull out the inner rail from the rail assembly. The inner rails attach to the chassis and the outer rails install in the rack.



- 3 Locate and align the 3 upper edge holes on each side of the chassis with the 3 corresponding holes on each of the inner rails.



- 4 After placing the rail against the chassis, secure the inner rail with the three silver-colored M4 screws.
- 5 Repeat steps 2 through 4 to install the other rail onto the chassis .


Rack Installation

After you install the inner rails on the chassis, you're ready to install the outer rails of the rail assemblies to the rack.



Note The rails are designed to fit in the racks that have a depth of 26" to 33".

- 1 In the package, locate a pair of front and rear brackets. Note that the brackets are marked with **F** (front) and **R** (rear).
- 2 Secure the front bracket (marked with the letter, **F**) to the outer rail with two **black round-head screws**, inserting a washer between the bracket and each screw.





Notice that the front bracket has a cutaway area at its sides.



- 3 Locate the two buttons on the outer rail and attach the rear bracket, using the washers and remaining black round-head screws.

Yes, the **R** is upside down, relative to the **F** on the front bracket.




- 4 Measure the depth of your rack and adjust the length of the rails accordingly.
- 5 Repeat the same steps for the other outer rail.


- 6 Secure both outer rail assemblies to the rack with the screws and washers appropriate to your rack. Make sure that you secure both the front and rear brackets to the outside of the rack.




- 7 Pull each rail into its fully extended and locked position.



- 8 After you securely install the rail assemblies on the rack, slide the NX-2600 or NX-2610 into the assemblies, as shown below:




When the latch catches in the slot, push with your finger to release it.



The NX-2600 or NX-2610 may not slide into the rack smoothly or easily when installed the first time. Therefore, some adjustment to the slide assemblies might be needed for easiest installation.

- 9 When installed, the appliance bracket should be flush with the rail bracket.
- 10 Secure the appliance to the rack by inserting and tightening the screws in the middle holes on the front side of the appliance, as shown:



NX-3500

You can mount the chassis in either a 2-post telco rack or a 4-post server rack.

The chassis weighs 34 lbs. (15.4 kg) and has dimensions of 3.5" H x 17.9" W x 22.3" D. Realistically, it requires one person to hold the chassis in position while another attaches the screws.

♦ To mount the NX-3500 into the telco or server rack

Attach the chassis to the rack, using the screws provided.

The appliance ships with four 12-24 x 5/8" screws — 2 for each ear.



NX-3600

You can mount the chassis in either a 2-post telco rack or a 4-post server rack.



CAUTION Due to the appliance's weight and size, Silver Peak recommends that **two people** mount the chassis in a 4-post server rack, using the chassis rails that ship with the appliance.

Please make sure that the chassis covers and chassis rails are installed on the chassis before you install the chassis into the rack.

Installing the Chassis Rails

Before installing the chassis rail, be sure to remove all external devices and connectors from the appliance.

- 1 Included in the shipping package are a pair of rail assemblies. In each rail assembly, locate the inner and the outer rail.

Outer rail mounts to the rack




Inner rail mounts to the appliance

- 2 Pull the black latch in the direction of the arrow to release the inner rail, and pull it out from the rail assembly. The inner rails attach to the chassis and the outer rails install in the rack.


Toward rear of chassis

Front end of chassis




IMPORTANT: The latch arrows point toward the front end of the chassis.

- 3 Locate the square hole at the end of the inner rail and the catch at the front side of the appliance chassis.



- 4 Place hole over catch and move the inner rail toward the rear to secure.




- 5 For each side, you'll need to use the following screws:
- One 10#-32 screw to anchor the rail's alignment to the chassis. Here, the size and shape of the hole in the rail is the same as the one in the appliance chassis.
 - Three 10#-32-R screws to secure the rail through keyhole slots. Relative to the anchor screw, one screw is between the alignment screw and the appliance faceplate.

Rear of chassis



Front of chassis




For keyhole slots — 10#-32R

For anchoring alignment — 10#-32

The remaining two screws are between the alignment screw and the rear of the chassis.

Rear of chassis



- 6** Repeat steps 2 through 5 to install the other rail onto the chassis.

Rack Installation

After you install the inner rails on the chassis, you're ready to install the outer rails of the rail assemblies to the rack.



Note The rails are designed to fit in the racks that have a depth of 24" to 36".


- 1 Locate the outside rails.

The end with the longer bar belongs in the **back** of the rack, and ...

...the end with the shorter bar belongs in the **front** of the rack.



- 2 As necessary, measure the depth of your rack and adjust the length of the rails accordingly.
- 3 Secure the front and rear brackets of both outer rail assemblies to the **inside** of the rack.



One screw comes attached to each front piece. If it's not in the top screw hole, move it to that position.


Secure the lower hole with the appropriate screw for your rack.

We used a 4-3-80-068 screw at the lower portion of each front and rear bracket. Notice that the lower screw is larger than the upper screw.


- 4 Pull each rail to its fully extended and locked position.




- 5 After you securely install the rail assemblies on the rack, slide the appliance into the assemblies.



When the appliance can move no further...



...pull with your finger (in the direction of the arrow) to release it.



The appliance may not slide into the rack smoothly or easily when installed the first time. Therefore, some adjustment to the slide assemblies might be needed for easiest installation.

- 6 When installed, the appliance bracket should be flush with the rail bracket.

- 7 Secure the appliance to the rack by inserting and tightening the screw in the lower holes on the front side of the appliance.



Here, we used an M5 x 20 for each appliance ear.

NX-5500, NX-5504, NX-7500, NX-7504, and NX-8504




CAUTION Due to the appliance's weight and size, Silver Peak recommends that **two people** mount the chassis in a 4-post server rack, using the chassis rails that ship with the appliance.

Please make sure that the chassis covers and chassis rails are installed on the chassis before you install the chassis into the rack.


Installing the Chassis Rails

Before installing the chassis rail, be sure to remove all external devices and connectors.


- 1 Included in the shipping package are a pair of rail assemblies. In each rail assembly, locate the inner rail and the outer rail.
- 2 Press the locking tab to release the inner rail from its locking position and pull out the inner rail from the rail assembly. The inner rails attach to the chassis and the outer rails install in the rack.



- 3 Locate the 5 rail buttons on each side of the chassis and locate the 5 corresponding hole on each of the inner rails. Note that one end of the hole is larger than the other end of the hole.



- 4 Align the larger end of each hole against its corresponding button. Once all aligned, push the holes toward their corresponding buttons, placing the rail on the chassis.
- 5 After placing the rail against the chassis, pull the rail forward until the rail buttons lock in the small ends of the corresponding holes.
- 6 Secure the rail to the chassis with a **Type G screw**.



- 7 Repeat steps 2 through 6 to install the other rail onto the chassis.

Rack Installation


After you install the inner rails on the chassis, you are ready to install the outer rails of the rail assemblies to the rack.





Note The rails are designed to fit in the racks that have a depth of 28" to 33".

- 1 In the package, locate a pair of front (short) and rear (long) brackets. Note that the brackets are marked with **Up/Front Arrows** (front) and **Up/Rear Arrows** (rear).
- 2 Secure the front (short) bracket (marked with the **Up/Front** arrows) to the outer rail with two **Type G screws**.


- 3 Locate the two buttons on the outer rail and attach the rear (long) bracket to it by sliding the opening of the rear rail through the button.



- 4 Measure the depth of your rack and adjust the length of the rails accordingly.
- 5 Repeat the same steps to install the other outer rail on the chassis.
- 6 Secure both outer rail assemblies to the rack with **Type H screws** and **Type I washers**.





- 7 After you securely install the rail assemblies on the rack, slide the NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504 into the assemblies, as shown below:



The NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504 may not slide into the rack smoothly or easily when installed the first time. Therefore, some adjustment to the slide assemblies might be needed for easiest installation.

- 8 Secure the appliance to the rack by tightening the screws on the front side of the NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504, as shown:



NX-5600, NX-7600, NX-8600, and NX-9610




CAUTION Due to the appliance's weight and size, Silver Peak recommends that **two people** mount the chassis in a 4-post server rack, using the chassis rails that ship with the appliance.

Please make sure that you install the chassis covers and chassis rails onto the chassis before you install the chassis into the rack.

Installing the Chassis Rails

Before installing the chassis rail, be sure to remove all external devices and connectors from the appliance.


- 1 Included in the shipping package are a pair of rail assemblies. In each rail assembly, locate the inner and the outer rail.



- 2 Pull the black latch in the direction of the arrow to release the inner rail, and pull it out from the rail assembly. The inner rails attach to the chassis and the outer rails install in the rack.

Toward rear of chassis

Front end of chassis




- 3 Locate the square hole at the end of the inner rail and the catch at the front side of the appliance chassis.




- 4 Place hole over catch and move the inner rail toward the rear to secure.



- 5 For each side, you'll need to use the following screws:
 - One 10#-32 screw to anchor the rail's alignment to the chassis. Here, the size and shape of the hole in the rail is the same as the one in the appliance chassis.
 - Three 10#-32-R screws to secure the rail through keyhole slots. Relative to the anchor screw, one screw is between the alignment screw and the appliance faceplate.



The remaining two screws are between the alignment screw and the rear of the chassis.



- 6 Repeat steps 2 through 5 to install the other rail onto the chassis.

Rack Installation

After you install the inner rails on the chassis, you're ready to install the outer rails of the rail assemblies to the rack.



Note The rails are designed to fit in the racks that have a depth of 24" to 36".


- 1 Locate the outside rails.

The end with the longer bar belongs in the **back** of the rack, and ...

...the end with the shorter bar belongs in the **front** of the rack.



- 2 As necessary, measure the depth of your rack and adjust the length of the rails accordingly.
- 3 Secure the front and rear brackets of both outer rail assemblies to the **inside** of the rack.




The lower screws are already attached to the outer rails' brackets. Secure the upper hole with the appropriate screws for your rack. Here, we used a 4-3-80-068 screw at the top of each front and rear bracket.


- 4 Pull each rail to its fully extended and locked position.




- 5 After you securely install the rail assemblies on the rack, slide the appliance into the assemblies.



When the appliance can move no further...



...pull with your finger (in the direction of the arrow) to release it.



The appliance may not slide into the rack smoothly or easily when installed the first time. Therefore, some adjustment to the slide assemblies might be needed for easiest installation.

- 6 When installed, the appliance bracket should be flush with the rail bracket.

- Secure the appliance to the rack by inserting and tightening the screw in the lower holes on the front side of the appliance.



Here, we used an M5 x 20 for each appliance ear.

NX-2700, NX-3700, NX-5700, NX-7700, NX-8700 and NX-9700




CAUTION Due to the appliance's weight and size, Silver Peak recommends that **two people** mount the chassis in a 4-post server rack, using the chassis rails that ship with the appliance.

Please make sure that you install the chassis covers and chassis rails onto the chassis before you install the chassis into the rack.

Installing the Chassis Rails

Before installing the chassis rail, be sure to remove all external devices and connectors from the appliance.


- 1 Included in the shipping package are a pair of rail assemblies. In each rail assembly, locate the inner and the outer rail.




- 2 Pull the black latch in the direction of the arrow to release the inner rail, and pull it out from the rail assembly. The inner rails attach to the chassis and the outer rails install in the rack.

Toward rear of chassis


Front end of chassis




- 3 Locate the square hole at the end of the inner rail and the catch at the front side of the appliance chassis.



- 4 Place hole over catch and move the inner rail toward the rear to secure.




- 5 For each side, you'll need to use the following screws:
 - One loose 10#-32 screw to anchor the rail's alignment to the chassis. Here, the size and shape of the hole in the rail is the same as the one in the appliance chassis.
 - Three 10#-32-R screws (previously attached to the chassis) for securing the rail through keyhole slots. Relative to the anchor screw, one screw is between the alignment screw and the appliance faceplate.



The remaining two screws are between the alignment screw and the rear of the chassis.

Rear of chassis



- 6 Repeat steps 2 through 5 to install the other rail onto the chassis.

Rack Installation

After you install the inner rails on the chassis, you're ready to install the outer rails of the rail assemblies to the rack.



Note The rails are designed to fit in the racks that have a depth of 24" to 36".


- 1 Locate the outside rails.

The end with the longer bar belongs in the **back** of the rack, and ...

...the end with the shorter bar belongs in the **front** of the rack.



- 2 As necessary, measure the depth of your rack and adjust the length of the rails accordingly.
- 3 Secure the front and rear brackets of both outer rail assemblies to the **inside** of the rack.




One screw comes attached to each front piece. If it's not in the top screw hole, move it to that position.

Secure the lower hole with the appropriate screw for your rack. Here, we used a 4-3-80-068 screw at the lower portion of each front and rear bracket. Notice that the lower screw is larger than the upper screw.


- 4 Pull each rail to its fully extended and locked position.




- 5 After you securely install the rail assemblies on the rack, slide the appliance into the assemblies.



When the appliance can move no further...




...pull with your finger (in the direction of the arrow) to release it.



The appliance may not slide into the rack smoothly or easily when installed the first time. Therefore, some adjustment to the slide assemblies might be needed for easiest installation.

- 6 When installed, the appliance bracket should be flush with the rail bracket.

- 7 Secure the appliance to the rack by inserting and tightening the screw in the lower holes on the front side of the appliance.



Here, we used an M5 x 20 for each appliance ear.

Connecting the Power and Verifying LEDs


This section describes how to connect the power and verify by LEDs that the power is on:

- **NX-1700** See page 57.
- **NX-2500** See page 57.
- **NX-2600 and NX-2610** See page 58.
- **NX-3500** See page 59.
- **NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504** See page 60.
- **NX-5600, NX-7600, NX-8600, or NX-9610** See page 62.
- **NX-2700, NX-3700, NX-5700, NX-7700, NX-8700 and NX-9700** See page 63.

After you connect the power, each appliance powers up and goes into Bypass (Fail-to-Wire) mode.


NX-1700

- 1 Connect the power cable to the back of the NX-1700 appliance.
- 2 Connect the other end of the power cable to your local power source.



NX-2500

- 1 Connect the power cable to the back of the NX-2500 appliance.
- 2 Connect the other end of the power cable to your local power source.

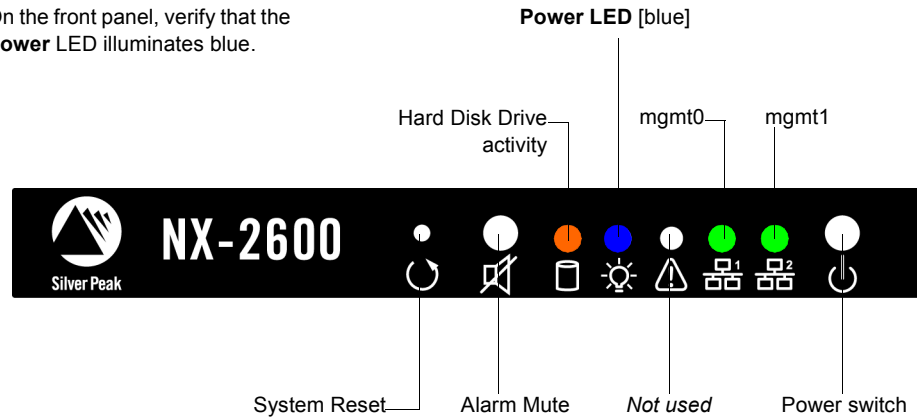


NX-2600 and NX-2610

- 1 Connect the power cable to the back of the NX-2600 or NX-2610 appliance.
- 2 Connect the other end of the power cable to your local power source.

NX-2600 and NX-2610


On the front panel, verify that the **Power LED** illuminates blue.



NX-3500

- 1 Connect each power cable to the back of the Silver Peak chassis, on the left side, and then connect the other end of the power cable to your local power source. For redundancy, plug in both cords, preferably to different sources.

Rear View of NX-3500 chassis - LEFT side




If both power cords aren't connected, an audible alarm sounds. To silence, plug in all power sources.

- 2 Verify that the Power LEDs illuminate solid green.

The Silver Peak NX-3500 has three Power LEDs:

- one Power LED on the front of the chassis, and
- one by each of the two power cord receptacles on the back panel (as described in Step 1).


Partial View of NX-3500 Front Panel



NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504

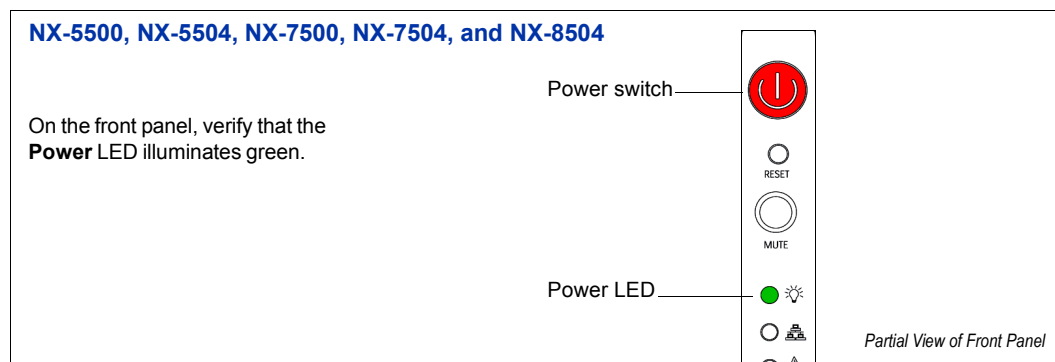
Connect the power as follows:

- 1 Connect each power cable to the back of the Silver Peak chassis, on the left side, and then connect the other end of the power cable to your local power source. For redundancy, plug in all power cords, preferably to different sources.

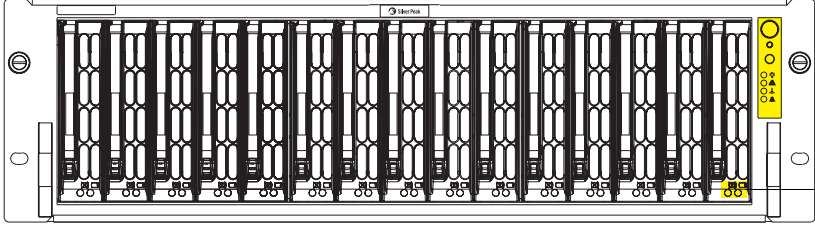


If all power cords are not connected, an audible alarm sounds. To silence, make sure to plug in all power sources.

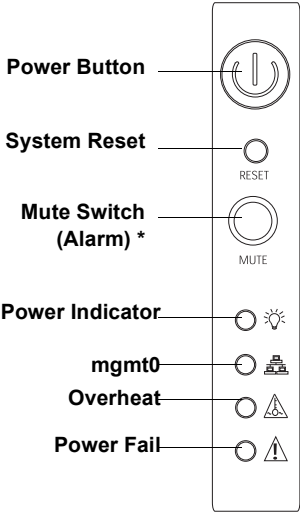
- 2 Verify that the **Power LEDs** illuminate solid green.



The Silver Peak NX-5500, NX-5504, NX-7500, NX-7504, and NX-8504 have one **Power LED** on the front of the chassis and one by each of the 3 power cord receptacles on the back panel (as described in the previous step).



Partial View of NX-5500, NX-5504, NX-7500, NX-7504, or NX-8504 Front Panel



Power Button

System Reset

Mute Switch (Alarm) *

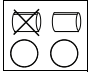
Power Indicator

mgmt0

Overheat

Power Fail

LED	Color	Condition	Description
HDD	Amber	Blink	HDD Activity
		Off	No Activity
Power	Green	On	System On
		Off	System Off
mgmt0 [interface]	Green	On	Linked
		Blink	mgmt0 Activity
		Off	Disconnected
Overheat	Red	On	System Overheat
		Off	System Normal
System Alert/ Power Failure	Red	On	1 or more PWR modules failure
		Off	System Normal



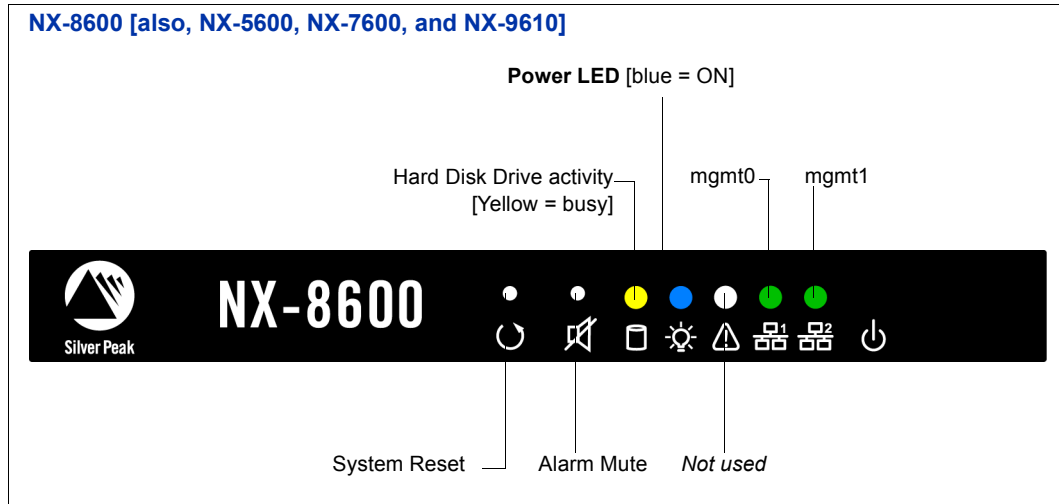
HDD

* Mute Switch (Alarm) — After the Alarm sounds, press this button to reactivate the function

NX-5600, NX-7600, NX-8600, or NX-9610

The LEDs for all three appliances are configured identically.


- 1 Connect the power cable to the back of the appliance.
- 2 Connect the other end of the power cable to your local power source.



- 3 On the front panel, verify that the **Power** LED illuminates blue.

NX-2700, NX-3700, NX-5700, NX-7700, NX-8700 and NX-9700

- 1 Connect the power cables to the back of the appliance.



- 2 Connect the other end of each power cable to your local power source.
- 3 On the front panel, verify that the **Power** LED illuminates blue..

Front view of NX-9700 / NX-8700

Illuminates **red** when a power supply is disconnected or off

Alarm mute

Power LED [blue = ON]

System reset

Power button

[plugging the power cords in automatically powers up the appliance]



Installing the Appliance into the Network

This section describes how to install the appliance into your network and verify connectivity before you run the initial configuration wizard from within a browser.

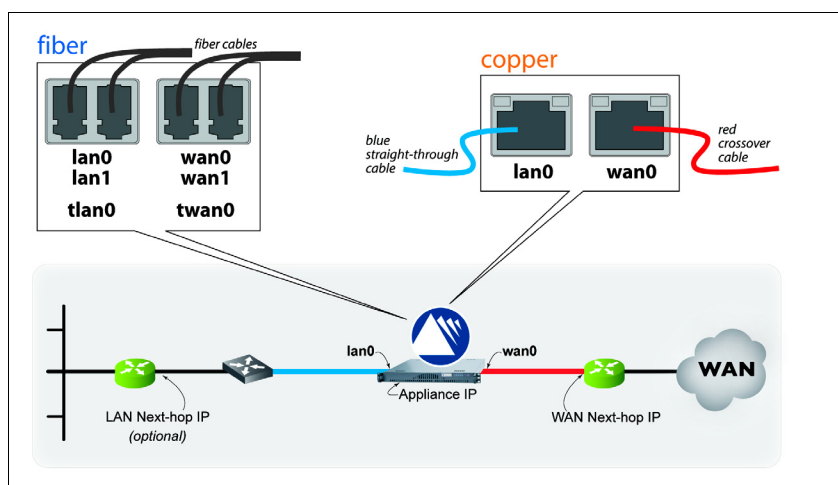


Tip For a **brief overview** of network deployments, see *“Typical Network Deployments” on page 11*. For **detailed in-line and out-of-path deployment scenarios and configuration instructions**, see the *Silver Peak NX Series Appliances Network Deployment Guide*.

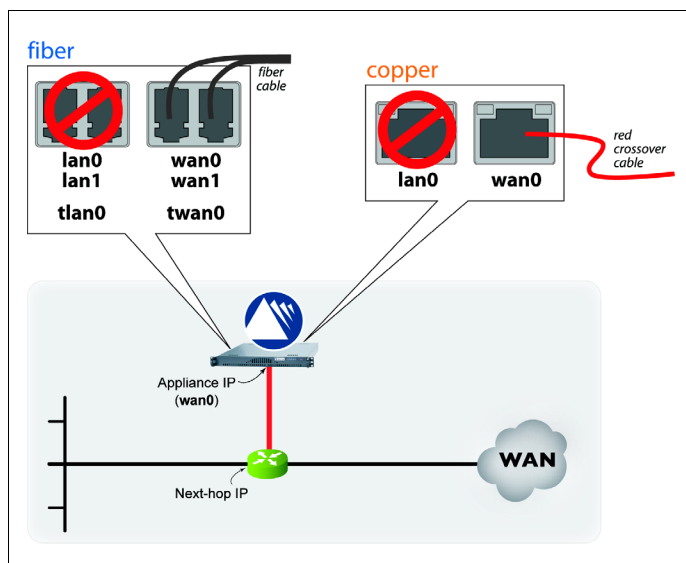
♦ To install the Silver Peak Appliance into the network

- 1 Identify the **LAN0**, **WAN0**, and **MGMT0** ports on the back of the appliance. Connect the supplied network cables from the appliance to the identified network equipment ports.

Bridge Mode



Router Mode



- Silver Peak follows the convention of using a blue (straight-through) cable for the **MGMT0** interface, a red (crossover) cable for the **WAN0** interface, and a yellow (straight-through) cable for the **LAN0** interface.

- In most cases, the WAN crossover cable is appropriate. If you have difficulty verifying connectivity in the next step, then switch to a WAN straight-through cable (not supplied).
- All four copper RJ-45 Gigabit Ethernet interfaces (**LAN0**, **WAN0**, **MGMT0**, and **MGMT1**) in the Silver Peak appliance support auto MDI/MDI-X and are auto-sensing.
- For fiber interfaces, consult this table:

	1 Gbps Fiber Interfaces		10 Gbps Fiber Interfaces	
	lan0 / wan0 Fiber Support	lan0 / wan0 Fail-to-Close	tlan0 / twan0 Fiber Support	tlan0 / twan0 Fail-to-Close
NX-8700	<ul style="list-style-type: none"> • 4 interfaces • LC connectors • Support multi-mode 50μ fiber / 62.5μ fiber 	yes	<ul style="list-style-type: none"> • 2 interfaces • LC connectors • Support multi-mode 50μ fiber 	no
NX-9610		no		no
NX-9700		yes		no

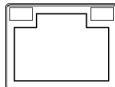
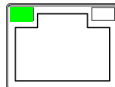
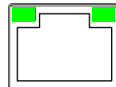

2 Verify connectivity between network devices.

This applies to Bridge mode (in-line deployment) only.

- a** Ping a host on the remote (WAN) side of the Silver Peak appliance from a host on the local (LAN) side of the appliance.
- b** If you are not able to verify connectivity, check the cabling. Do not proceed to the next step if connectivity still cannot be verified.

- c Check the state of the various management interface port LEDs, per the following charts.

	mgmt0 & mgmt1				
	Not connected	10 Mbps	100 Mbps	1000 Mbps	Auto
NX-1700 Speed = solid -- Link/Activity: blinking = traffic	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed 	
NX-2500 Speed = solid -- Link/Activity: blinking = traffic	Speed Link/ Activity 	Speed Link/ Activity 	Speed Link/ Activity 	Speed Link/ Activity 	
NX-2600 NX-2610 Speed = solid -- Link/Activity: blinking = traffic	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed
NX-3500 [blinking only]					
NX-5500 NX-5504 NX-7500 NX-7504 NX-8504 Link = solid -- Activity: blinking = traffic	Link Activity 	Link Activity 	Link Activity 	Link Activity 	Link: ON for 1000 Mbps and 100 Mbps only. (All full duplex)
NX-5600 NX-7600 NX-8600 NX-9610 Speed = solid -- Link/Activity: blinking = traffic	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed 	Link/ Activity Speed

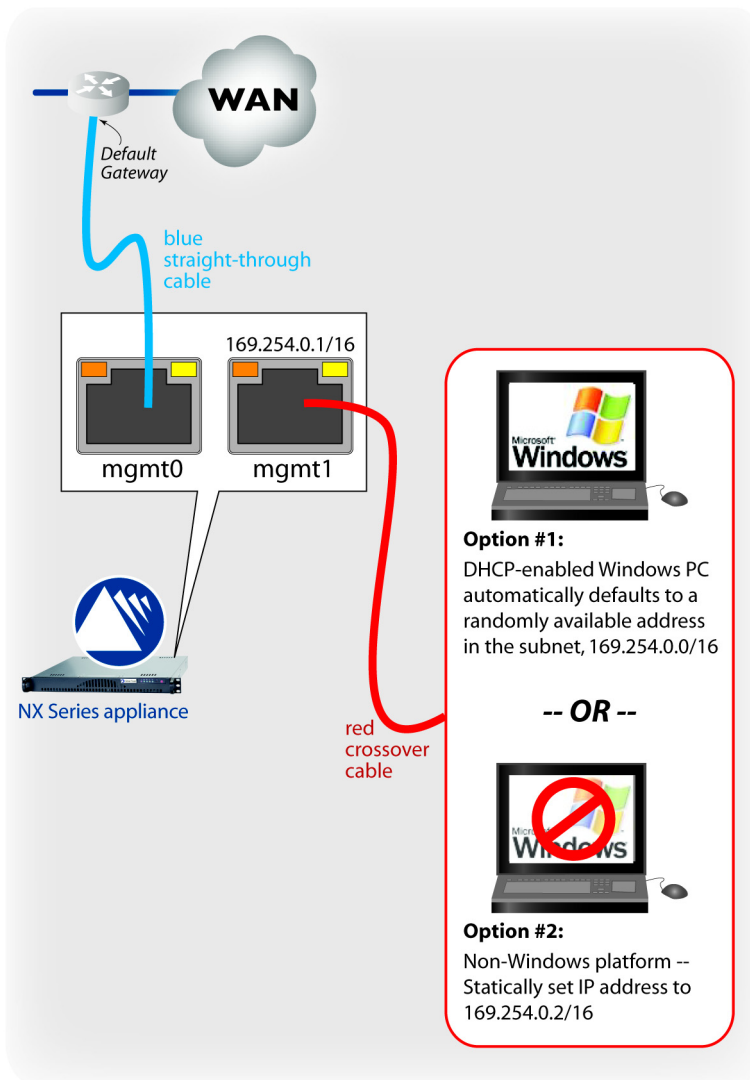
	mgmt0 & mgmt1				
	Not connected	10 Mbps	100 Mbps	1000 Mbps	Auto
<div><div><div>NX-2700</div><div>NX-3700</div><div>NX-5700</div><div>NX-7700</div><div>NX-8700</div><div>NX-9700</div></div><div>Speed = solid</div><div>--</div><div>Link/Activity: solid = link blinking = traffic</div></div> <div><div>Link/ Activity</div><div>Speed</div></div>	<div><div>Link/ Activity</div><div>Speed</div></div>	<div><div>Link/ Activity</div><div>Speed</div></div>	<div><div>Link/ Activity</div><div>Speed</div></div>		

Cabling for Configuration Management

This section describes how to cable your laptop to the appliance for initial configuration.

After you power up the appliance, you need to establish a management connection and then perform the initial configuration by accessing the wizard in the browser-based graphical user interface (GUI).

Silver Peak assigns the **MGMT1** ethernet interface a default IP address of **169.254.0.1** and a subnet mask of **/16** (the same as **255.255.0.0**).



Both RJ-45 Gigabit Ethernet interfaces (**MGMT0** and **MGMT1**) on the NX Series appliances support MDI/MDI-X and auto-negotiation.

- 1 Identify the **MGMT0** and **MGMT1** interfaces on the rear of the appliance.
- 2 Using the supplied **blue** straight-through cable, connect the appliance's **MGMT0** interface to the identified network equipment interface.
- 3 Using the supplied **red** crossover cable, connect the appliance's **MGMT1** interface to your PC.

Running the Initial Configuration Wizard

The Silver Peak Systems Appliance Manager has the following system requirements:

- Operating Systems: Windows XP, Windows 2000, or Windows Vista
 - Web Browser: Internet Explorer 6.0, Mozilla Firefox 2.0
- 1 Verify that you meet all the system requirements. If necessary, update your operating system and/or browser.
 - 2 Start a browser session to connect from your laptop to **169.254.0.1**. The initial configuration page appears.

When you first log on, each appliance's name follows the pattern, **silverpeak-[6 digits of MAC address]**.

The hostname is limited to a maximum of 24 characters.

The screenshot shows the 'Administration - Wizard (Page 1/2)' interface. At the top, the Silver Peak logo is on the left, and 'Logged in as: admin (logout)' is on the right. Below the logo, the appliance name 'silverpeak-85e108' and IP '10.0.60.12' are displayed. The main form area contains several sections: 'Hostname' with a text box containing 'silverpeak-85e108'; 'Appliance ID' with a dropdown set to '1' and a note '1..65534 (The Appliance ID is a network-unique number)'; 'Mode' with radio buttons for 'Router' and 'Bridge' (selected); 'Admin Password (optional)' with 'Old', 'New', and 'Confirm' fields; 'mgmt0' network settings including a 'DHCP' checkbox, 'IP Address / Netmask' (0.0.0.0 / 1), 'Next-hop IP Address' (0.0.0.0), and 'speed/duplex' (auto/auto); and 'Date/Time' settings including 'Date' (2007/02/15), 'Time' (14:26:29), and 'Time Zone' (GMT). Navigation buttons '< Back', 'Next >', and 'Cancel' are at the bottom. A copyright notice 'Copyright © 2004-2007 Silver Peak Systems, Inc. All rights reserved.' is at the very bottom.

After you finish configuring, you can also see the results of this field's entry on the **Configuration - IP Route** page. To access it, go to **Configuration > Networking > IP Route**.



Note If you want to access this page in the future, just go to the **Administration** menu, and select **Initial Config Wizard**.

- 3 Based on the information you collected in the worksheet earlier in the chapter, complete the first page of this two-page wizard.

4 Click **Apply**. The last page of the wizard appears.

The page content depends on whether you selected **Bridge** or **Router** for **Mode**, on the first page.

How the page looks if you selected **Bridge** mode...

lan1 and wan1 are inaccessible if you have a 2-port appliance.

In Bridge mode, you can see this field's entry on the **Configuration - IP Datapath Route** page. To access it, go to **Configuration > Networking > IP Datapath Route**.

How the page looks if you selected **Router** mode...

The Second IP address and its Next-hop IP are for configuring dual-home router mode.

Skip Wizard takes you to the Appliance Manager Home page without saving a new configuration. On initial configuration, only use if Support advises you to.

- 5 Complete the fields, and click **Apply**. When the appliance asks permission to reboot, allow it.

To begin implementing immediately, see the detailed instructions for the specific scenarios in the *Silver Peak NX Series Appliances Network Deployment Guide*.



CHAPTER 3

The Appliance Manager

This chapter describes the structure, content, and features of the Appliance Manager. It also familiarizes you with elements, hierarchy, and conventions employed by the Appliance Manager's interface.

In This Chapter

- **Accessing the Appliance Manager** See page 74.
- **Guided Tour of the Appliance Manager** See page 76.

Accessing the Appliance Manager

The Silver Peak Systems Appliance Manager has the following system requirements:

- **Operating Systems:** Windows XP, Windows 2000, or Windows Vista
- **Web Browsers:** Internet Explorer, Mozilla Firefox


Verify that you meet all the system requirements. If necessary, update your operating system and/or browser.

- 1 In the browser, enter **http://<mgmt0 IP Address>/** to access the Appliance Manager. This is the address you configured on the first page of the initial configuration wizard.



Note The default web setting allows for both **http** and **https**, so you can use either at this stage.

A security alert pop-up window displays.

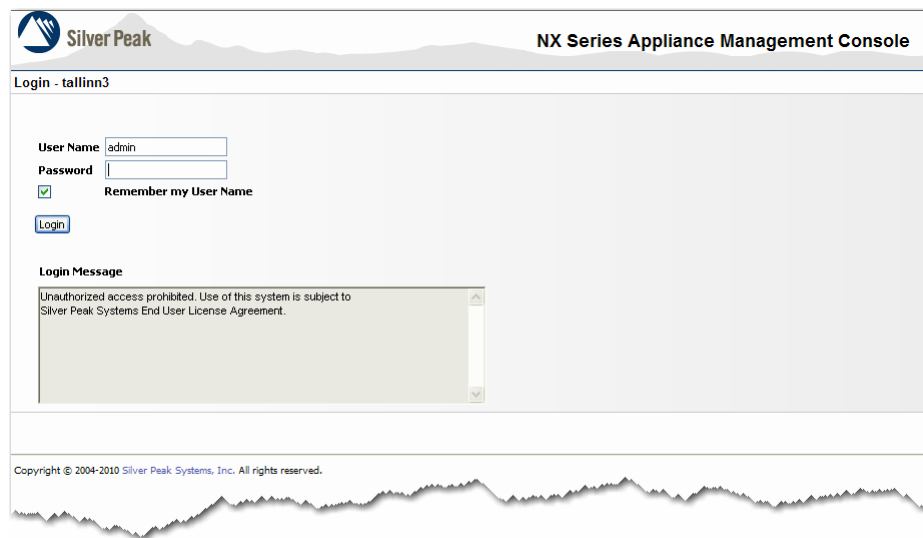


Note If a message displays telling you that the page cannot be displayed, you may be trying to access the appliance across a firewall. In that case, it may be necessary to open TCP ports 80 and 443 to/from the management IP address.

- 2 Click **OK** to proceed. Another **Security Alert** window displays.

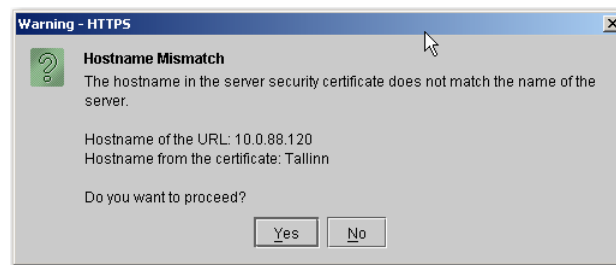


- 3 Click **Yes**. The login page displays.



The image shows the Silver Peak NX Series Appliance Management Console login page. The page has a header with the Silver Peak logo and the title "NX Series Appliance Management Console". Below the header, the text "Login - tallinn3" is displayed. The login form includes fields for "User Name" (containing "admin") and "Password". There is a checkbox labeled "Remember my User Name" which is checked. A "Login" button is located below the password field. A "Login Message" section contains a warning: "Unauthorized access prohibited. Use of this system is subject to Silver Peak Systems End User License Agreement." The footer of the page states "Copyright © 2004-2010 Silver Peak Systems, Inc. All rights reserved."

If your appliance name and IP address are not mapped in your DNS server, you may see a message like the following one. If so, click **Yes** to proceed.



- 4 In the login screen, the **User Name** defaults to **admin**. In the **Password** field, enter the password you set for the **admin** user in the configuration wizard. The **Home** page displays.

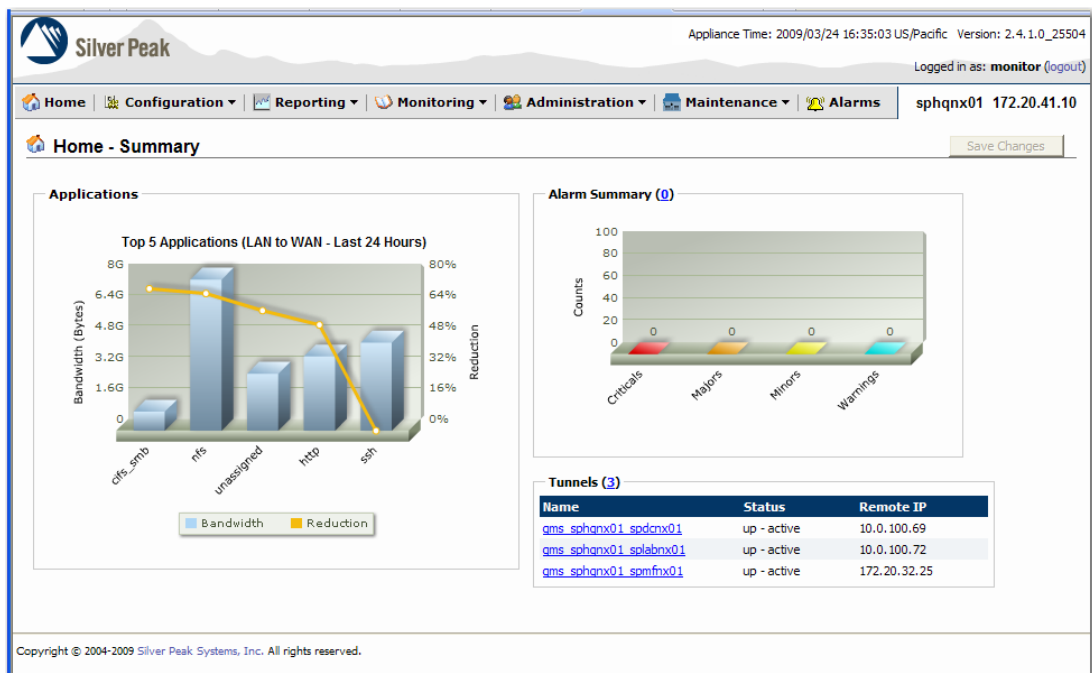
Guided Tour of the Appliance Manager

This section describes navigational and screen elements that are common throughout the user interface. It also describes common operations such as selecting and deleting items from a table, sorting columns, and refreshing screens:

- **The Appliance Manager Home Page** See page 76.
- **Banners** See page 80.
- **Menu Structure** See page 80.
- **Managing Tabular Data** See page 88.
- **Netmask Notations** See page 88.
- **Date and Time Conventions** See page 90.
- **Secure Access Methods** See page 90.
- **Guidelines for Creating Passwords** See page 90.
- **Object Names** See page 90.
- **Saving Your Configuration** See page 91.
- **Definition Help** See page 91.

The Appliance Manager Home Page

The Appliance Manager **Home - Summary** page displays summary information for applications, current alarms, and tunnels. Additionally, it provides links to the main pages from which each summary derives.



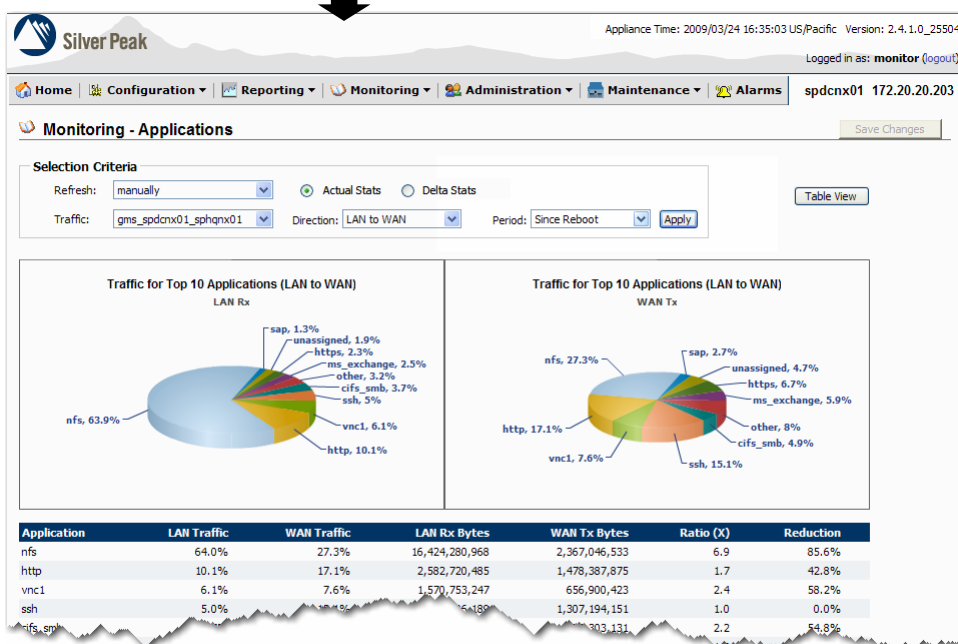
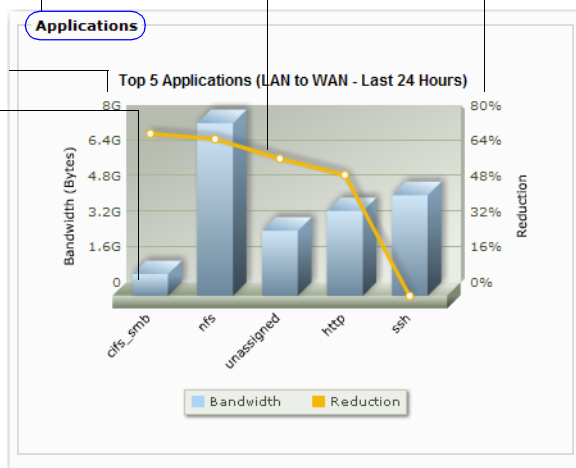
Applications Section

The **Applications** section lists the **Top 5 Applications**, from LAN to WAN, for the last 24 hours.

Clicking the link, **Applications**, opens the **Monitoring - Applications** page

Right Y-axis shows the percent reduction in bandwidth used. This axis is paired with the line.

Left Y-axis shows the bandwidth used. This axis is paired with the bars.




For more information, see [Chapter 4, "Configuring Host Settings."](#)

Current Alarms Area

Whereas each page's banner summarizes all current alarms, the **Home** page's **Alarm Summary** area individually lists all unacknowledged and uncleared alarms.

Summarizes the total number of alarms and links to the **Alarms - Current Alarms** page.



The screenshot shows the Silver Peak Appliance Manager interface. The top navigation bar includes links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as admin. The main content area is titled "Alarms - Current Alarms" and includes a sub-section "Alarm Summary (1)" with a bar chart showing counts for Criticals (0), Majors (1), Minors (0), and Warnings (0). Below this is a table of alarms with columns for Seq No., Date/Time, Type, Severity, Source, Description, Clear, and Ack. A single alarm is listed with Seq No. 9, Date/Time 2007/11/06 09:55:18, Type TUN, Severity Major, Source Hawk2Heron, and Description Tunnel Health is Unhealthy. At the bottom of the table are buttons for Apply, Clear All, Ack All, and Unack All.

Seq No.	Date/Time	Type	Severity	Source	Description	Clear	Ack
9	2007/11/06 09:55:18	TUN	Major	Hawk2Heron	Tunnel Health is Unhealthy	<input type="checkbox"/>	<input type="checkbox"/>



For more information, see [Chapter 16, "Monitoring Alarms."](#)

Tunnels Section

The **Tunnels** section lists all tunnels by name. It includes each tunnel's status, along with the IP address that terminates the tunnel's remote end. Clicking the number in parentheses in the header takes you directly to the **Configuration - Tunnels** page.

Summarizes the total number of tunnels, and links to the **Configuration - Tunnels** page.

Appliance Time: 2009/03/24 16:35:03 US/Pacific Version: 2.4.1.0_25504
Logged in as: admin (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms

Hawk-3500 10.0.42.76

Configuration - Tunnels

Total Tunnels: 1

Name	Status	Admin	Local IP	Remote IP	Max BW(cfg/cur) Kbps	Min BW(cfg/cur) Kbps	IPSec Enabled	Up Time
<input type="checkbox"/> Hawk2Heron	up - degraded	up	192.168.1.2	192.168.0.2	AUTO/10000	32/32	no	7h 34m 18s

Remove Selected Add

© 2004-2009 Silver Peak Systems, Inc. All rights reserved.

Possible states for tunnels are as follows:

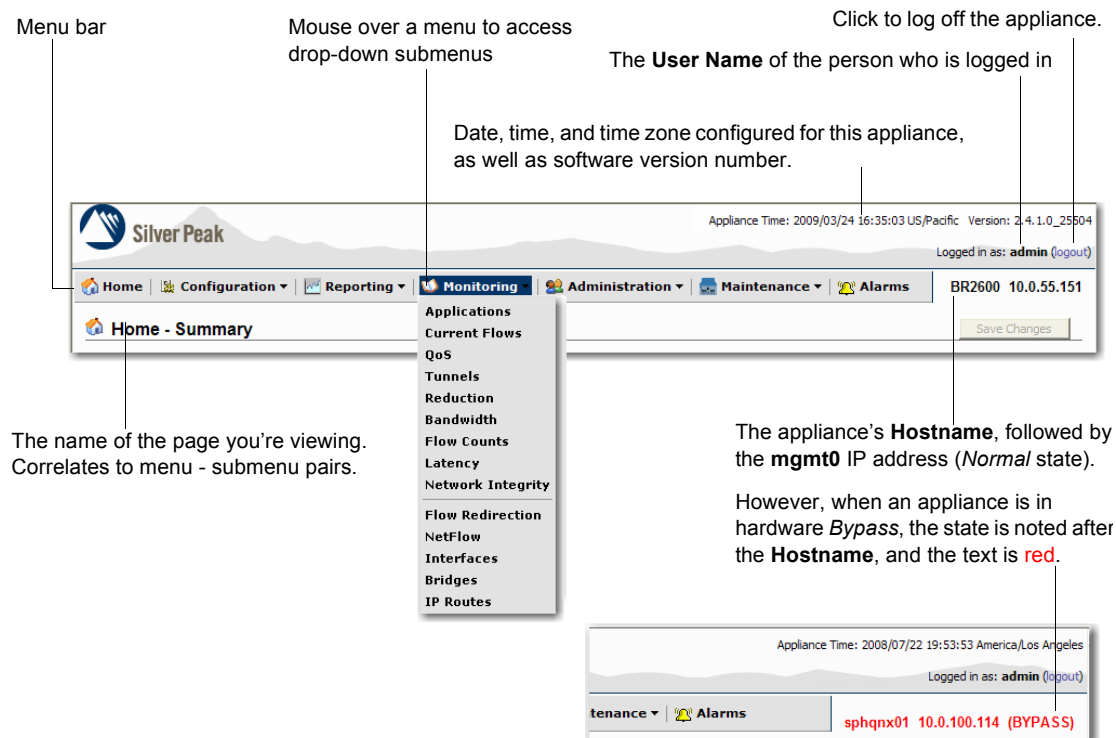
Tunnel Status	Description
Down	The tunnel is down. This can be because the tunnel administrative setting is down , or the tunnel can't communicate with the appliance at the other end.
Down – In progress	The tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel.
Down – Misconfigured	The tunnel is down because it's misconfigured. It needs the correct remote Appliance ID and IP address.
Down – Bad VRRP IP	The tunnel is down because the configured VRRP IP for the remote appliance is invalid.
Up – Active	The tunnel is up and active. Traffic destined for this tunnel is forwarded to the remote appliance.
Up – Degraded	<p>The tunnel is up, but it has excess WAN latency. This happens when the tunnel control packet round-trip time (RTT) exceeds the configured thresholds.</p> <p>This status displays when either of two alarms is raised: Tunnel Health is Unhealthy or Tunnel Health is Marginal.</p> <p>For more information about specific alarm text, see “Types of Alarms” on page 403.</p>
UNKNOWN	The tunnel status is unknown. This can be because the appliance is unable to retrieve the current tunnel status. Try again later.



For more information, see [Chapter 5, “Creating Tunnels.”](#)

Banners

The banner is common to all of the Appliance Manager pages:



Menu Structure

This section lists each main menu with its submenus, and describes its contents.

Configuration Menus

These menus are for configurations related to host settings, networking, applications, and deployments.

Table 3-1 Configuration Menus

Menu Name	Description
System	<p>Configure basic system parameters.</p> <p>Identifies the appliance's hostname, ID, location, and contact person. Here you specify the following:</p> <ul style="list-style-type: none"> whether or not to encrypt data on the appliance hard drives whether or not to put the appliance in System Bypass whether the deployment is <i>in-line</i> (Bridge Mode) or <i>out-of-path</i> (Router Mode), along with IP addresses for the appliance(s) and next-hop(s) if in Bridge mode, whether to propagate the link down if using a 4-port appliance, whether or not to use etherchannel bonding the maximum WAN bandwidth. <p>For more information, see Chapter 4, "Configuring Host Settings."</p>
Date/Time	<p>Configure date and time.</p> <p>Specifies time zone, and whether time is set manually or via one of the user-specified NTP servers.</p> <p>For more information, see Chapter 4, "Configuring Host Settings."</p>

Table 3-1 Configuration Menus (Continued)

Menu Name	Description
Tunnels	<p><i>Configure tunnels.</i></p> <p>Specifies all tunnels, their states and endpoints, their minimum and maximum bandwidths (configured and current), whether or not IPSec is enabled, and how long the tunnel has been up.</p> <p>When modifying a tunnel, you can also access and change the default values for MTU, mode (gre / udp), FEC, reorder wait, and traffic classes.</p> <p><i>For more information, see Chapter 5, "Creating Tunnels."</i></p>
Pass-through	<p><i>Configure pass-through traffic.</i></p> <p>Specifies the minimum and maximum bandwidth for the pass-through traffic. Includes a table for managing bandwidth and queuing by traffic class.</p> <p><i>For more information, see Chapter 8, "Bandwidth Management & QoS Policy."</i></p>
Access Lists	<p><i>Configure access control lists.</i></p> <p>Each ACL consists of one or more rules. Rules specify MATCH criteria that you can associate with Route, QoS, and Optimization policies [maps]. ACLs are convenient when you want to name a set of MATCH criteria and/or reuse them in multiple rules or maps.</p> <p><i>For more information, see Chapter 6, "Theory of Operations."</i></p>
Route Policy	<p><i>Configure route map(s) and specify one map to be the active route policy.</i></p> <p>Route maps associate MATCH criteria, or an existing ACL, with the following SET actions:</p> <ul style="list-style-type: none"> • tunnel • tunnel down action [what to do when the associated tunnel goes down] <p><i>For more information, see Chapter 7, "Route Policy."</i></p>
QoS Policy	<p><i>Configure QoS map(s) and specify one map to be the active QoS policy.</i></p> <p>QoS maps associate MATCH criteria, or an existing ACL, with SET actions. QoS actions tell the appliance how to handle values found in the LAN and WAN ToS (Type of Service) fields.</p> <p><i>For more information, see Chapter 8, "Bandwidth Management & QoS Policy."</i></p>
Optimization Policy	<p><i>Configure optimization map(s) and specify one map to be the active optimization policy.</i></p> <p>Optimization maps associate MATCH criteria, or an existing ACL, with any combination of the following SET actions:</p> <ul style="list-style-type: none"> • Network Memory • payload compression • TCP acceleration • CIFS acceleration <p><i>For more information, see Chapter 9, "Optimization Policy."</i></p>

Table 3-1 Configuration Menus (Continued)

Menu Name	Description
Networking	<p><i>Configure various networking parameters.</i></p> <p>Interfaces: Identifies the specifics of the two management (mgmt0, mgmt1) and two network (lan0, wan0) interfaces. Allows the user to choose whether to statically set IP addresses or to use DHCP.</p> <p>A 4-port appliance includes two additional interfaces — lan1 and wan1.</p> <p>VLAN: When the appliance is in Bridge mode, specifies the configuration for VLANs. This enables tagging of packets and use of VLANs as discriminators in policy and/or ACL MATCH criteria.</p> <p>IP Routes: Specifies or provides information about the following:</p> <ul style="list-style-type: none"> Static routes and the default gateway control how management traffic is routed out of the appliance. Specifies next-hop address(es) for LAN-side networks that are not directly connected to a Bridge-mode (in-line) appliance. <p>DNS: Specifies up to three Domain Name Servers, along with any user-specified domains.</p> <p>VRRP: For high availability, allows you to configure traffic redirection, with or without redundancy, by configuring VRRP (Virtual Router Redundancy Protocol) on a common virtual interface.</p> <p>WCCP: For high availability, allows you to configure traffic redirection, with or without redundancy, by configuring WCCP (Web Cache Communication Protocol).</p> <p>Flow Redirection: Ensures that you don't have asymmetrical TCP flows between peer appliances.</p> <p>NetFlow: For collecting and sending traffic flow data to a centralized NetFlow collector.</p> <p><i>For more information, see Chapter 4, "Configuring Host Settings."</i></p> <p><i>For detailed deployment procedures, see the Silver Peak NX Series Appliances Network Deployment Guide.</i></p>
Application	<p><i>Configure custom applications or application groups.</i></p> <p>Built-in: Lists the standard applications included, by name and port number(s).</p> <p>User Defined: For configuring custom applications, based on any combination of port, port ranges, IP address, DSCP, or VLAN.</p> <p>Groups: Lets you create and name sets of applications. You can use these in the Application field of any ACL's or policy's MATCH criteria.</p> <p><i>For more information, see Chapter 6, "Theory of Operations."</i></p>

Reporting Menus

Each of these displays historical statistics for traffic over a selectable time interval. These are useful for evaluating ongoing performance and optimization results.

Table 3-2 Reporting Menus

Menu Name	Description
Applications	<p><i>View historical application statistics for tunnel, pass-through, or all traffic.</i></p> <p>Provides pie charts and/or a table to display data over a selectable time interval.</p> <p><i>For more information, see “Viewing Application Historical Statistics” on page 248.</i></p>
Reduction	<p><i>View the reduction of (data) bytes afforded by the Silver Peak appliance.</i></p> <p>For all optimized traffic or on an individual tunnel basis, provides charts and/or a table to display data over a selectable time interval.</p> <p><i>For more information, see “Viewing Reduction Statistics” on page 251.</i></p>
Bandwidth	<p><i>View the bandwidth improvements afforded by the Silver Peak appliance.</i></p> <p>For all optimized traffic or on an individual tunnel basis, provides charts and/or a table to display data over a selectable time interval.</p> <p><i>For more information, see “Viewing Bandwidth Statistics” on page 253.</i></p>
Flow Counts	<p><i>View flow counts.</i></p> <p>For all optimized traffic or on an individual tunnel basis, provides data over a selectable time interval, for the various TCP and non-TCP flows.</p> <p><i>For more information, see “Viewing Flow Counts” on page 255.</i></p>
Latency	<p><i>View latency data.</i></p> <p>On a per-tunnel basis, provides data about round-trip latency (minimum, maximum, and average).</p> <p><i>For more information, see “Viewing Latency” on page 257.</i></p>
Network Integrity	<p><i>View packet loss and out-of-order packets.</i></p> <p>On a per-tunnel basis, displays</p> <ul style="list-style-type: none"> • packet loss (minimum, maximum, and average) before /after Forward Error Correction • out-of-order packets (minimum, maximum, and average) before /after Packet Order Correction <p><i>For more information, see “Viewing Network Integrity” on page 258.</i></p>
All Reports	<p><i>Based on the filters chosen, displays all Reports for the traffic, direction, and time interval selected.</i></p> <p>This includes reports for applications, reduction, bandwidth, flows, latency, loss, and out-of-order packets.</p> <p><i>For more information, see “Viewing a Summary of All Historical Reports” on page 261.</i></p>

Monitoring Menus

These reports are realtime in nature and thereby more suitable for use in troubleshooting. Each graph provides for access to the raw data via the **Table View** button.

Table 3-3 Monitoring Menus

Menu Name	Description
Applications	<p><i>Monitor real-time application statistics for individual tunnel, pass-through shaped, or pass-through unshaped traffic.</i></p> <p>Provides the choice of a pie chart or a table to display application statistics for the WAN or LAN side. The pie chart displays the percentage of traffic for the top 10 applications. The table expands the data to include all available applications.</p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see "Viewing Application Realtime Statistics" on page 270.</i></p>
Current Flows	<p><i>Monitor current flows, based on multiple selectable criteria.</i></p> <p>Retrieves a listing of up to 100 realtime connections based on selectable filter criteria. These include end points, ports, route, and application.</p> <p><i>For more information, see "Viewing Current Flows" on page 273.</i></p>
QoS	<p><i>Monitor tunnel-specific QoS statistics.</i></p> <p>Selectable for all tunnels or for individual tunnels, specifies the total number of bytes and packets transmitted and received, based on traffic class and/or WAN QoS [DSCP] marking. Also provides granular information about types of packets, as in dropped, invalid, duplicate, etc.</p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see "Viewing Tunnel QoS Statistics" on page 286.</i></p>
Tunnels	<p><i>Monitor tunnel real-time statistics.</i></p> <p>LAN/WAN Statistics: Specifies the total number of bytes and packets transmitted and received, as well as invalid, lost, and duplicate packets.</p> <p>Flows/Latency/Packet Correction Statistics: Provides data about accelerated and non-accelerated TCP flow, round-trip latency (minimum, maximum, and average), packet loss before and after applying Forward Error Correction, and out-of-order packets before and after Packet Order Correction.</p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see "Viewing Tunnel Realtime Statistics" on page 288.</i></p>
Reduction	<p><i>Monitor the reduction of (data) bytes.</i></p> <p>Selectable for optimized traffic or for individual tunnels, displays the total number of bytes (and packets) received, processed, and transmitted by a tunnel in the outbound (LAN-to-WAN) and inbound (WAN-to-LAN) direction.</p> <p><i>For more information, see "Viewing Reduction Statistics" on page 293.</i></p>
Bandwidth	<p><i>Monitor bandwidth usage.</i></p> <p>Selectable for all tunnels or for individual tunnels, compares the traffic's bandwidth usage on the LAN and WAN.</p> <p><i>For more information, see "Viewing Bandwidth Statistics" on page 295.</i></p>
Flow Counts	<p><i>Monitor flow counts.</i></p> <p>For all optimized traffic or on an individual tunnel basis, provides realtime data for the various TCP and non-TCP flows.</p> <p><i>For more information, see "Viewing Flow Counts" on page 255.</i></p>

Table 3-3 Monitoring Menus (Continued)

Menu Name	Description
Latency	<p><i>View latency data.</i></p> <p>On a per-tunnel basis, provides data about round-trip latency (minimum, maximum, and average).</p> <p><i>For more information, see “Viewing Latency Statistics” on page 299.</i></p>
Network Integrity	<p><i>Monitor packet loss and out-of-order packets.</i></p> <p>On a per-tunnel basis, displays</p> <ul style="list-style-type: none"> • packet loss (minimum, maximum, and average) before /after Forward Error Correction • out-of-order packets (minimum, maximum, and average) before /after Packet Order Correction <p><i>For more information, see “Viewing Network Integrity Statistics” on page 300.</i></p>
Flow Redirection	<p><i>Monitor how many flows were redirected to maintain TCP symmetry.</i></p> <p>Between an appliance and its peer (if one exists), displays how many flows, packets, and bytes were redirected to ensure that flows are not asymmetrical.</p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see “Viewing Flow Redirection Statistics” on page 302.</i></p>
NetFlow	<p><i>Monitor the number of records sent to the NetFlow collector(s).</i></p> <p>Displays the number of control packets sent and received between this appliance and the other peers in the cluster. Also displays how much traffic was redirected to and from the other peers.</p> <p><i>For more information, see “Viewing NetFlow Statistics” on page 304.</i></p>
Interfaces	<p><i>Monitor network LAN/WAN interface statistics.</i></p> <p>For the physical lan0, wan0, mgmt0, and mgmt1 interfaces, specifies the number of bytes, as well as the number of packets discarded, errors encountered, packets overrun, frames sent, carriers used, and multicast packets.</p> <p>For a 4-port appliance, data also accrues for the lan1 and wan1 interfaces.</p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see “Viewing Interface Statistics” on page 305.</i></p>
Bridges	<p><i>Monitor the interfaces used and link health in bridge mode.</i></p> <p>In bridge mode deployments, this screen provides details that help explain the amount of traffic and state of the interfaces in a bridged topology.</p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see “Viewing Bridge Mode Statistics” on page 307.</i></p>
IP Routes	<p><i>Monitor next-hop reachability.</i></p> <p>You can view the statistics accrued since the last reboot or since resetting the counter to zero.</p> <p><i>For more information, see “Viewing IP Routes” on page 308.</i></p>

Administration Menus

These menus are related to managing various appliance administrative functions.

Table 3-4 Administration Menus

Menu Name	Description
Logging	<p><i>Configure logging thresholds and view event and alarm logs.</i></p> <p>Log Settings: For configuring system logging thresholds. Allows you to set parameters related to the local log, including specifying a minimum severity level logged and how often to start a new log file, based on frequency or size. Also, you can designate a remote log server with its own severity threshold.</p> <p>Event Log Viewer: Displays timestamped messages for all system-level activity.</p> <p>Alarm Log Viewer: Displays timestamped messages for all alarm activity.</p> <p><i>For more information, see the following:</i></p> <ul style="list-style-type: none"> • “Configuring Log Settings” on page 310 • “Understanding the Events Log” on page 315 • “Viewing a Log of All Alarms” on page 316.
Debug Files	<p><i>Manage file system.</i></p> <p>Provides user data disk information, and lets you manage the following collections of generated files:</p> <ul style="list-style-type: none"> • Log • Debug Dump • Snapshot • TCP Dump Result • Tech Support <p><i>For more information, see “Managing Debug Files” on page 318.</i></p>
Pre-position	<p><i>Enable an appliance to place received FTP content into Network Memory.</i></p> <p>Allows you to pre-position—or push—files from a file server (FTP client) to the appliance (FTP server). The end result is that the file data is prepopulated into Network Memory.</p> <p><i>For more information, see “Pre-Positioning Data for Enhanced Acceleration Benefits” on page 325.</i></p>
SNMP	<p><i>Configure SNMP trap destinations.</i></p> <p>You can specify which hosts receive SNMP traps.</p> <p><i>For more information, see “Configuring SNMP” on page 327.</i></p>
User Management	<p><i>Manage users and their access.</i></p> <p>Users: Lists all local users — enabled or not. Lets you assign or change user names, groups, and passwords.</p> <p>Banners: For configuring a Login Message and/or Message Of The Day.</p> <p>Authentication: For configuring authentication methods and authorization information.</p> <p>RADIUS: For configuring RADIUS settings and servers.</p> <p>TACACS+: For configuring TACACS+ settings and servers.</p> <p><i>For more information, see “Managing User Accounts” on page 332.</i></p>
Web	<p><i>Configure web server access settings.</i></p> <p>For configuring web protocol settings (http, https), and web user settings (inactivity time-out, maximum number of user sessions).</p> <p><i>For more information, see “Configuring Settings for Web Protocols and Web Users” on page 355.</i></p>

Table 3-4 Administration Menus (Continued)

Menu Name	Description
Initial Config Wizard	<i>Provides access to the Initial Config Wizard.</i> <i>For more information, see “Initial Configuration Wizard” on page 356.</i>
Support	<i>Provides information that Technical Support requires when you contact them, along with the contact info.</i> <i>For more information, see “Support” on page 358.</i>

Maintenance Menus

These menus are related to maintaining the appliance software, databases, disks, Network Memory, and system restarts.

Table 3-5 Administration Menus

Menu Name	Description
System Information	<i>Lists appliance-specific system information.</i> Includes hostname, appliance ID, model, system status, uptime, date/time, software version, serial number, mode, appliance IP address, and whether or not disk encryption is enabled. <i>For more information, see “Viewing System Information” on page 360.</i>
Software Upgrade	<i>Upgrade, install, and manage software images.</i> For loading software images into the two partitions by installing them from a local hard disk, a URL, an SCP (Secure Copy) server, or an FTP server. Also lets you select which partition will be active at the next reboot of the appliance. <i>For more information, see “Upgrading the Appliance Manager Software” on page 361.</i>
Configuration Management	<i>Manage configuration files.</i> Allows you to save or load an appliance configuration file, view a saved configuration, revert to a saved file, activate a file, rename a file, or delete a file. <i>For more information, see “Managing the Appliance Configuration File” on page 372.</i>
Disk Management	<i>Manage disk.</i> Provides read-only RAID information, and provides a table that lists each disk, its status, size, and serial number. <i>For more information, see “Replacing a Hard Disk Drive” on page 410.</i>
Network Connectivity	<i>Test network connectivity.</i> Provides access to three tests — ping , tracert , and tcpdump . <i>For more information, see “Testing Network Connectivity” on page 386.</i>
Erase Network Memory	<i>Clears Network Memory without rebooting the appliance.</i> <i>For more information, see “Erasing Network Memory” on page 399.</i>
Restart System	<i>Restart system with options.</i> Allows you to select from the following: <ul style="list-style-type: none"> • Reboot the appliance • Reboot the appliance and erase Network Memory™ • Shutdown the appliance until the power is turned on again <i>For more information, see “Restarting the Appliance” on page 400.</i>

Alarm Menus

Table 3-6 Alarm Menus

Menu Name	Description
[Current] Alarms	<p>View <i>current alarms</i>.</p> <p>Lists all critical, major, minor, and warning alarms.</p> <p>For more information, see Chapter 16, "Monitoring Alarms."</p>

Managing Tabular Data

By default, a table sorts on the first column. Additionally, the Appliance Manager allows you to select any column for sort by clicking on its heading. The selected column's heading displays an arrow to the right.

The arrowhead shows where the highest values are. In this case, they're at the bottom of the table. To reverse the order, click the arrow.

	IP1	Port1	IP2	Port2 ⬇	Tx Action	Application	L/W Ratio (X)	W/L Ratio (X)	CIFS SMB	Protocol	Up Time	Details
<input type="checkbox"/>	172.10.16.85	39683	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	17m 18s	
<input type="checkbox"/>	172.10.16.85	39681	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	17m 24s	
<input type="checkbox"/>	172.10.16.85	39663	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	19m 16s	
<input type="checkbox"/>	172.10.16.85	39675	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	18m 4s	
<input type="checkbox"/>	172.10.16.85	39659	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	19m 44s	
<input type="checkbox"/>	172.10.16.85	39685	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	16m 51s	
<input type="checkbox"/>	172.10.16.85	39661	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	19m 39s	
<input type="checkbox"/>	172.10.16.85	39679	172.27.100.4	22	none	ssh	N/A	0.5	no	tcp	17m 30s	
<input type="checkbox"/>	172.10.15.86	123	172.27.11.156	123	dc1_2_miami	ntp	0.8	0.9	no	udp	1d 3h 30m 18s	
<input type="checkbox"/>	172.10.15.86	2121	172.27.11.156	139	dc1_2_phoenix	cifs_smb	0.6	0.2	no	tcp [accel]	4m 38s	
<input type="checkbox"/>	172.10.15.86	2108	172.27.11.156	139	dc1_2_phoenix	cifs_smb	0.7	0.2	no	tcp [accel]	19m 38s	

Netmask Notations

When designating netmasks, the Appliance Manager uses the **slash notation** (/) for netmasks or bitmasks.

Mode

☒ Router
 ☐ Bridge

Appliance IP Address / Netmask:

10.10.10.10 / 24

Next-hop IP Address:

172.30.2.3

/ notation, also known as the "slash" notation

Displays **byte notation** as read-only. When you change the / notation, the read-only byte notation updates.

The following table shows the relationship between the / (“slash”) notation, the byte notation, and the corresponding binary numbers (with a dot every eight digits) for the 32-bit addresses. Also included is a count of how many Class A/B/C networks the larger networks encompass.

/ Notation	Binary	Byte Notation	# Class
/0	00000000.00000000.00000000.00000000	0.0.0.0	256 A
/1	10000000.00000000.00000000.00000000	128.0.0.0	128 A
/2	11000000.00000000.00000000.00000000	192.0.0.0	64 A
/3	11100000.00000000.00000000.00000000	224.0.0.0	32 A
/4	11110000.00000000.00000000.00000000	240.0.0.0	16 A
/5	11111000.00000000.00000000.00000000	248.0.0.0	8 A
/6	11111100.00000000.00000000.00000000	252.0.0.0	4 A
/7	11111110.00000000.00000000.00000000	254.0.0.0	2 A
/8	11111111.00000000.00000000.00000000	255.0.0.0	1 A
/9	11111111.10000000.00000000.00000000	255.128.0.0	128 B
/10	11111111.11000000.00000000.00000000	255.192.0.0	64 B
/11	11111111.11100000.00000000.00000000	255.224.0.0	32 B
/12	11111111.11110000.00000000.00000000	255.240.0.0	16 B
/13	11111111.11111000.00000000.00000000	255.248.0.0	8 B
/14	11111111.11111100.00000000.00000000	255.252.0.0	4 B
/15	11111111.11111110.00000000.00000000	255.254.0.0	2 B
/16	11111111.11111111.00000000.00000000	255.255.0.0	1 B
/17	11111111.11111111.10000000.00000000	255.255.128.0	128 C
/18	11111111.11111111.11000000.00000000	255.255.192.0	64 C
/19	11111111.11111111.11100000.00000000	255.255.224.0	32 C
/20	11111111.11111111.11110000.00000000	255.255.240.0	16 C
/21	11111111.11111111.11111000.00000000	255.255.248.0	8 C
/22	11111111.11111111.11111100.00000000	255.255.252.0	4 C
/23	11111111.11111111.11111110.00000000	255.255.254.0	2 C
/24	11111111.11111111.11111111.00000000	255.255.255.0	1 C
/25	11111111.11111111.11111111.10000000	255.255.255.128	
/26	11111111.11111111.11111111.11000000	255.255.255.192	
/27	11111111.11111111.11111111.11100000	255.255.255.224	
/28	11111111.11111111.11111111.11110000	255.255.255.240	
/29	11111111.11111111.11111111.11111000	255.255.255.248	
/30	11111111.11111111.11111111.11111100	255.255.255.252	
/31	11111111.11111111.11111111.11111110	255.255.255.254	
/32	11111111.11111111.11111111.11111111	255.255.255.255	

Date and Time Conventions

The Appliance Manager uses these conventions across all application modules:

- If the screen shows data related to today's date, then only the time is shown.
- If the data is from a past or future date, then both the time and date are included.
- The timestamp is based on the 24-hour (military) clock.

Times shown are based on the local time zone.

Secure Access Methods

By default, the Silver Peak appliance ships preconfigured for web access over **http** with Secure Sockets Layer (SSL) encryption. In other words, the URL begins with **https://**. To change the configuration, see the **Administration - Web** page.

Guidelines for Creating Passwords

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

Object Names

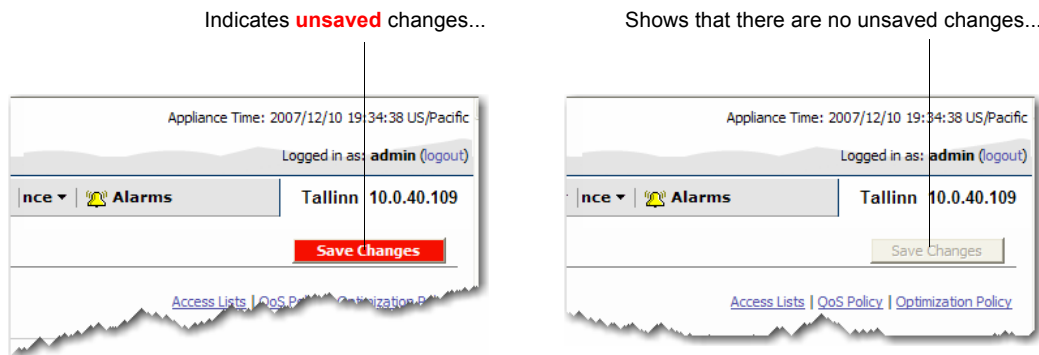
When you create a name for an object, such as a tunnel, access control list, or a route map, you can use one of the following types of characters:

- alphabetical (upper or lower case)
- numerical
- dash (-)
- underscore (_)
- dot (.)

Saving Your Configuration

As you **Apply** page settings, the Appliance Manager applies the values to the running configuration file. As long as the **Save Changes** button — located in the upper right part of the page — has a **red** background, you have not permanently saved the edits.

After you click **Save Changes**, the button reverts to a light background until you make more changes.




You can keep applying your changes as you work across multiple pages and then save the running configuration when you're done.

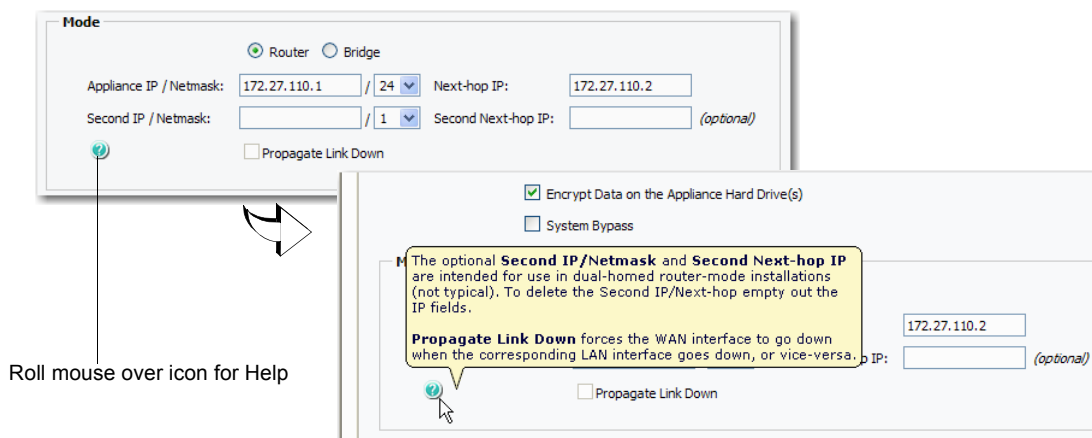
There are a few items, however, that are not part of, or don't affect, the configuration file. These include:

- log files (alarms, events, troubleshooting) and their display
- software partitions (and, hence, upgrade or install operations)
- system restarts

So, for those, you won't see a reminder to save the configuration, either on screen or in this document. But, when in doubt, save. It can't hurt.

Definition Help

When you see the  icon near a term or field, roll the mouse over it for definition Help.





Configuring Host Settings

This chapter describes how to configure or modify appliance system parameters, including the WAN bandwidth available to the appliance. Additionally, it describes how to set the date and time, add DNS servers, work with the routing tables, modify network interface parameters, and configure export to NetFlow collectors.

In This Chapter

- **Overview** See page 94.
- **Configuring Appliance Identity and Max System Bandwidth** See page 94.
- **Selecting a System Deployment** See page 96.
- **Configuring Network Parameters** See page 103.
- **Configuring Flow Exports for NetFlow** See page 114.

Overview

When you first installed the appliance and logged in via the browser, you were presented with the Initial Configuration Wizard. What you configured depends on whether you selected Bridge mode (for *in-line* deployment) or Router mode (for *out-of-path* deployment). Then, typically, you move on to using the Appliance Manager.

The basic parameters that you can and/or need to configure now include:

- **Configuring Appliance Identity and Max System Bandwidth** See page 94.
- **Selecting a System Deployment** See page 96.
- **Configuring Network Parameters** See page 103.
- **Configuring Flow Exports for NetFlow** See page 114.

Configuring Appliance Identity and Max System Bandwidth

When you access the **Configuration - System** page, you can do all of the following:

- Add contact information, such as the appliance owner's name, address, phone number, and/or e-mail address
- Add a user-friendly note about the appliance's location, whether it be a street address or a company's department
- Choose not to encrypt data on the appliance hard drive(s), contrary to the default selection
- Change modes: **Bridge**, for in-line deployment, or **Router**, for out-of-path deployment. After changing the mode, you must reboot the appliance for the change to take effect.
- Choose to propagate a down link, or not. **Propagate Link Down** forces the WAN interface to go down when the corresponding LAN interface goes down, or vice versa. By default, this option is enabled.
- View and/or modify the total bandwidth configured for the WAN interface
- Select **System Bypass** to put the appliance into hardware bypass mode.
- Configure etherchannel bonding, to bond pairs of Ethernet ports (**lan0 + lan1, wan0 + wan1**) into a single port with one IP address. For more information, see "*Configuring Gigabit Etherchannel Bonding*" on page 99.

The **Hostname** and **Appliance ID** fields display what you entered when you first connected to the appliance and were taken to the initial configuration wizard. The same is true for the mode (**Router** or **Bridge**) you selected (lower down this page).

You can modify those entries here, except for the **Hostname**. The only way to change that within the GUI is to go to the **Administration** menu and select **Initial Config Wizard**. The appliance must reboot for the change to take effect.

Encrypting hard disk data is the appliance default.


When you select **System Bypass**, the appliance mechanically isolates itself from the network, allowing traffic to flow without intervention.

When you deselect **System Bypass**, the appliance again allows traffic to be directed to tunnels for optimization.

For 4-port devices, enables gigabit etherchannel bonding.

Propagate Link Down forces the WAN interface to go down when the corresponding LAN interface goes down, or vice versa. Checked by default.

The **WAN Bandwidth** section refers to the bandwidth at the router's WAN interface.



For guidance, refer back to *"Best Practices for Bandwidth Management"* on page 184.

The default value for **wan0**, shown in the **Max Bandwidth** field, is different for each appliance model. To maintain the best traffic flow, refer to the section, *"Best Practices for Bandwidth Management"* on page 184.

The **NX-9700**, **NX-9610**, and **NX-8700** run at either 1G (by default) or 10G. To run at 10G, you must cable the appropriate interfaces, select its **10 Gigabit** checkbox, and reboot.

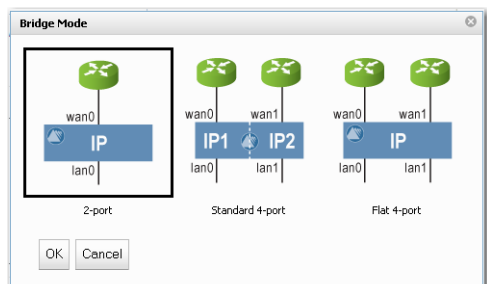
Notice that in the GUI, its interfaces are labeled **twan0** and **tlan0**.

Selecting a System Deployment

Along with a view of the deployment options, this section addresses the following:

- **Sorting Through the Deployment Options** See page 98.
- **Configuring Gigabit Etherchannel Bonding** See page 99.

Changing deployment mode requires a reboot.



BRIDGE MODE

This screen appears when you click **Advanced** after choosing **Bridge**.

Bridge 2-port

- Most common bridge mode deployment
- The appliance is in-line between a single WAN router and a single WAN-side switch.

Bridge Standard 4-port

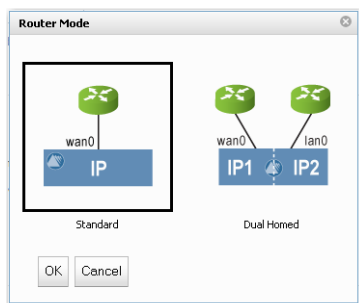
- Most common 4-port bridge configuration
- 2 routers / 2 subnets / 1 appliance
- 2 ISPs or 2 services (MPLS, IPSec VPN, MetroEthernet, etc.) or both

Bridge Flat 4-port

- Less common 4-port bridge configuration
- 2 routers / 1 subnet / 1 appliance
- 2 ISPs or 2 services (MPLS, IPSec VPN, MetroEthernet, etc.) or both


Bridge 4-port Bonding

- Increased throughput on a very high-end appliance and/or interface-level redundancy
- "Super" version of Bridge 2-port
- Uses 2 ports from same WAN router




ROUTER MODE

This screen appears when you click **Advanced** after choosing **Router**.



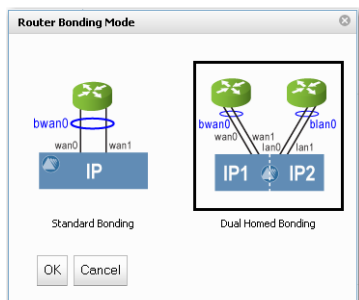
Router Standard

- Most common one-armed configuration
- Single subnet / single router



Router Dual Homed

- For flow redirection on the WAN side
- 2 routers / 2 subnets / 1 appliance
- 2 ISPs or 2 services (MPLS, IPSec VPN, MetroEthernet, etc.) or both
- Can set up 2 instances of either WCCP or Policy-Based Routing




ROUTER BONDING MODE

This screen appears when you click **Advanced** after choosing **Router** and **Bonding**.


Using 4-port appliances for more throughput

[1 Gbps + 1 Gbps = 2 Gbps]



Router Standard Bonding

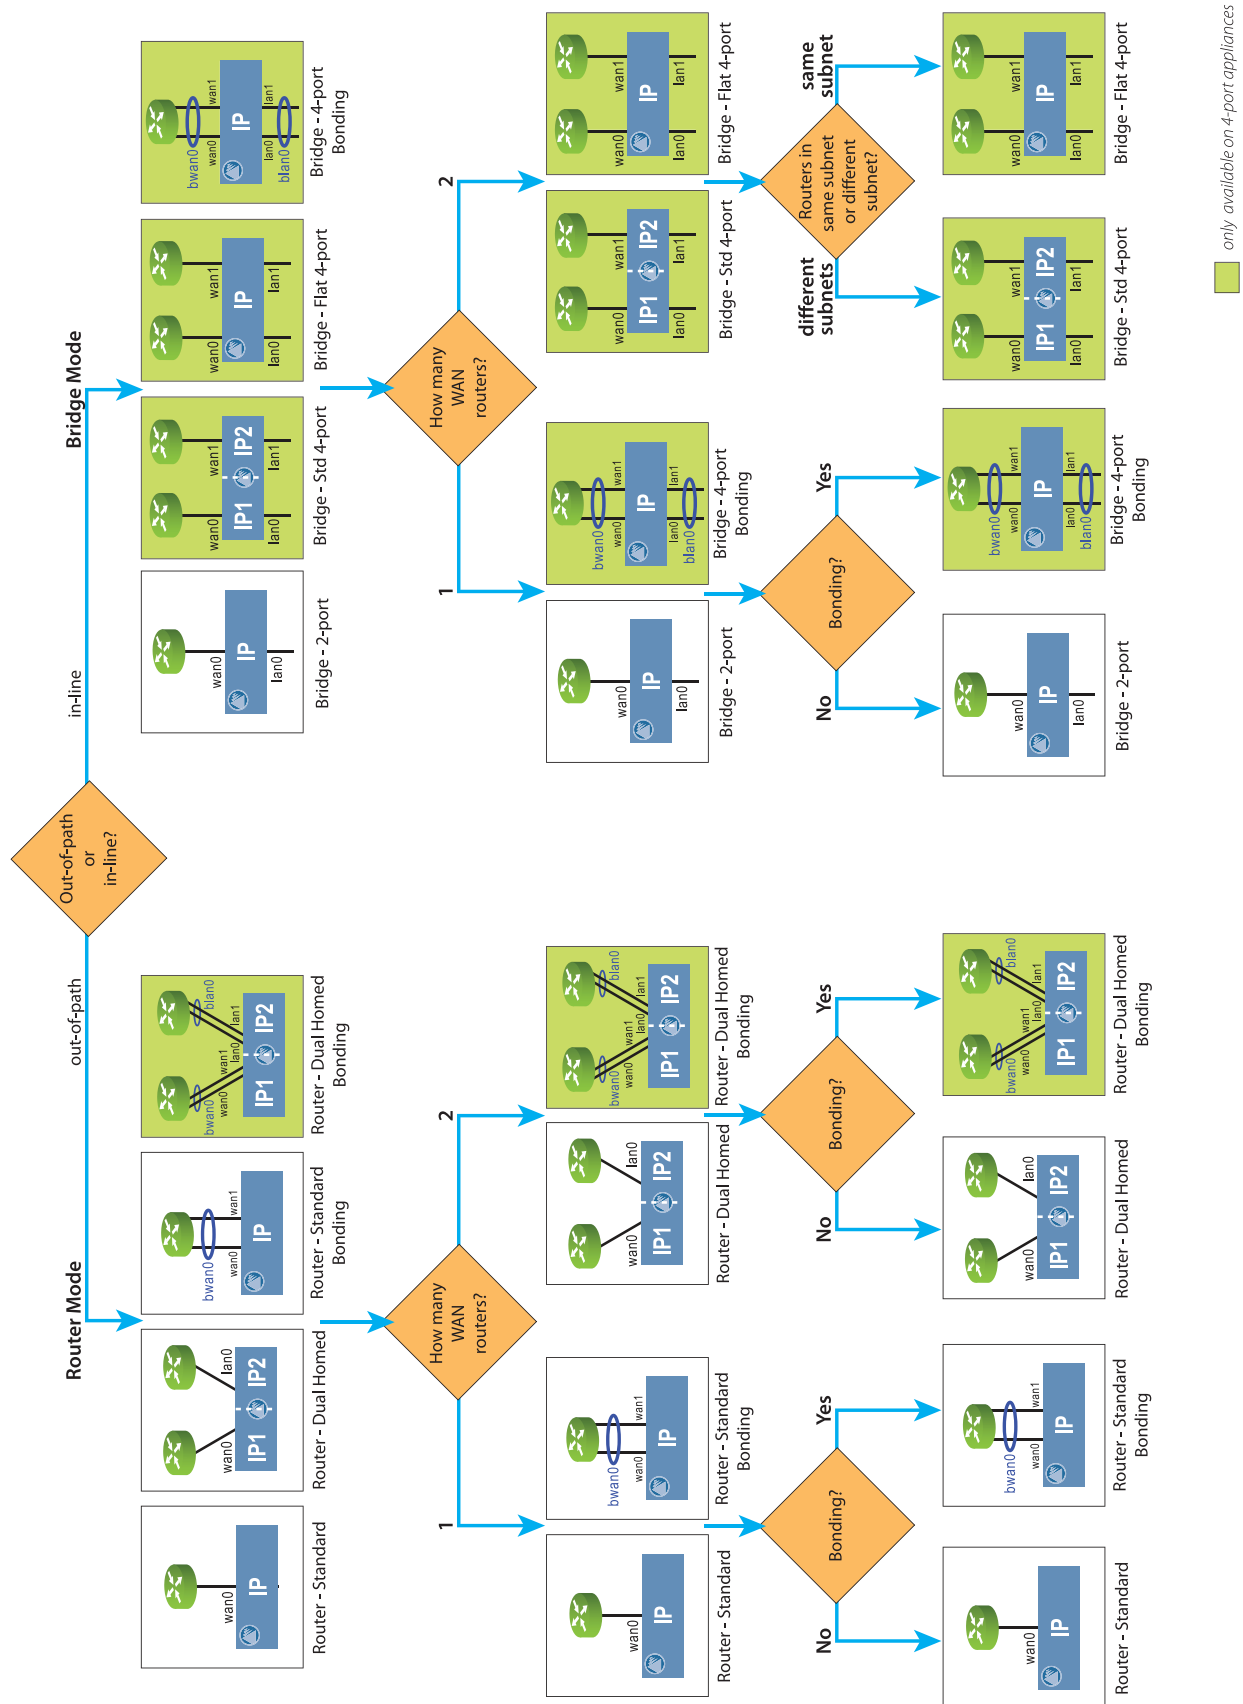
- Increased throughput on a very high-end appliance and/or interface-level redundancy
- “Super” version of Router Standard
- Uses 2 ports at same WAN router



Router Dual Homed Bonding

- Increased throughput on a very high-end appliance and/or interface-level redundancy
- “Super” version of Router Dual Homed
- Uses 2 ports at each WAN router

Sorting Through the Deployment Options



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 NX Series 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 - System** page, and select one of the three available bonding modes:
 - a To select the **Bridge** mode bonding option for **4-port Bonding**, click **Bonding** and **Bridge**.
 - b To view and select from both **Router** mode bonding options, click **Bonding**, **Router**, and **Advanced**. Choose from **Standard Bonding** or **Dual Homed Bonding**, and click **OK**.
- 2 Click **Apply**. The **Confirm Dialog** box appears.



- 3 Click **Reboot**. The appliance reboots.
- 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>

```

Setting the Date and Time

You need to set the system date and time by either entering it manually or assigning an NTP server to the appliance. If you decide to use an NTP server, you must first manually set the date and time to be close to accurate, before pointing to NTP.

- 1 From the **Configuration** menu, select **Date/Time**.

- If this button is **red**, click it to **save** your changes.
- If this button is **greyed out**, you have no unsaved changes.

This opens a calendar when clicked.

The screenshot shows the Silver Peak web interface. At the top, the 'Configuration' menu is selected. The 'Date/Time' configuration page is displayed. In the 'Date/Time Setting' section, the 'Manually' radio button is selected. The 'Time Zone' is set to 'US/Pacific'. The 'Date' field is '2009/03/31' with a calendar icon. The 'Time' field is '16:09:33'. Below these are three 'NTP Servers' fields. A 'Save Changes' button is at the top right. A line points from the text 'This opens a calendar when clicked.' to the calendar icon in the date field.

- 2 At the **Time Zone** field, use the drop-down menu to select your appliance's time zone.

♦ To set the date and time manually

- 1 In the **Date/Time Setting** section, click **Manually**.
- 2 From the **Time Zone** field, select the appliance's local time zone.
- 3 In the **Date** field, enter the date in the format **YYYY/MM/DD**, that is Year/Month/Date.
- 4 In the **Time** field, enter the time based on a 24-hour clock, using the format **HH/MM/SS**, that is Hour/Minute/Second.
- 5 Click **Apply**.

♦ **To add an NTP server**

- 1 In the **Date/Time Setting** area, click **NTP Time Synchronization**.

The screenshot shows the Silver Peak Configuration - Date/Time page. The 'Date/Time Setting' section has 'NTP Time Synchronization' selected. The 'Time Zone' is 'US/Pacific'. The 'Date' is '2007/11/08' and the 'Time' is '13:46:01'. The 'NTP Servers' section has three rows for 'Server 1', 'Server 2', and 'Server 3', each with an IP address field and a 'Version' dropdown menu. The 'Apply' and 'Cancel' buttons are at the bottom.

- 2 From the **Time Zone** field, select the appliance's local time zone.
- 3 In the **NTP Servers** section, do the following:

The close-up screenshot shows the 'NTP Servers' section. It has three rows for 'Server 1', 'Server 2', and 'Server 3'. 'Server 1' has the IP address '216.218.254.202' and 'Version' set to '4'. 'Server 2' and 'Server 3' have empty IP address fields and 'Version' set to '4'. The 'Apply' and 'Cancel' buttons are at the bottom.

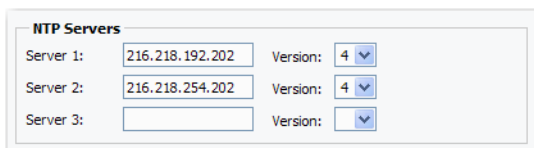
- a Enter the IP address for an NTP server of your choice.
- b From the **Version** field, select **3** or **4**, as appropriate.
- 4 If you want, you can enter two additional NTP servers. The Appliance Manager selects from the servers in the order listed.
 - If **Server 1** is unavailable, it tries **Server 2**.
 - If both **Server 1** and **Server 2** are unavailable, it tries **Server 3**.
- 5 Click **Apply**.
- 6 Click **Save Changes**.

♦ **To edit an NTP server**

- 1 In the **NTP Servers** section, edit the server's IP address and/or select a different version number.
- 2 Click **Apply**.
- 3 Click **Save Changes**.

♦ **To remove an NTP server**

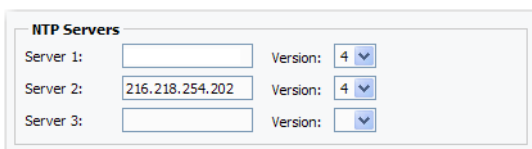
In this example, we remove **Server 1** from the list and observe the results.



The screenshot shows the 'NTP Servers' configuration panel. It contains three rows, each with a 'Server' label, an IP address text box, and a 'Version' dropdown menu. Server 1 has IP '216.218.192.202' and Version '4'. Server 2 has IP '216.218.254.202' and Version '4'. Server 3 has an empty IP box and a dropdown menu.

Server	IP Address	Version
Server 1	216.218.192.202	4
Server 2	216.218.254.202	4
Server 3		

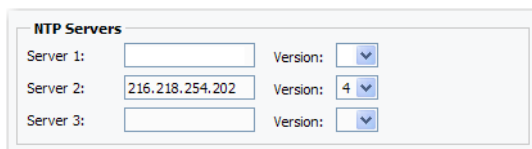
- 1 In the **NTP Servers** area, do the following:
 - a In the **Server** field, cut or delete the IP address you want to remove.



The screenshot shows the 'NTP Servers' configuration panel. Server 1's IP address field is now empty. Server 2's IP is '216.218.254.202' and Version is '4'. Server 3's IP is empty and Version is a dropdown menu.

Server	IP Address	Version
Server 1		4
Server 2	216.218.254.202	4
Server 3		


- b In the **Version** field, select the blank space from the drop-down menu.



The screenshot shows the 'NTP Servers' configuration panel. Server 1's Version dropdown menu is now set to a blank space. Server 2's IP is '216.218.254.202' and Version is '4'. Server 3's IP is empty and Version is a dropdown menu.

Server	IP Address	Version
Server 1		
Server 2	216.218.254.202	4
Server 3		

- 2 Click **Apply**. The remaining server(s) move up the list to keep the sequence unbroken.



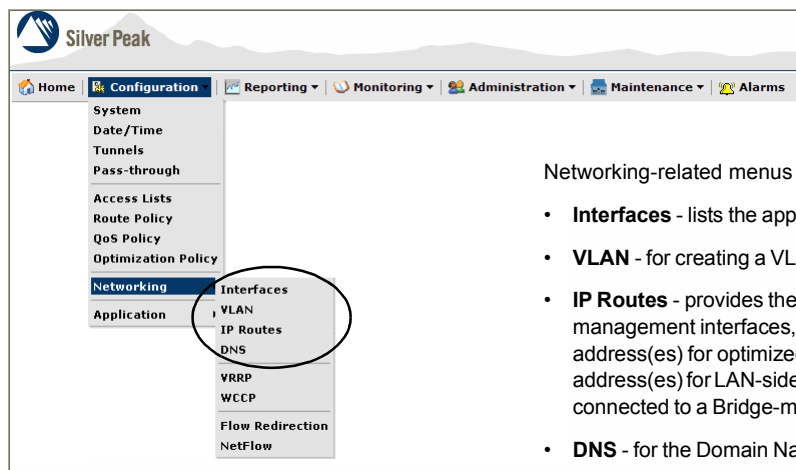
The screenshot shows the 'NTP Servers' configuration panel after the changes have been applied. Server 1 now contains the IP '216.218.254.202' and Version '4'. Server 2's IP is empty and Version is a dropdown menu. Server 3's IP is empty and Version is a dropdown menu.

Server	IP Address	Version
Server 1	216.218.254.202	4
Server 2		
Server 3		

- 3 Click **Save Changes**.

Configuring Network Parameters

This section discusses the networking-related menus found under the **Configuration > Networking** menu.



Networking-related menus include:

- **Interfaces** - lists the appliance's physical interfaces.
- **VLAN** - for creating a VLAN interface
- **IP Routes** - provides the default gateway for the management interfaces, the WAN next-hop address(es) for optimized traffic, and the next-hop address(es) for LAN-side networks that are not directly connected to a Bridge-mode (in-line) appliance.
- **DNS** - for the Domain Name Server (DNS)

These menus are discussed in the following sections:

- **Modifying the Physical Interface Parameters** See page 104.
- **Configuring IP Routes** See page 107.
- **Routing Management Traffic** See page 108.
- **Routing LAN-side Traffic to the Next Hop** See page 111.
- **Adding Domain Name Servers** See page 113.

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

Modifying the Physical Interface Parameters

This section describes how to modify the interface parameters, such as speed, duplex, and MTU (Maximum Transmission Unit).

Here you can choose whether you want to statically set the IP addresses (recommended) or whether you want to use DHCP (Dynamic Host Configuration Protocol). Using DHCP automatically retrieves the IP address from the DHCP server and displays it in the table.



WARNING DHCP 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.

◆ To access the Configuration - Interfaces page

Go to the menu and select **Configuration > Networking > Interfaces**.

At first, the interface values displayed result from the data entered during the initial installation.

This screen shows a 2-port NX appliance in **Bridge** mode, the in-path deployment scenario. Because the appliance "bridges" the **wan0** and **lan0** interfaces with the Appliance IP address, no separate **wan0** and **lan0** IP addresses display. If this were a 4-port appliance, then this would also list the **wan1** and **lan1** interfaces

Status is the interface's actual state at the moment.

Admin denotes the state you specify for the interface.

If DHCP is not configured, the IP addresses are static.

Speed and Duplex default to **Auto**.

Name	Status	Admin	IP Address	Netmask	DHCP	Speed	Duplex	MTU	MAC
wan0	up	up	--	--	no	1000Mb/s (auto)	full (auto)	1500	00:0C:BD:00:80:F8
lan0	up	up	--	--	no	1000Mb/s (auto)	full (auto)	1500	00:0C:BD:00:80:F9
mgmt0	up	up	10.0.42.76	255.255.255.224	yes	1000Mb/s (auto)	full (auto)	1500	00:E0:81:2F:B7:D6
mgmt1	down	up	169.254.0.1	255.255.0.0	no	--	--	1500	00:E0:81:2F:B7:D7

In **Router** mode (for out-of-path deployments), **no** cable connects to the LAN port. Therefore, although the **mgmt0**, **mgmt1**, and **wan0** interfaces have IP addresses, the **lan0** interface does **not**.

Name	Status	Admin	IP Address	Netmask	DHCP	Speed	Duplex	MTU	MAC
wan0	Down	up	172.30.2.34	255.255.255.0	no	--	--	1500	00:0C:BD:00:7F:4A
lan0	Down	down	--	--	no	--	--	1500	00:0C:BD:00:7F:4B
mgmt0	Up	up	10.0.55.90	255.255.252.0	no	100Mb/s (auto)	full (auto)	1500	00:E0:81:2F:B5:98
mgmt1	Down	up	169.254.0.1	255.255.0.0	no	--	--	1500	00:E0:81:2F:B5:99

When you click an interface's **Name**, the page displays the **Configure interface [interface name]** area.

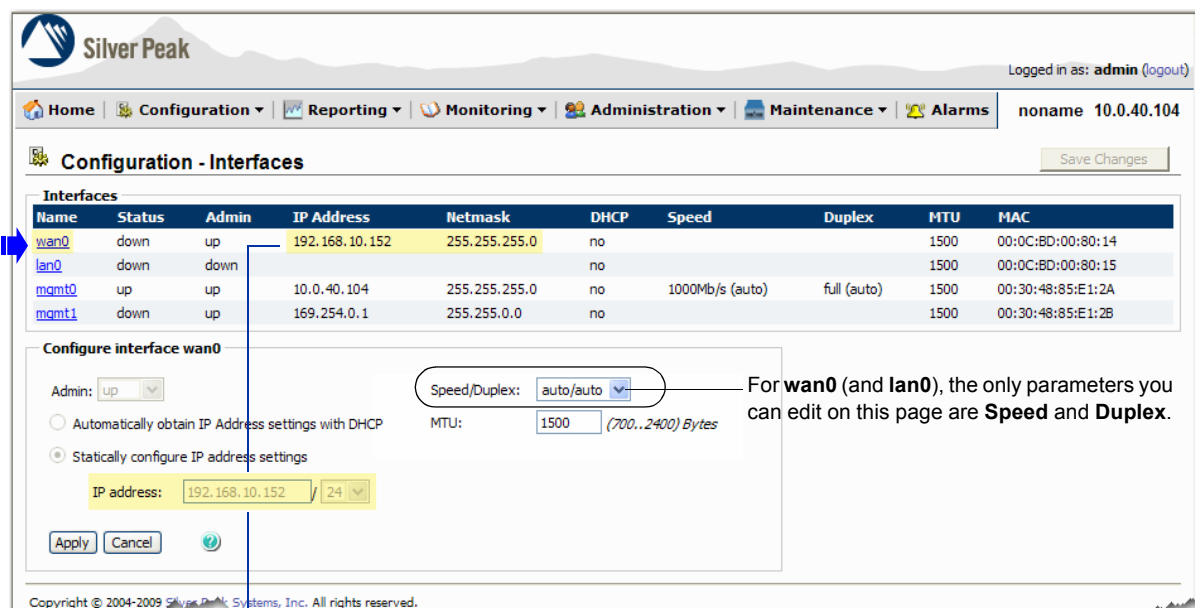
Interface name displays here

The screenshot shows a configuration window titled "Configure interface mgmt0". The title bar has "mgmt0" circled. Inside the window, there are several settings: "Admin" is set to "up", "Speed/Duplex" is set to "auto/auto", "MTU" is set to "1500" with a note "(700..2400) Bytes", and "IP address" is set to "10.0.41.119/27". There are two radio buttons: "Automatically obtain IP Address settings with DHCP" (selected) and "Statically configure IP address settings". At the bottom, there are "Apply", "Cancel", and a status icon.

- If an interface has no assigned IP address, then the field is blank.
- In Bridge (in-line) mode, no IP addresses display for **lan0** and **wan0**.
- In Router (out-of-path) mode, no IP address displays for **lan0**. Also the **wan0** IP address is the same as the **Appliance IP** address.

◆ **To configure or edit an interface**

- 1 In the **Interfaces** table, click any hyperlinked interface **Name**. The **Configuration - Interfaces** area displays. This example displays the edit area for the **wan0** interface.



Configuration - Interfaces

Save Changes

Name	Status	Admin	IP Address	Netmask	DHCP	Speed	Duplex	MTU	MAC
wan0	down	up	192.168.10.152	255.255.255.0	no			1500	00:0C:BD:00:80:14
lan0	down	down			no			1500	00:0C:BD:00:80:15
mgmt0	up	up	10.0.40.104	255.255.255.0	no	1000Mb/s (auto)	full (auto)	1500	00:30:48:85:E1:2A
mgmt1	down	up	169.254.0.1	255.255.0.0	no			1500	00:30:48:85:E1:2B

Configure interface wan0

Admin:

☐ Automatically obtain IP Address settings with DHCP

☒ Statically configure IP address settings

IP address:

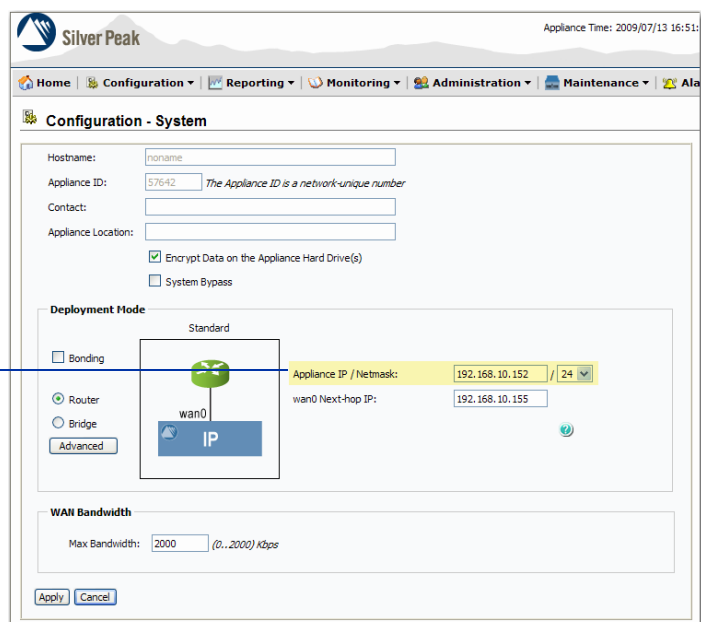
Speed/Duplex: MTU: (700..2400) Bytes

Apply Cancel

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.



Tip To change **wan0**'s IP address (which is the same as the **Appliance IP Address**), you need to access the **Configuration - System** page.



Configuration - System

Hostname:

Appliance ID: The Appliance ID is a network-unique number

Contact:

Appliance Location:

☒ Encrypt Data on the Appliance Hard Drive(s)

☐ System Bypass

Deployment Mode

☐ Bonding

☒ Router

☐ Bridge

Advanced

Standard

Appliance IP / Netmask:

wan0 Next-hop IP:

WAN Bandwidth

Max Bandwidth: (0..2000) Kbps

Apply Cancel

- 2 Edit the interface parameters, as needed.
- 3 Click **Apply**, and click **Save Changes**.

Configuring IP Routes

The **Configuration - IP Routes** page is the locus for information about:

- the default gateways for **management** interfaces
- **WAN next-hop** address(es) for optimized traffic, and
- **next-hop** address(es) for **LAN-side** networks that are not directly connected to a Bridge-mode (in-line) appliance.

Management Route(s) provide the default gateway(s) for the management interfaces.

This is a routing table.

For more information, see "Routing Management Traffic" on page 108.

WAN Next-hop(s) provide next-hop address(es) for optimized traffic.

When two WAN next-hops are configured Active/Active in 4-port bridge mode:

- LAN0 ingress traffic is routed to the first WAN next-hop
- LAN1 Ingress traffic is routed to the second WAN next-hop

When two WAN next-hops are configured Active/Active in dual-homed router mode:

- WAN0 ingress traffic is routed to the first WAN next-hop
- LAN0 Ingress traffic is routed to the second WAN next-hop

The screenshot shows the Silver Peak Configuration - IP Routes page. At the top, there's a navigation bar with links like Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The main content area is divided into three sections: Management, WAN, and LAN. Each section contains a table of IP routes. The Management table lists default gateways for mgmt0 and mgmt1 interfaces. The WAN table shows next-hop IP addresses and their roles. The LAN table shows destination IP netmasks and next-hop IP addresses with metrics. There are buttons for adding, removing, and setting default gateways for each section. A modal window for adding a LAN route is also shown, with fields for destination IP netmask, next-hop IP, and metric.

For edits to WAN next-hops, this returns you to the **Configuration - System** page.

The screenshot shows the Deployment Mode configuration page. It has a section for Deployment Mode with options for Bonding, Router, and Bridge. The Bridge mode is selected. Below this, there are fields for Appliance IP / Netmask, wan0 Next-hop IP, and lan0 Next-hop IP. There is also a checkbox for Propagate Link Down.

LAN Route(s) provide next-hop address(es) for LAN-side networks that are not directly connected to a bridge-mode appliance.

Redundant (backup) LAN routes can be created by specifying another nextphop with a higher metric value.

Examples:

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

For more information, see "Routing LAN-side Traffic to the Next Hop" on page 111.

Routing Management Traffic

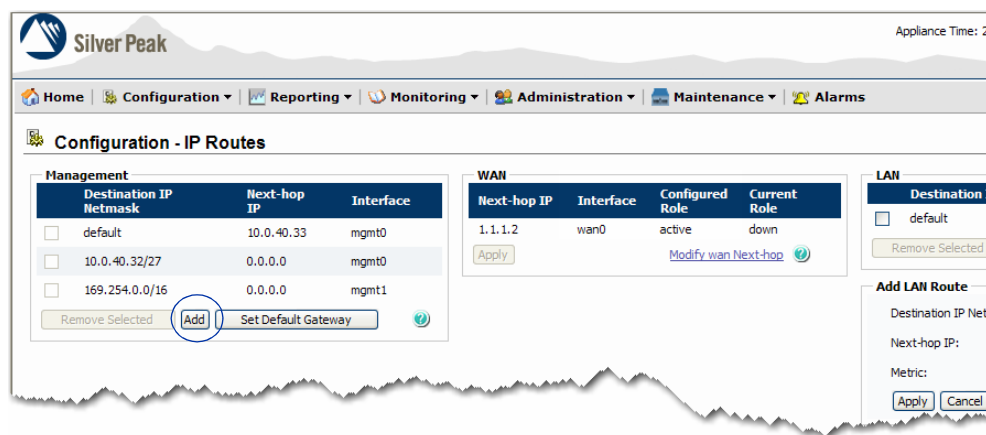
Static routes on the **Configuration - IP Route** page control how *management traffic* is routed out of the appliance.

The **IP Routes** table shows the configured static routes, as well as the system's dynamically created routes. These dynamic routes are automatically added by the appliance, as a result of configuring the appliance's management interfaces.

Dynamic routes cannot be deleted or added by the user.

This section describes the following:

- **To add a new static management route** See page 108.
 - **To configure a new Next-hop IP Address for a management interface** See page 109.
 - **To remove a management route from the routing table** See page 110.
- ♦ **To add a new static management route**
- 1 To add a static route, click **Add**.



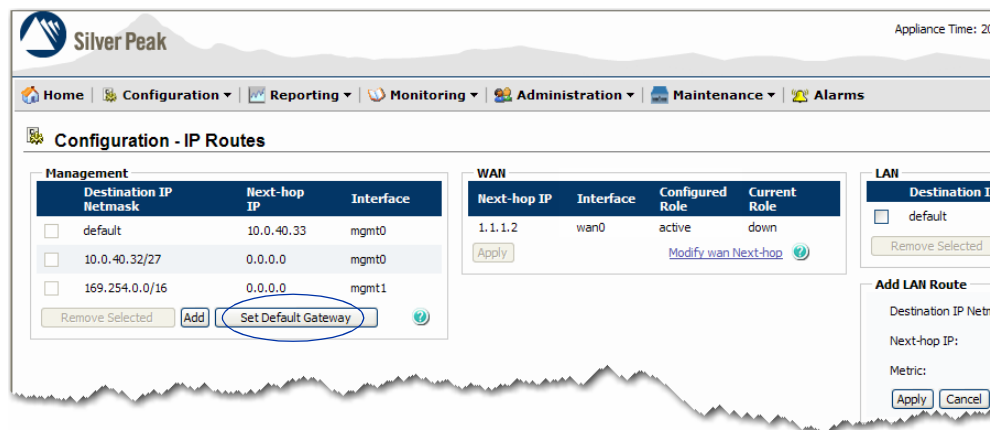
The **Add Static Management Route** area appears.

- 2 In the **Add Static Management Route** area, enter the values for the **Destination IP**, **Netmask**, and the associated **Next-hop IP**. And select either **mgmt0** or **mgmt1**.

- 3 Click **Apply**. The new static route displays in the table.
- 4 Click **Save Changes**.

◆ To configure a new Next-hop IP Address for a management interface

- 1 Click **Set Default Gateway**.



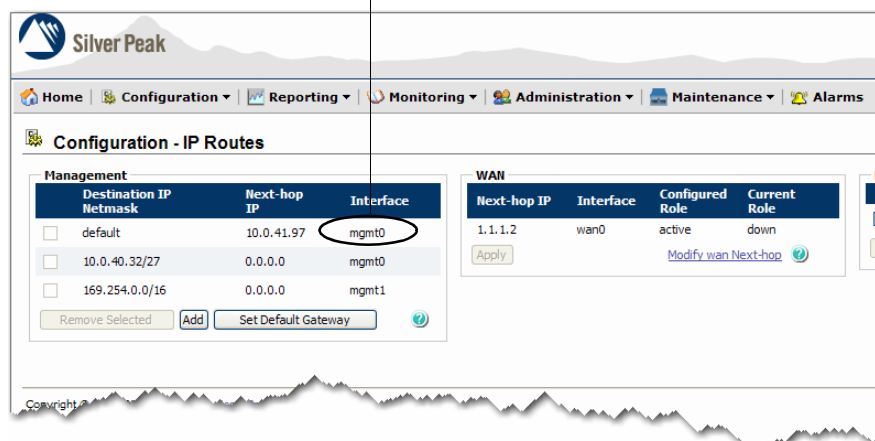
The **Set Default Gateway** area appears.

- 2 In the **Next-hop IP** field, edit the IP address for the next router hop and select the management interface.



- 3 Click **Apply**.

In the routing table, the new Next-hop IP address replaces the old one.



- 4 Click **Save Changes**.

♦ **To remove a management route from the routing table**

- 1** Select the route in the routing table. The Appliance Manager greys out routes that you cannot delete.
- 2** Click **Remove Selected**. The Appliance Manager asks you to confirm this action.
- 3** Click **OK**.
- 4** Click **Save Changes**. The Appliance Manager permanently deletes the route.

Routing LAN-side Traffic to the Next Hop

This section of the **Configuration - IP Routes** page provides next-hop address(es) for LAN-side networks that are not directly connected to a Bridge-mode (in-line) appliance.

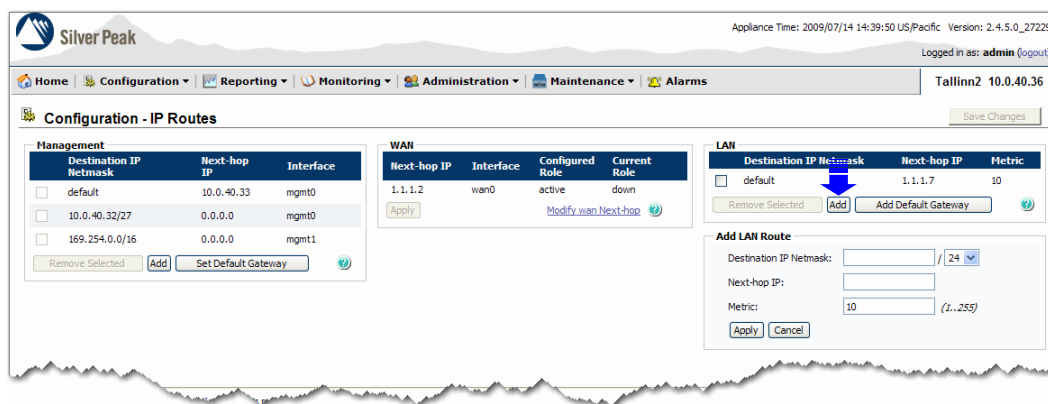
This is only applicable when the appliance is in Bridge mode. That is, it's an in-line deployment.

This section describes the following:

- **To add a LAN-side route** See page 111.
- **To add a default LAN-side gateway** See page 112.
- **To remove a LAN route from the LAN table** See page 112.
- **A note about datapath connectivity** See page 112.

◆ To add a LAN-side route

- 1 In the **LAN** area, click **Add**. The **Add LAN Route** area appears below it.



- 2 In the **Destination IP Netmask** field, enter the traffic's desired destination and subnet mask.
- 3 In the **Next-hop IP** field, enter the IP address of the nearest router to which you want to direct the traffic.
- 4 If you're creating a redundant (backup) LAN route, specify a higher metric value for the backup route.
- 5 Click **Apply**.

◆ **To add a default LAN-side gateway**

- 1 In the **LAN** area, click **Add Default Gateway**. The **Add Default Gateway** area appears below it.

The screenshot shows the Silver Peak Configuration - IP Routes page. The page has a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as admin. The page title is Configuration - IP Routes. There are three main sections: Management, WAN, and LAN. The LAN section is active, showing a table with columns: Destination IP Netmask, Next-hop IP, and Metric. A blue arrow points to the 'Add Default Gateway' button in the LAN section.

Destination IP Netmask	Next-hop IP	Interface	Configured Role	Current Role
default	10.0.40.33	mgmt0		
10.0.40.32/27	0.0.0.0	mgmt0		
169.254.0.0/16	0.0.0.0	mgmt1		

Next-hop IP	Interface	Configured Role	Current Role
1.1.1.2	wan0	active	down

Destination IP Netmask	Next-hop IP	Metric
default	1.1.1.7	10

Buttons: Remove Selected, Add, Add Default Gateway, Apply, Cancel.

- 2 In the **Next-hop IP** field, enter the IP address of the nearest router to which you want to direct the traffic.
- 3 If you're creating a redundant (backup) LAN gateway, specify a higher metric value for the backup.
- 4 Click **Apply**.

◆ **To remove a LAN route from the LAN table**

- 1 Select the route in the routing table. The Appliance Manager greys out routes that you cannot delete.
- 2 Click **Remove Selected**. The Appliance Manager asks you to confirm this action.
- 3 Click **OK**.
- 4 Click **Save Changes**. The Appliance Manager permanently deletes the route.

A note about datapath connectivity

If there is either a **LAN Next-Hop Unreachable** or **WAN Next-Hop Unreachable** alarm after upgrading to Version 2.4.3.1, resolve the alarm(s) immediately by configuring the gateway(s) to respond to ICMP pings from the Silver Peak NX Appliance IP Address.

Adding Domain Name Servers

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 **silver-peak.com**, instead of using the routable IP address.

On the **Configuration - DNS** page, you can configure up to three DNS servers.

Enter the IP Address of the **Primary DNS IP address**, and click **Apply**.
Entering additional ones, in case the first one fails, is optional.

The screenshot shows the 'Configuration - DNS' page in the Silver Peak web interface. The page has a header with the Silver Peak logo, 'Appliance Time: 2009/04/09 11:16:29 US/Pacific', 'Version: 2.4.2.0_25736', and 'Logged in as: admin (logout)'. Below the header is a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The main content area is titled 'Configuration - DNS' and has a 'Save Changes' button in the top right. The 'Name Servers' section contains three input fields: 'Primary DNS IP address' (172.20.40.15), 'Secondary DNS IP address' (172.20.40.10 (Optional)), and 'Tertiary DNS IP address' (Optional). Below these are 'Apply' and 'Cancel' buttons. The 'Domain Names' section has a 'Domain Name' input field, 'Add -->' and 'Remove <--' buttons, and a list box containing 'speak.local'. A line points from the text 'To add a domain name, enter it in the Domain Name field and click Add.' to the 'Domain Name' input field. Another line points from the text 'The domain name then displays in the field to the right.' to the list box.

To add a domain name, enter it in the **Domain Name** field and click **Add**.

The domain name then displays in the field to the right.

Finally, save the changes.

Configuring Flow Exports for NetFlow

On the **Configuration - NetFlow** page, you can configure your appliance to export statistical data to NetFlow collectors.

The screenshot shows the Silver Peak web interface. At the top, there's a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as 'admin'. The main content area is titled 'Configuration - NetFlow'. It contains a 'Netflow Configuration' section with the following fields:

- Flow Exporting Enabled: ☐
- Active Flow Timeout: (1..30) mins
- Traffic Type: ☒ WAN Tx, ☐ WAN Rx, ☐ LAN Tx, ☐ LAN Rx
- Collector 1 IP:
- Collector 1 Port:
- Collector 2 IP:
- Collector 2 Port:

There are 'Apply' and 'Cancel' buttons at the bottom left of the configuration section, and a 'Save Changes' button at the top right. The footer shows 'Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.'

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.

♦ To add a NetFlow collector

- 1 Select **Configuration > Networking > NetFlow**. The **Configuration - NetFlow** page appears.
- 2 To enable flow exporting, select **Flow Exporting Enabled**.
- 3 Accept or change the default value (in minutes) for **Active Flow Timeout**. The range is 1 to 30 minutes.
- 4 From **Traffic Type**, select the interface(s) you want. The default is **WAN Tx**.
- 5 In the **Collector 1 IP** field, enter the IP address of the device to which you're exporting the NetFlow statistics. The default **Collector Port** is **2055**.
If you want to add a second collector, you can.
- 6 Click **Apply**.
- 7 Click **Save Changes**.

♦ To delete a NetFlow collector

- 1 Delete the specific IP address(es) from the **Collector [1 or 2] IP** field.
- 2 Click **Apply**.
- 3 Click **Save Changes**.



Creating Tunnels

This chapter describes how to create and manage tunnels. The appliance only optimizes traffic that the Route Policy directs to a tunnel.

Additionally, it addresses tunnel compatibility mode, which enables two nodes with mismatched software versions to keep the tunnel up and offer basic services.

The discussion of creating tunnels for high availability with VRRP and WCCP is beyond the scope of this document. For those specifics, see the [Silver Peak NX Series Appliances Network Deployment Guide](#).

In This Chapter

- **Overview** See page 116.
- **Creating a Traffic-Carrying Tunnel** See page 117.
- **Editing a Tunnel** See page 123.
- **Deleting a Tunnel** See page 125.
- **Tunnel Compatibility Mode** See page 126.

Overview

Silver Peak Systems appliances are interconnected by tunnels. With this mechanism, Silver Peak uses proprietary techniques to optimize traffic flows and enhance application performance.

When a Silver Peak appliance receives a packet, three separate policies — **Route**, **QoS**, and **Optimization** — use MATCH criteria to delineate flows and apply SET actions.

How Policies Affect Tunnel Traffic

The **Route Policy**'s MATCH criteria and SET actions determine if a flow is directed to a specific tunnel. If it is, then:

- The appliance encapsulates the flow's packets, according to the tunnel configuration. The default is **GRE** without **IPSec**. Other options include:
 - **GRE** with **IPSec**
 - **UDP** without **IPSec**
- The tunnel may be shaped to a specified maximum bandwidth to avoid overrunning downstream bottlenecks.
 - The bandwidth shaping is configured on the **Configuration - Tunnels** page, and the **QoS Policy** assigns the traffic class.
 - The **QoS Policy** honors or changes the DSCP markings to request appropriate per-packet treatment by the network.
- The **Optimization Policy** applies optimization, compression, and acceleration techniques to enhance application performance.

Tunnel Characteristics

Each Silver Peak tunnel:

- Is bidirectional (or consists of a pair of unidirectional tunnels). The tunnel does not become operational until connectivity is established in both directions.
- Is specified by a source IP address and destination IP address, owned by the two terminating Silver Peak appliances.
- Can have the terminating appliances automatically negotiate for maximum bandwidth. Or, you can set it manually.
- By default, uses the Generic Routing Encapsulation (**GRE**) protocol — without **IPSec** — to interconnect Silver Peak appliances. You can, however, choose to enable **IPSec** bidirectionally. If you choose **UDP**, then **IPSec** is not available.
- Runs a keepalive protocol so that a tunnel failure can be detected rapidly and appropriate recovery actions initiated.

Creating a Traffic-Carrying Tunnel

If this is the appliance's first tunnel, Silver Peak recommends that you put the local and remote appliances in **System Bypass** until you've created the tunnel(s) and tuned the policies — Route, QoS, and Optimization — for the local and remote appliances. Then, when you're done, take the appliances out of **System Bypass**.

This serves a number of purposes:

- It keeps things “quiet” until you're done. Specifically because tunnels default to **Admin Up**.
- It tests Fail To Wire.

This step is recommended, but not required.

◆ To put the appliance in System Bypass when creating the first tunnel

- 1 From the **Configuration** menu, select **System**. The **Configuration - System** page appears.

Silver Peak

Appliance Time: 2009/04/14 13:59:!!

Home Configuration Reporting Monitoring Administration Maintenance Alarms

Configuration - System

Hostname:

Appliance ID: The Appliance ID is a network-unique number

Contact:

Appliance Location:

☒ Encrypt Data on the Appliance Hard Drive(s)

☐ System Bypass

Deployment Mode

☐ Bonding

☐ Router

☒ Bridge

2-port

W0

IP

L0

Appliance IP / Netmask: /

WAN Next-hop IP:

LAN Next-hop IP: (optional)

☒ Propagate Link Down

WAN Bandwidth

Max Bandwidth: (0..155000) Kbps

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

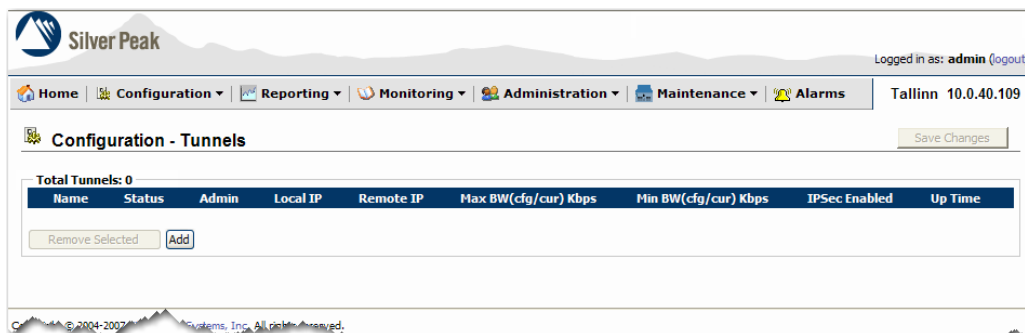
- 2 Before creating the first tunnel, select **System Bypass** and click **Apply**.
- 3 Repeat for the appliance at the remote end.

- ◆ **To create a traffic-carrying tunnel**

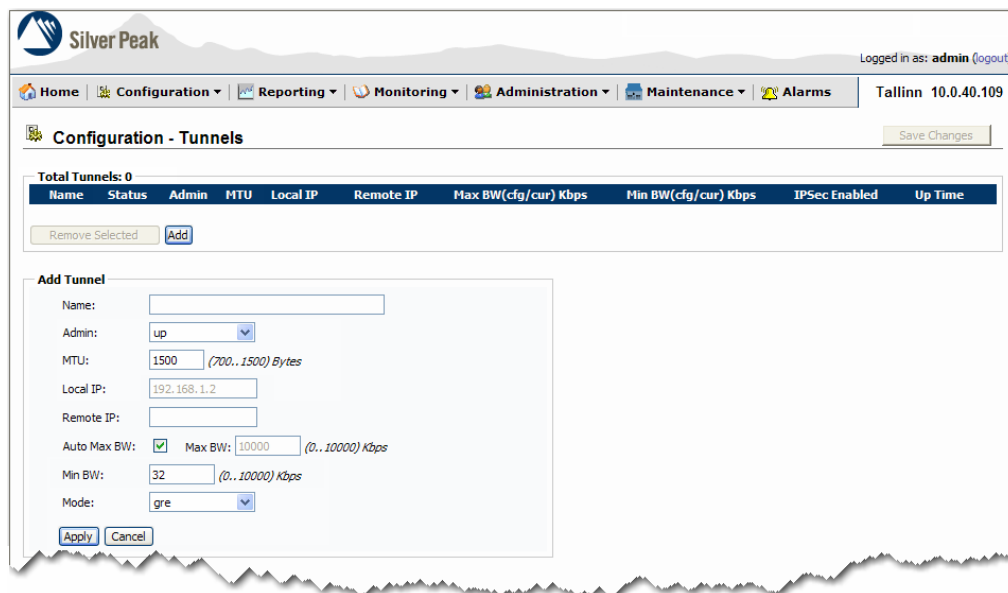
Follow this procedure.

Step 1 Create the tunnel

- 1 On the **Configuration** menu, click **Tunnels**. The **Configuration - Tunnels** page appears.



- 2 Click **Add**. On the lower half of the page, the **Add Tunnel** area appears.



- a** In the **Name** field, assign a locally significant name. Silver Peak recommends using the naming convention of *SiteA-to-SiteB*. There's a 32-character limit.

The valid character set for names (tunnel, acl, etc.) includes:

- the alphabet (both upper and lower case)
- numbers
- hyphen (-)
- underscore (_), and
- dot (.)
- Spaces are **not** supported.

- b** In the **MTU** field, accept the default of **1500**.
- c** In the **Local IP** field, the Appliance Manager prefills the IP address for the local appliance. This coincides with the address for the **wan0** interface.
- d** In the **Remote IP** field, enter the Appliance IP address that belongs to the remote appliance.
- e** Leave the **Auto Max BW** checked, as is. This setting enables the appliances to negotiate the maximum tunnel bandwidth.



For more information about this feature, see [“How Tunnel Auto BW Works” on page 189](#).

In the **Min BW** field, Silver Peak recommends that you always use the default value, **32** kbps.



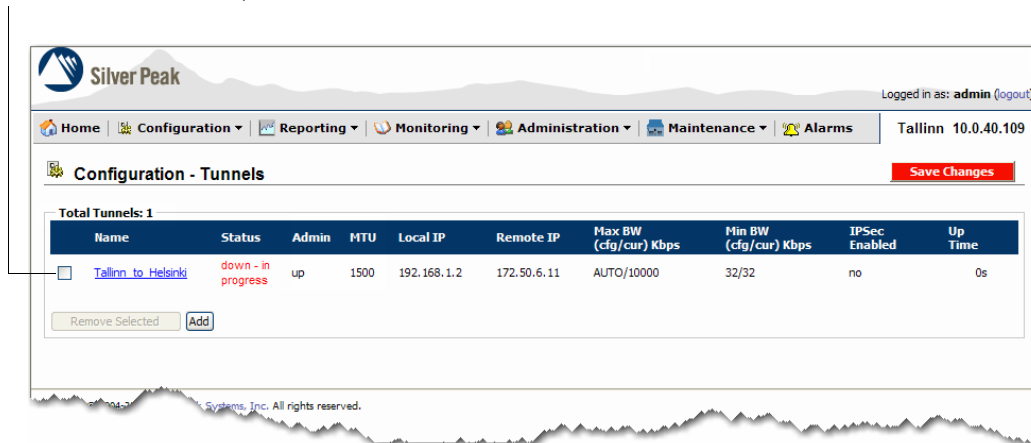
For more information about prudently setting bandwidths, see [Chapter 8, “Bandwidth Management & QoS Policy.”](#)

- f** If you want to enable IPsec, select **ipsec** from the **Mode** list. Remember that you'll need to enable this at both ends of the tunnel. In most cases, you'll leave this unchanged, opting for the default GRE tunnel.

For IPsec, the Appliance Manager autogenerates a pre-shared key. However, if you want to enter your own string, a field becomes available for that when you go back in to modify the tunnel.

- 3 Click **Apply**. The new tunnel displays in the table.

After you click **Apply**, the new tunnel displays here.
To view or edit the details, click the tunnel's name.



- 4 Click **Save Changes**.

Step 2 Configure the tunnel at the remote appliance

Login to the Appliance Manager of the remote Silver Peak appliance (at *BranchA*), and repeat [Step 1](#).

Step 3 [Suggested]

Create the entries you need for the policies — Route, QoS, and Optimization.

Entries are comprised of MATCH criteria paired with policy-specific SET actions:

- MATCH criteria are described in [Chapter 6, “Theory of Operations.”](#)
- Specific SET actions are described in each of the policy chapters:
 - [Chapter 7, “Route Policy”](#)
 - [Chapter 8, “Bandwidth Management & QoS Policy”](#)
 - [Chapter 9, “Optimization Policy”](#)

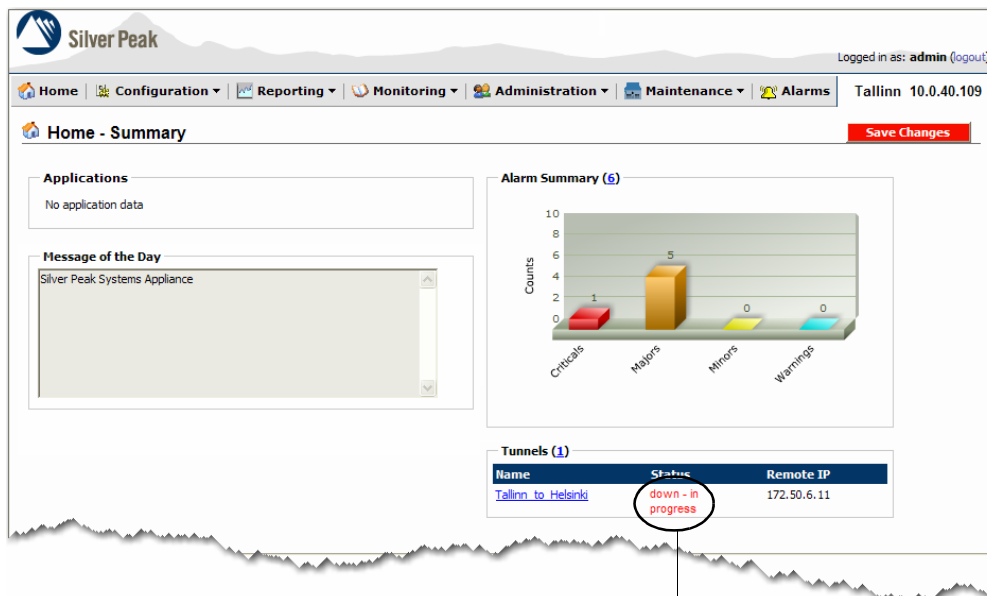


Note The Route Policy's default entry auto optimizes all TCP traffic through a tunnel it automatically determines is appropriate. Because of this, you need only add Route Policy entries for UDP, IP, or pass-through traffic. You don't need to add entries in the Optimization Policy or QoS Policy. For more information, see [Chapter 7, “Route Policy.”](#)

Step 4 At each end, verify the status of the tunnel.

You can check the status while at either site by looking for the tunnel listing on either:

- the **Home** page



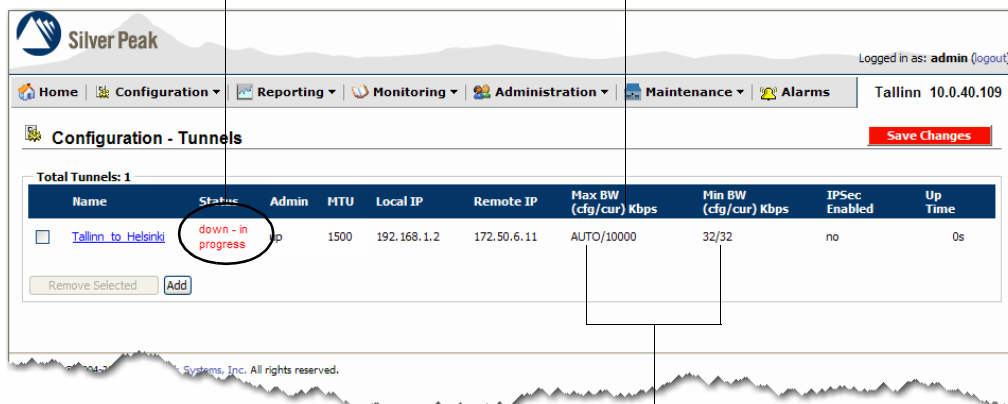
The **Status** display changes from **Down** to **Down - In progress** to **Up - Active**.

or

- the **Configuration - Tunnels** page.

...the **Status** display changes from **down** to **down - in progress** to **up - active**.

In the **Max BW** column, **AUTO/10000** means that the tunnel is set to autonegotiate the value and is currently at 10,000 kbps.



Both the **Max BW** and **Min BW** columns show the configured bandwidth followed by the current bandwidth — **(cfg/cur)**.

Possible states for tunnels are as follows:

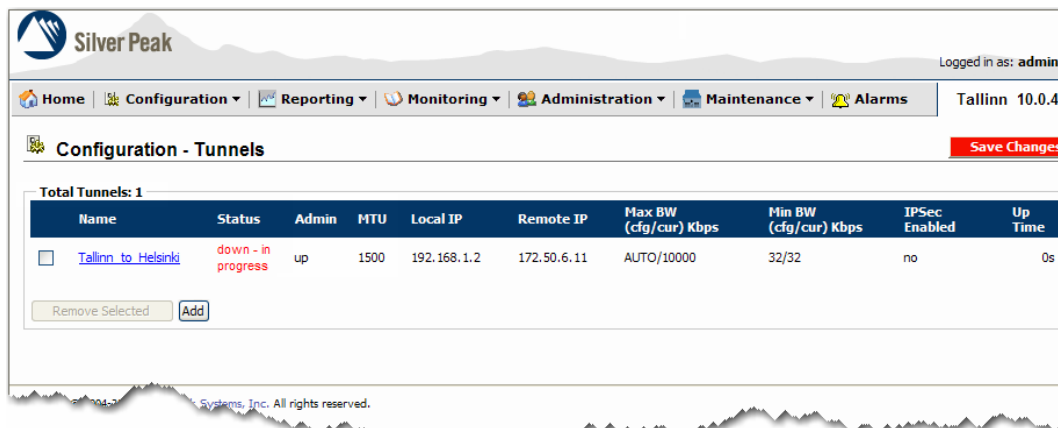
Tunnel State	Description
down	The tunnel is down. This can be because the tunnel administrative setting is down, or the tunnel can't communicate with the appliance at the other end.
down – in progress	The tunnel is down. Meanwhile, the appliance is exchanging control information with the appliance at the other end, trying to bring up the tunnel.
down – misconfigured	The tunnel is down because it's misconfigured. It needs the correct remote Appliance ID and IP address.
down – bad VRRP ip	The tunnel is down because the configured VRRP IP for the remote appliance is invalid.
up – active	The tunnel is up and active. Traffic destined for this tunnel will be forwarded to the remote appliance.
up – degraded	The tunnel is up, but it has excess WAN latency. This happens when the tunnel control packet round-trip time (RTT) exceeds the default threshold of 850 milliseconds. For more information about specific alarm text, see “Types of Alarms” on page 403 .
UNKNOWN	The tunnel status is unknown. This can be because the appliance is unable to retrieve the current tunnel status. Try again later.

The modifier, **– Idle**, can be added to any tunnel state (for example, **Up – Active – Idle**). **Idle** means that there has been no traffic in either direction on the tunnel for five minutes, and that as a result, the periodic sending of keepalives has been reduced to once a minute.

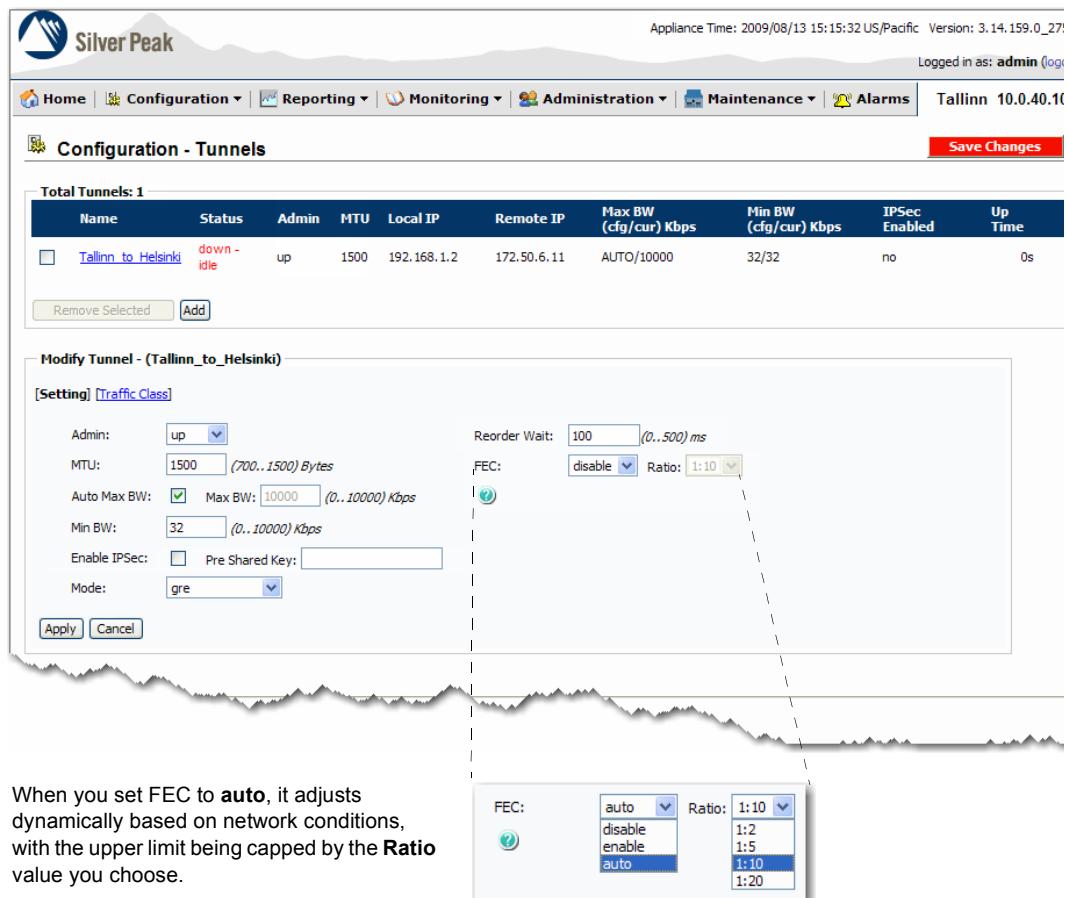
Editing a Tunnel

You can edit a tunnel according to the following procedure.

- 1 From the **Configuration** page, select **Tunnels**. The **Configuration - Tunnels** page appears.



- 2 To edit a tunnel, click its name. The tunnel's details appear under two links — **Setting** and **Traffic Class**.



Setting includes some parameters that weren't visible when you created the tunnel: **Pre Shared Key**, **FEC**, and **Reorder Wait**:

- **Reorder Wait**: Sets 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.

Set to zero to disable Packet Order Correction.

- **FEC**: Enable Forward Error Correction to reconstruct lost packets (as reported by the far-end appliance).

The **auto** setting dynamically adjusts the FEC setting based on network conditions. When set to **auto**, the value selected in the **Ratio** field becomes the upper limit, or "cap", for adjustments.



Note 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).

Note that FEC may introduce out-of-order packets if the reorder wait time is not set high enough.

- **Pre-Shared Key**: When left blank, a preshared key will be generated by the appliance.

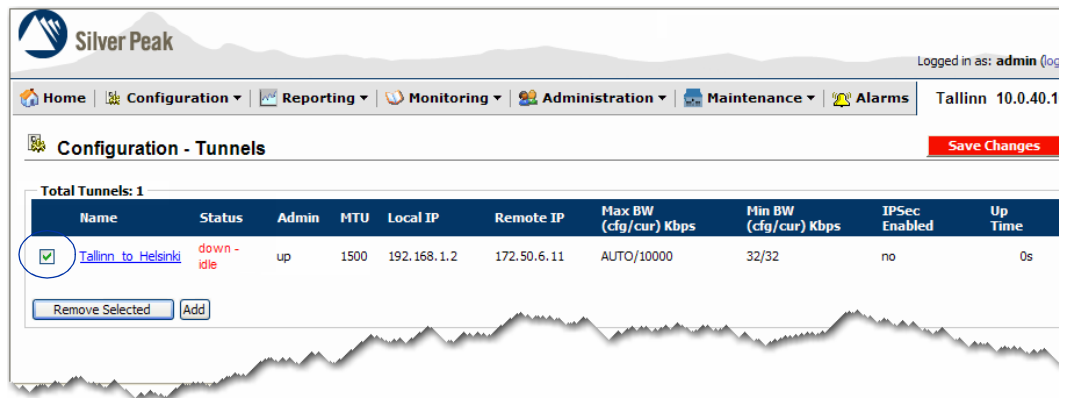
For information about the **Traffic Class** tab, see [Chapter 8, "Bandwidth Management & QoS Policy."](#)

- 3 After making your edits, click **Apply**.
- 4 Click **Save Changes**.


Deleting a Tunnel

- ♦ **To delete a tunnel**

- 1 On the **Configuration - Tunnels** page, click the check box to the left of the tunnel you want to delete.



- 2 Click **Remove Selected**. A confirmation message appears.



- 3 Click **OK**.
- 4 Click **Save Changes**.

Tunnel Compatibility Mode

Tunnel Compatibility Mode enables two nodes with mismatched software versions to keep the tunnel up and offer some basic services. Check the *Release Notes* to verify software version compatibility.

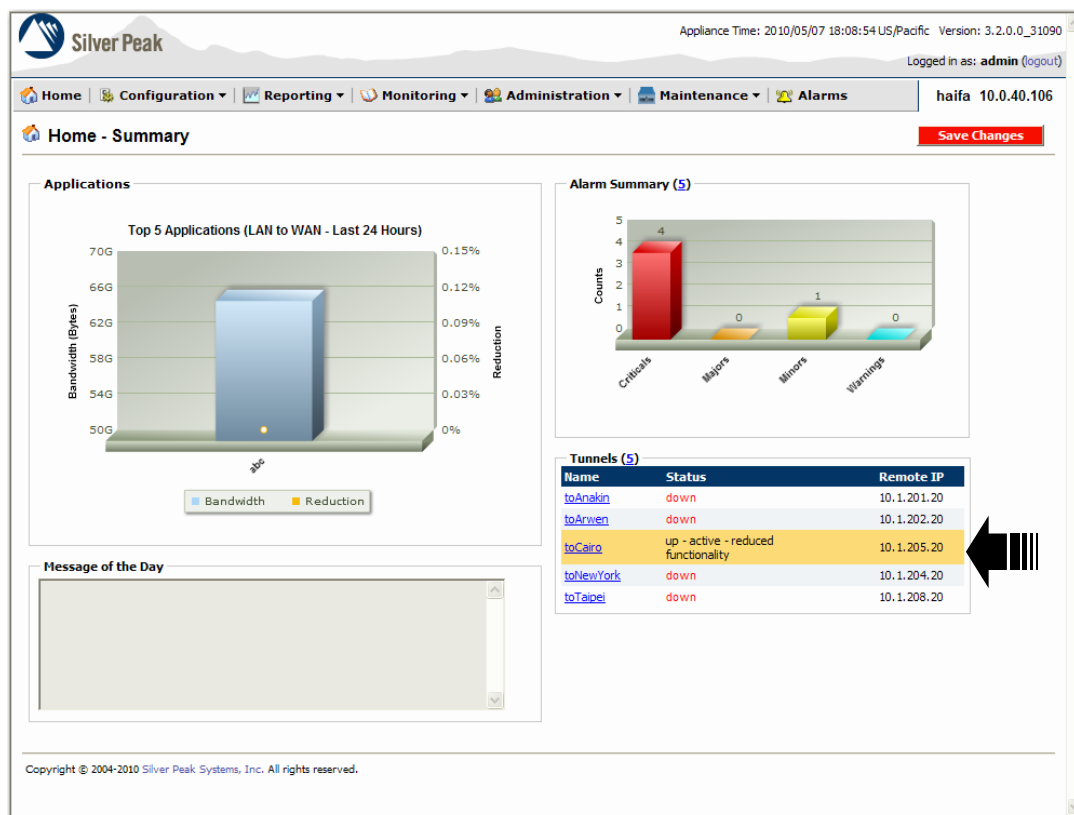
Preserved Functionalities

- Tunnel encapsulation (including IPSec, if applicable)
- QoS shaping and marking
- Forward Error Correction [FEC]
- Packet Order Correction [POC]
- packet coalescing
- statistics gathering
- network path behavior

Disabled Optimizations

- Network Memory
- Payload Compression
- TCP Acceleration
- CIFS Acceleration

Flows in such a tunnel are highlighted in orange and marked, *reduced functionality*.



The Appliance Manager displays *reduced functionality* on the home page and the following other pages:

- Alarms - Current Alarms
- Monitoring - Tunnels
- Monitoring - Current Flows
- Monitoring - Current Flow Detail



Theory of Operations

This chapter describes how the Silver Peak appliance optimizes traffic by allowing you to define flows with MATCH criteria and to direct flows with policy maps.

It also describes techniques for streamlining your network management by using Access Control Lists (ACLs), user-defined applications, and application groups.

In This Chapter

- **Processing Traffic Flows** See page 128.
- **Understanding MATCH Criteria** See page 130.
- **Using ACLs (Access Control Lists)** See page 135.
- **How Policies and ACLs Filter Traffic** See page 145.
- **Managing Applications and Application Groups** See page 147.

Processing Traffic Flows

Silver Peak appliances are interconnected by tunnels, which transport optimized traffic flows.

Policies control how the appliance filters LAN-side packets into flows — and whether an individual flow is ultimately:

- directed to a tunnel, shaped, and optimized
- processed as shaped, pass-through (unoptimized) traffic
- processed as unshaped, pass-through (unoptimized) traffic
- continued to the next applicable Route Policy entry if a tunnel goes down, or
- dropped.

The Appliance Manager has separate policies for **routing**, **optimization**, and **QoS** functions. You can create multiple versions (maps) in each policy, but only the activated map is applied.

Maps use prioritized rules, known as *entries*, to sort traffic. Entries pair MATCH criteria—which define and delineate a flow—with SET action(s) that are specific to the policy type. MATCH criteria are based on the *5-tuple*, along with some additional criteria:

- A basic 5-tuple is defined by a combination of protocol, source IP, destination IP, source port, and destination port.
- To save work, you can single out many applications by name, instead of citing a combination of protocol paired with source and/or destination port(s).
- MATCH criteria also let you filter on an incoming flow's DSCP markings
- If you've configured any VLANs in Appliance Manager, then you can also discriminately MATCH based on those VLAN tags. Otherwise, all traffic is optimized according to existing policies, independent of VLANs.

What Maps and Policies Do

Policies prescribe how the appliance handles the LAN packets it receives. The Appliance Manager provides three kinds of policies that act independently of each other on any given flow. These are the *Route policy*, the *QoS policy*, and the *Optimization policy*. Each addresses a specific set of questions.

A **Route policy** asks:

- What traffic do I want to optimize?
- What traffic do I want to send to the WAN without optimization (pass-through traffic)?
- Do I want to shape all pass-through traffic or not?
- What traffic do I want to drop?
- How do I route traffic if a tunnel goes down?

A **QoS policy** asks:

- How do I want to prioritize traffic?
- For a flow directed to Tunnel_A, what traffic class should it be assigned?
- For a flow designated as shaped, pass-through traffic, what pass-through traffic classes should it be assigned?
- How should the DSCP markings be treated? Trust the incoming LAN-to-WAN marking or re-mark for the LAN or WAN?

An **Optimization policy** asks:

- What optimization techniques do I want to apply to a given flow?
- Are there any flows that don't need all available optimization techniques?

Within each of the three policies, you can have multiple maps — with only one map being active at any given time. Each map can have multiple entries.

Let's say the traffic mix you see during after-hours network maintenance tasks differs significantly from what you see during regular business hours. To address that, you can create a separate, customized Route map for each scenario. Obviously, only one of these maps can be active at any time — so, the map you activate is the *policy*.

Default Behaviors

When you first configure the appliance, there are three policies. There is one active map in each of the three policies. Each map has one fixed default entry, listed as follows:

- By default, the **Route map** automatically optimizes all TCP flows, and the remaining non-TCP traffic passes through as unoptimized, but shaped to the maximum system bandwidth. To optimize non-TCP traffic, you simply add an entry to the Route map for the traffic you want to optimize.
- The **QoS map** puts all traffic in the default traffic class and trusts the existing DSCP markings.
- The **Optimization map** applies all optimizations — Network Memory, payload compression, TCP acceleration, and CIFS acceleration — to tunnelized flows.

No matter how many entries you add to a map, the default entry always has the last priority. Therefore, if the LAN traffic doesn't meet any user-prescribed MATCH criteria, ultimately the default entry's SET actions apply.

So, to begin optimizing traffic, you only need to configure a maximum of two items per appliance:

1 A tunnel to the remote appliance.

If automatically optimizing all TCP flows takes care of all your traffic, then you don't need to do anything else. All non-TCP traffic is sent to the WAN as shaped, pass-through traffic.

2 To optimize non-TCP traffic, create an entry in the active Route map (found in the Route Policy), specifying which non-TCP flows to direct to the specified tunnel. When they get there, they're already optimized and shaped.

All traffic that doesn't match is sent to the WAN as shaped, pass-through traffic.

Typically, you'll only make additional entries when you choose to treat some of the traffic differently. For example, you might decide to route, but not optimize, VoIP calls.

The next section delves into MATCH criteria, followed by discussions of how ACLs and Application Groups can simplify your workload.

Subsequently, the next three chapters discuss the SET actions for each policy.

Understanding MATCH Criteria

Within a policy, the appliance searches sequentially for the first match, and when it finds it, the appliance performs the associated SET action(s). If no user-configured entries match, then it applies the default entry's MATCH criteria and SET action.

Following are the SET actions available for each policy. The **default value** is highlighted in **blue**:

Policy	Parameters for SET actions	Options	For more information, see...
Route	Tunnel	<ul style="list-style-type: none"> [a specific tunnel] Auto optimized Pass-through shaped Pass-through unshaped Drop 	Chapter 7, "Route Policy"
	Tunnel Down Action	<ul style="list-style-type: none"> Pass-through shaped Pass-through unshaped Drop Continue 	
QoS	Traffic Class	<ul style="list-style-type: none"> specific to the tunnel's configuration. Default is the tunnel's Traffic Class 1. based on pass-through configuration. Default is Traffic Class 1. 	Chapter 8, "Bandwidth Management & QoS Policy"
	LAN QoS	trust-lan (plus other DSCP markings)	
	WAN QoS	trust-lan (plus other DSCP markings)	
Optimization	Network Memory	Default = ON	Chapter 9, "Optimization Policy"
	Payload Compression	Default = ON	
	TCP Acceleration	Default = ON	
	CIFS Acceleration	Default = ON	

The Appliance Manager lets you modify an active entry's MATCH criteria and/or SET actions while the entry is in use.

The rest of this section describes the basic building blocks of filtering traffic into flows:

- **Configuring MATCH Criteria in a Map or Policy** See page 131.
- **Using ACLs to Summarize MATCH Criteria** See page 132.
- **Specifying Applications and Protocols in MATCH Criteria** See page 133.

Configuring MATCH Criteria in a Map or Policy

Whereas each type of policy has different SET actions available, the options available for MATCH criteria are universal across all maps — Route, QoS, and Optimization.

All MATCH criteria include the following parameters:

- Priority
- Protocol
- Source IP Address and Netmask
- Destination IP and Netmask
- Application
- Source and Destination Ports
- DSCP
- VLAN [if VLAN tag is configured on the **Configuration - VLAN** page]

The default value for every MATCH statement field is “any”.

The following screen shows how the MATCH criteria display on the **Configuration - QoS Policy** page. The structure is the same on the Route and Optimization policy configuration pages.

There is **no implicit connection** between map names or priorities across different policies. In fact, the default map for each policy is named, **map1**.

Entries are assigned **Priority** in intervals of 10 (ten), making it easy to insert another entry later. Also, you can reorder an entry's **Priority** by editing the number.

Highlight of the MATCH criteria fields. Which ones you use, or are accessible, is a function of how you're trying to delineate flows within traffic.

The screenshot shows the Silver Peak web interface for configuring a QoS Policy. The page title is "Configuration - QoS Policy". Below the title, there's a "Map Name" dropdown set to "map1" and an "Activate" button. A "Save Changes" button is in the top right. Below the map name, there's a "New QoS Map" link and a "map1 (active)" label. The main part of the page is a table with columns for "Match Criteria" and "Set Actions". The "Match Criteria" columns are: Priority, ACL, Protocol, Src/Subnet, Dst/Subnet, Application, Src/Dst Port, DSCP, and VLAN. The "Set Actions" columns are: Traffic Class, LAN QoS, and WAN QoS. The table contains five entries with priorities 5, 10, 20, 30, and 40. The first entry has Protocol "ip", Src/Subnet "0.0.0.0/0", Dst/Subnet "0.0.0.0/0", Application "VOIP", and Traffic Class "2". The last entry has Protocol "ip", Src/Subnet "0.0.0.0/0", Dst/Subnet "0.0.0.0/0", Application "any", and Traffic Class "1".

Match Criteria									Set Actions		
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class	LAN QoS	WAN QoS
5		ip	0.0.0.0/0	0.0.0.0/0	VOIP		any	any.any	2	trust-lan	trust-lan
10		ip	0.0.0.0/0	0.0.0.0/0	Interactive		any	any.any	3	trust-lan	trust-lan
20		ip	0.0.0.0/0	0.0.0.0/0	Transactional		any	any.any	4	trust-lan	trust-lan
30		ip	0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	5	trust-lan	trust-lan
40		ip	0.0.0.0/0	0.0.0.0/0	any		any	any.any	10	trust-lan	trust-lan
default									1	trust-lan	trust-lan

Notice that between the **Priority** and **Protocol** columns, there's an **ACL** column. For more information, see "Using ACLs to Summarize MATCH Criteria" on page 132.

Incoming LAN traffic can only match one entry per policy. Therefore, best practice is to prioritize entries from most restrictive matches to the least restrictive.

Using ACLs to Summarize MATCH Criteria

Sometimes, you might want to use the same MATCH criteria in different maps. With the Appliance Manager, you can create “reusable” MATCH criteria in the form of an **ACL** (Access Control List). An ACL has one or more prioritized **rules**.

This ACL, named **Branch_to_HQ_Day**, has three prioritized, user-configured rules.

These links return you to the **Policy** page of your choice.

Configuration - Access Lists

Used by: Route Map: map1

Priority	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Action
10	ip	100.100.100.0/24	10.10.21.0/24	any		any	any.any	deny
20	ip	100.100.100.0/24	3.3.3.0/24	any		any	any.any	deny
30	ip	100.100.100.0/24	10.10.6.0/24	any		any	any.any	permit

Remove Selected Add Remove ACL

Branch_to_HQ_Night

Used by: ACL is NOT being used by any map

Each ACL is available from the drop-down list on any policy's Configuration page. Here, we specify the ACL, **Branch_to_HQ_Day**, as the MATCH criteria for this Route Policy entry.

Configuration - Route Policy

Map Name: map1 Activate

map1 (active)

Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Tunnel	Set
10	Branch_to_HQ_Day								Tallinn_to_Helsinki	
default									[auto optimized]	

Apply Cancel

If you use an ACL to specify MATCH criteria, then only its name displays on the policy's Configuration page.

Clicking the hyperlink takes you to the **Configuration - Access Lists** page (above), where you can view the details of its MATCH criteria.

ACLs are only applied when explicitly included in a policy.

Specifying Applications and Protocols in MATCH Criteria

Applications are pairings of protocol with a source and/or destination port. In the Appliance Manager, to filter traffic by application, the **Protocol** you choose determines:

- if you must select an application by name, or use its default, **any**
If you select ip from the Protocol field: See page 133.
- if you must specify the source and destination ports
If you select tcp or udp from the Protocol field: See page 134.
- if the protocol name is all that's required
If you select any other protocol (that is, not ip, tcp, or udp) from the Protocol field: See page 134.

You may specify any of the protocols in the following table. All protocols not shown in **bold** text require only the **Protocol** name in the MATCH criteria. Most of the time, you'll use **ip**, **tcp**, or **udp**:

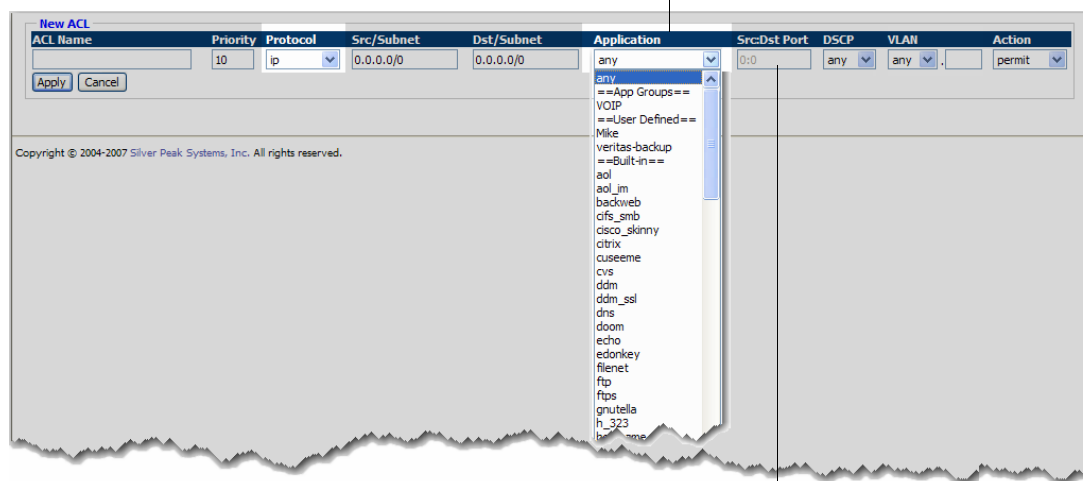
ip	ah	etherip	idpr-cmtip	ip-mobility	iso-ip	pim	vrrp
tcp	egp	fx	idrp	ipip	iso-tp4	rdp	1-255
udp	eigrp	gre	igmp	ipip4	l2tp	rsvp	
	encap	icmp	igp	ipx-in-ip	mhrp	sctp	
	esp	idpr	ip-comp	irtp	ospf	tlsp	

The following examples show the MATCH criteria fields in an ACL. These are identical to the MATCH criteria in policies.

♦ If you select **ip** from the **Protocol** field:

...then you can select a specific **Application**; the default is **any**.

The Appliance Manager filters for the application's source or destination port.



The **Src/Dst Port** field is not applicable/available.

- To include all traffic types, use the default value, **any**.
- To restrict traffic to a built-in application, a user-defined application, or a user-defined application group, select from the **Application** drop-down menu.

To create a user-defined application, see *“Defining Custom Applications” on page 152*.

To create an application group, see *“Creating and Using Application Groups” on page 156*.

♦ **If you select tcp or udp from the Protocol field:**

...then you must specify a **Source Port** and a **Destination Port**.

The screenshot shows the 'New ACL' configuration window. The 'Protocol' field is set to 'udp'. The 'Src/Subnet' and 'Dst/Subnet' fields are both set to '0.0.0.0/0'. The 'Src/Dst Port' field is set to '0:0'. The 'Application' field is disabled and shows 'any'. The 'DSCP' field is set to 'any', and the 'VLAN' field is set to 'any'. The 'Action' field is set to 'permit'. The 'Apply' and 'Cancel' buttons are visible at the bottom left.

The **Application** field is not applicable/available.

The screenshot shows the 'New ACL' configuration window. The 'Protocol' field is set to 'tcp'. The 'Src/Subnet' and 'Dst/Subnet' fields are both set to '0.0.0.0/0'. The 'Src/Dst Port' field is set to '1984:0'. The 'Application' field is disabled and shows 'any'. The 'DSCP' field is set to 'any', and the 'VLAN' field is set to 'any'. The 'Action' field is set to 'permit'. The 'Apply' and 'Cancel' buttons are visible at the bottom left.

Src:Dst Port

Means to match on...

- 0:0 any source port *and* any destination port
- 0:100 any source port *and* only destination port 100
- 100:0 only source port 100 *and* any destination port
- 100:100 only source port 100 *and* only destination port 100

This last case (100:100) is **not** OR. The only way to match on 100 for either source or destination port is to use two different MATCH entries (0:100, 100:0).

♦ **If you select any other protocol (that is, *not* ip, tcp, or udp) from the Protocol field:**

The screenshot shows the 'New ACL' configuration window. The 'Protocol' field is set to 'ah'. The 'Src/Subnet' and 'Dst/Subnet' fields are both set to '0.0.0.0/0'. The 'Src/Dst Port' field is set to '0:0'. The 'Application' field is disabled and shows 'any'. The 'DSCP' field is set to 'any', and the 'VLAN' field is set to 'any'. The 'Action' field is set to 'permit'. The 'Apply' and 'Cancel' buttons are visible at the bottom left.

The **Src:Dst Port** field is not applicable/available.

The **Application** field is not applicable/available.

Using ACLs (Access Control Lists)

Access Control Lists are useful for creating reusable MATCH criteria. In other words, you can use the same ACL as a MATCH condition in more than one policy — Route, QoS, or Optimization.


The following example shows an appliance with three ACLs, all of which are in use by a Route Map named **map1**. One ACL is expanded to show the details:

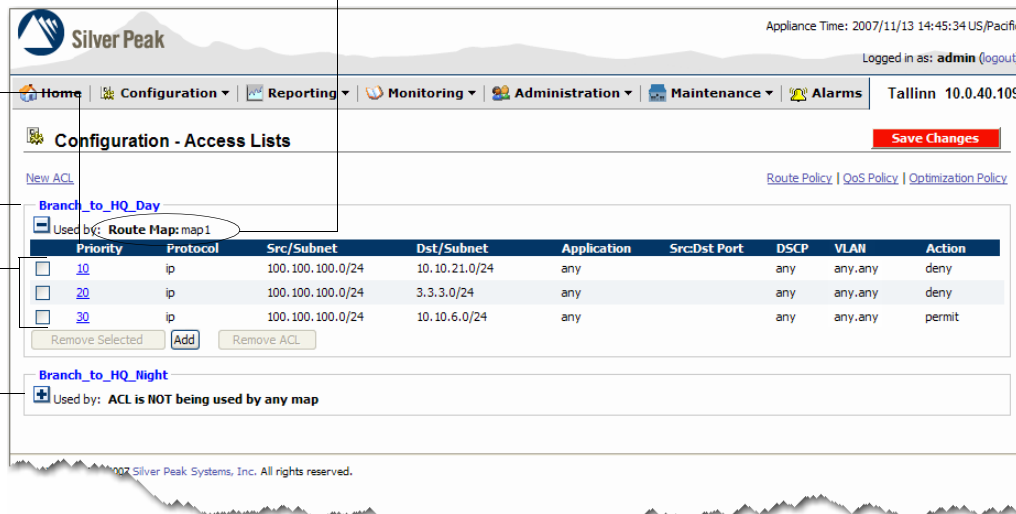
This page shows you if any map is using this ACL. However, it doesn't tell you whether or not that map is currently active. For that information, you need to check the specific policy's page.

The **Priority** uniquely identifies the rule and also specifies the order of the rule within the ACL.

ACL's name

ACL rules

When there is more than one ACL, each displays as a collapsed list until you click on the  to expand it.



Appliance Time: 2007/11/13 14:45:34 US/Pacific
Logged in as: admin (logout)
Tallinn 10.0.40.109

Configuration - Access Lists Save Changes

[New ACL](#) [Route Policy](#) [QoS Policy](#) [Optimization Policy](#)

Branch_to_HQ_Day
Used by: **Route Map: map1**

Priority	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Action
10	ip	100.100.100.0/24	10.10.21.0/24	any	any	any	any.any	deny
20	ip	100.100.100.0/24	3.3.3.0/24	any	any	any	any.any	deny
30	ip	100.100.100.0/24	10.10.6.0/24	any	any	any	any.any	permit

[Remove Selected](#) [Add](#) [Remove ACL](#)

Branch_to_HQ_Night
Used by: **ACL is NOT being used by any map**

© 2007 Silver Peak Systems, Inc. All rights reserved.

Silver Peak ACLs have the following characteristics:

- An Access Control List (ACL) consists of one or more ordered access control rules.
- Rules process sequentially, based on the unique priority number assigned during creation.
- You can reorder a rule by changing its priority.
- Each access control rule is composed of two parts:
 - The first portion specifies the filter based on source IP address, destination IP address, and application (or a pairing of protocol with source and destination port numbers). The rule applies to a particular packet only if the packet matches all conditions.
 - The second portion specifies the action: either **Permit** or **Deny**.
 - **Deny** prevents further processing of the flow by *that ACL, specifically*. The appliance then proceeds on 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**.

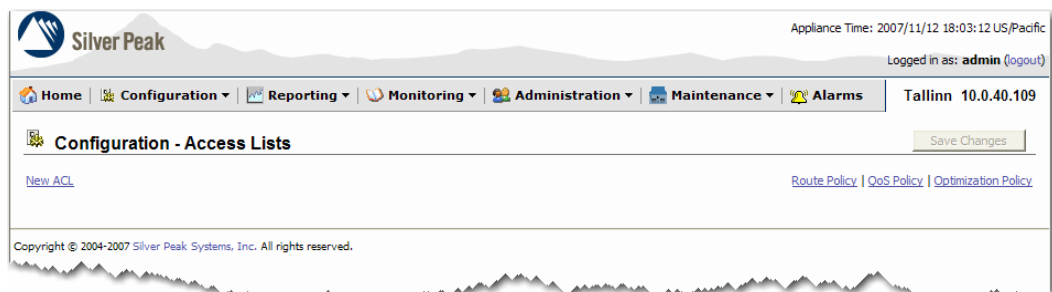
For more information, see “How Policies and ACLs Filter Traffic” on page 145.

- When you first create the ACL, it is **not** active and has no effect on the system. An ACL becomes active when it's associated with SET actions in a policy entry.
- When creating ACL rules, list the **Deny** statements first. Also, it's best to prioritize less restrictive rules ahead of more restrictive rules.
- You cannot remove an ACL if a policy is using it.

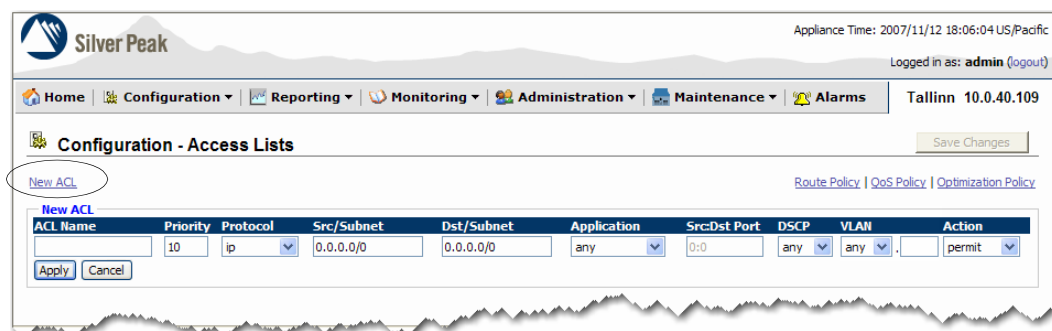
Creating an Access Control List (ACL)

♦ To create an ACL

- 1 From the **Configuration** menu, click **Access Lists**. The **Configuration - Access Lists** page appears.



- 2 Click **New ACL**. The **new ACL** area appears.



By default:

If you don't specify the...	then the rule matches...
application name	any application
source/destination port number	any value
DSCP value	any value

- 3 In the **ACL Name** field, enter a locally significant name for the ACL.
- 4 For each ACL rule, the Appliance Manager automatically assigns **Priority** in increments of ten, beginning with **10**. This gives you room to insert additional rules later, without having to reprioritize everything. There's no need to change the number now.



Note If you're using Silver Peak's Global Management System (GMS), your ACL will begin with **1010**, instead of **10**.

- 5 In specifying a type of traffic, the **Protocol** you choose determines whether you are limited to using the **Application** field, the **Src:Dst Port** field, or if you're restricted from using either one.

If you select this from the Protocol field...	then...
ip	...you can select a built-in or custom application from the Application field. To select a user-customized application that hasn't been created yet (evidenced by its not being in the drop-down menu), refer to <i>"Defining Custom Applications" on page 152</i> .
tcp	...you must specify a source port and a destination port in the Src:Dst Port field.
udp	For a review of formats and limitations, see <i>"If you select tcp or udp from the Protocol field:" on page 134</i> .
anything else	<ul style="list-style-type: none"> The Application field is not applicable/available. The Src:Dst Port field is not applicable/available.

- 6 In the **Source Network** and **Destination Network** fields, use the slash notation to specify the IP address and netmask.

For example, to stipulate an IP address of **172.50.2.0** with a netmask of **255.255.255.0**, enter **172.50.2.0/24**.

To view a table listing the notations, see *"Netmask Notations" on page 88*.

- 7 From the **DSCP** field, select the desired option. The default value is **any**.
- 8 From the **Action** field, select either **Permit** or **Deny**. If traffic satisfies the MATCH criteria, then:
- Deny** prevents further processing of the flow by *that ACL, specifically*. The appliance then proceeds on to the next entry in the policy — Route, QoS, or Optimization.
For an explanatory diagram, see *"Scenario #3 — Traffic matches ACL with Deny" on page 146*.
 - Permit** allows the matching traffic flow to proceed on to the policy entry's associated SET action(s).

The default value is **Permit**.

For illustrated detail, see *"How Policies and ACLs Filter Traffic" on page 145*.

- 9 Click **Apply**. You've now created your first Access Control List, containing one rule.

Configuration - Access Lists

Used by: ACL is NOT being used by any map

Priority	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Action
10	ip	172.30.10.0/24	172.50.10.0/24	aol		any	any.any	deny

Remove Selected Add Remove ACL

Save Changes

To add an additional rule to the same ACL, click **Add**, and repeat Steps 3 through 9.



Configuration - Access Lists

Used by: ACL is NOT being used by any map

Priority	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Action
10	ip	172.30.10.0/24	172.50.10.0/24	aol		any	any.any	deny
new 20	ip	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	permit

Apply Cancel

If you're adding a new rule, the word, **new**, displays at the beginning of the line. This distinguishes it from a rule you're modifying. In that case, a check box precedes the line's contents, instead.

Be aware that if you assign a priority number that already exists, the new rule **overwrites** the old one.

- 10 Click **Save Changes**.

Modifying an ACL Rule

You can modify ACLs (and policies) without deactivating the policy. Changes do not affect existing connections or flows, which become “stale”. Changes only affect new connections.



Note You can see a list of existing flows by going to the **Monitoring** menu and selecting **Current Flows**. Any flows that have become stale (because of a change in an ACL or policy) display with a pink background. For an example, see *“Viewing Current Flows” on page 273*.

♦ To modify an ACL Rule

- 1 In the **Priority** column, click the number of the rule you want to modify.

Clicking the rule's **Priority** makes all fields editable.

Configuration - Access Lists

Used by: Route Map: map1

Priority	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Action
10	ip	100.100.100.0/24	10.10.21.0/24	any		any	any.any	deny
20	ip	100.100.100.0/24	3.3.3.0/24	any		any	any.any	deny
30	ip	100.100.100.0/24	10.10.6.0/24	any		any	any.any	permit

Buttons: Remove Selected, Add, Remove ACL

This ACL, **Branch_to_HQ_Day**, is expanded to display all its rules.

- 2 Make the desired edits.

The route map, **map 1**, uses this ACL. However, it doesn't tell you whether or not that map is currently activated as a policy. For that information, you need to check the specific policy's page.

Configuration - Access Lists

Used by: Route Map: map1

Priority	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Action
10	ip	100.100.100.0/24	10.10.21.0/24	any		any	any.any	deny
20	ip	100.100.100.0/24	3.3.3.0/24	any		any	any.any	deny
30	ip	100.100.100.0/24	10.10.6.0/24	any		any	any.any	permit

Buttons: Apply, Cancel

- 3 Click **Apply**.
- 4 Click **Save Changes**.

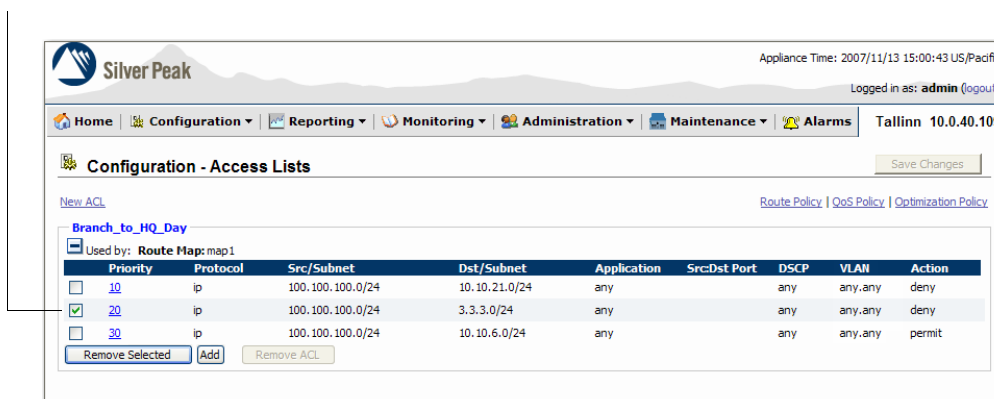
Removing an ACL Rule

You can remove ACL rules without deactivating the policy. Changes do not affect existing connections or flows. That is, connections will not reset if they no longer satisfy the MATCH criteria. Changes only affect new connections.


◆ To remove an ACL Rule

- 1 In the leftmost column, click the check box(es) of the ACL rule(s) you want to remove.

When a check box is selected, a green arrow displays.

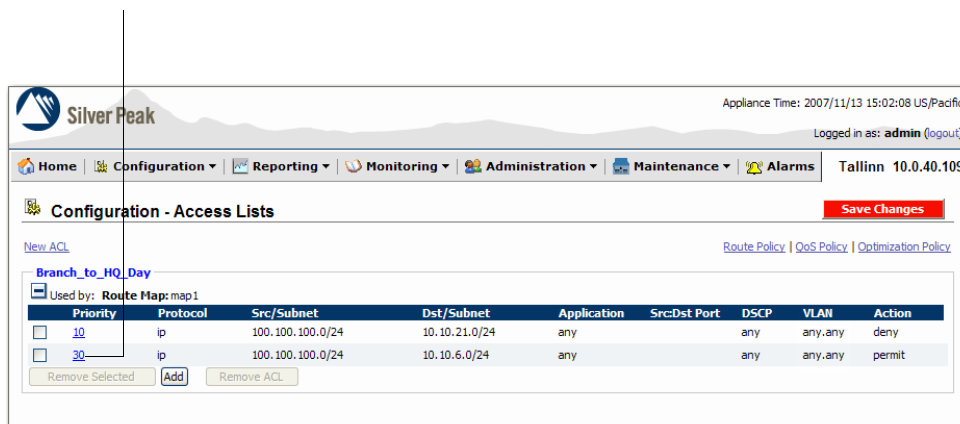


- 2 Click **Remove Selected**. A warning appears, asking you to confirm the deletion(s).



- 3 Click **OK**.

Removing an ACL rule does not cause any renumbering of the remaining rules.



- 4 Click **Save Changes**.

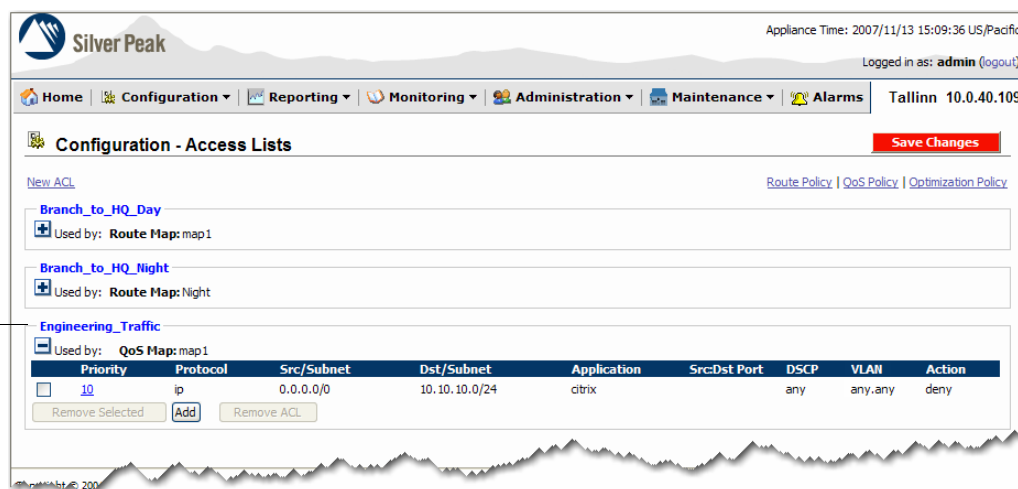
Removing an ACL

You cannot remove an ACL if it's associated with any map. To delete an ACL, first remove it from the associated map(s). Following is such an example.

♦ To remove an ACL associated with a map

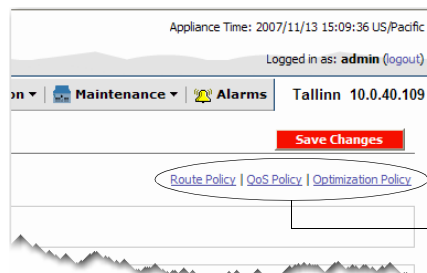
In this example, we'll remove the ACL, **Engineering_Traffic**.

- 1 On the **Configuration - Access Lists** page, examine the ACL you want to remove.



This ACL is associated with the QoS Map, **map1**. Notice that there's no way to select the ACL for removal. You have to dissociate it from the QoS map first.

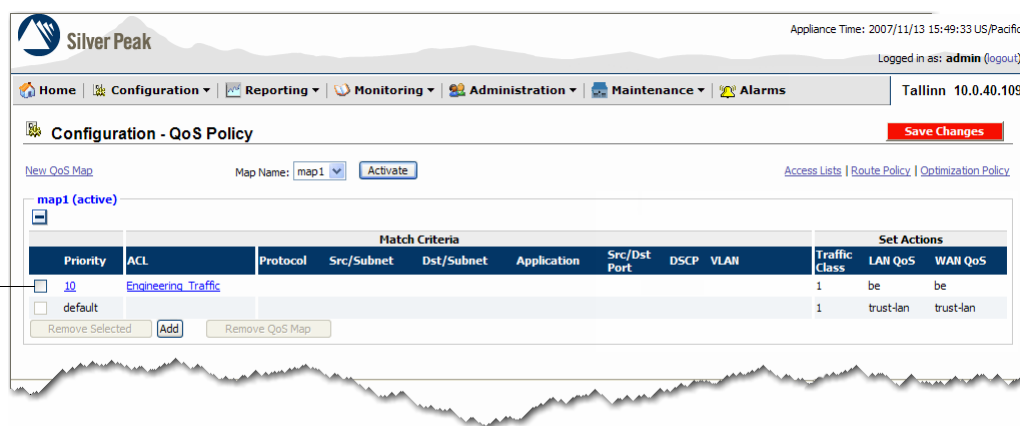
- 2 To access the page where the QoS maps are, click **QoS Policy**.



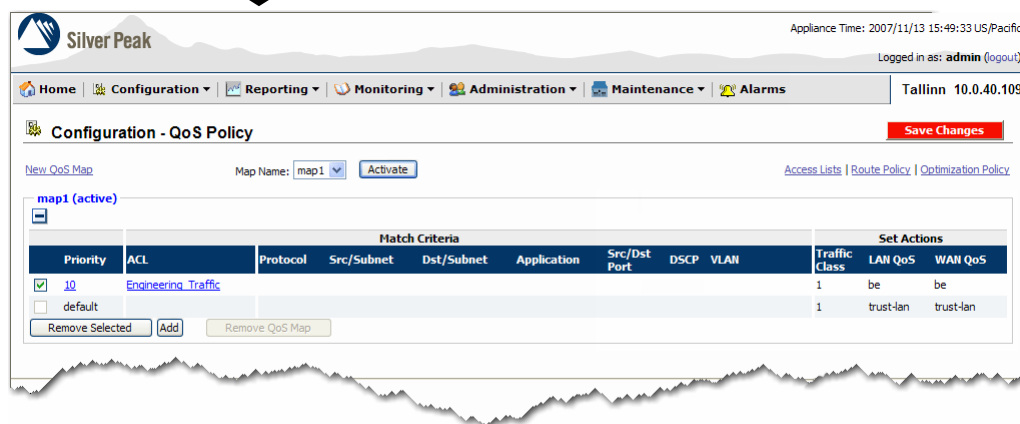
The Appliance Manager provides links so you can move between the **Configuration - Access Lists** page and each of the policy configuration pages.

When you click on **QoS Policy** and arrive at the **Configuration - QoS Policy** page, the **QoS Policy** link is absent (obviously) — replaced by the **Access Lists** link.


The **Configuration - QoS Policy** page appears.



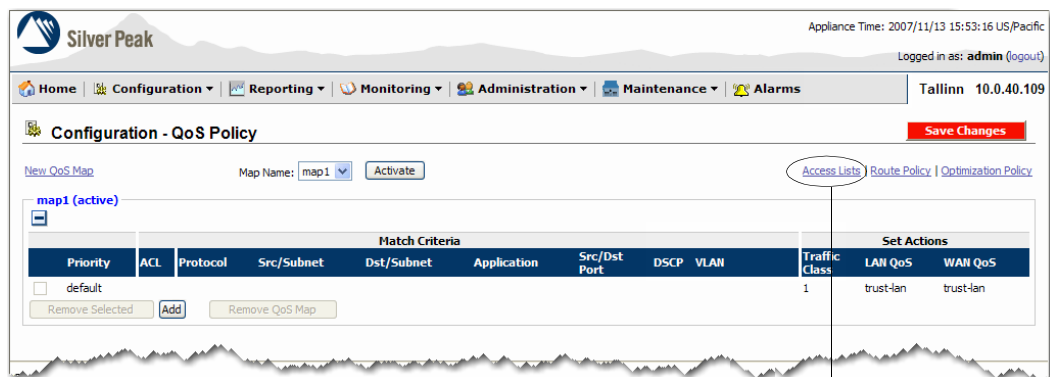
Click the check box for the entry that contains the ACL. The buttons change state.



- 3 Click **Remove Selected**. You'll be prompted for confirmation.

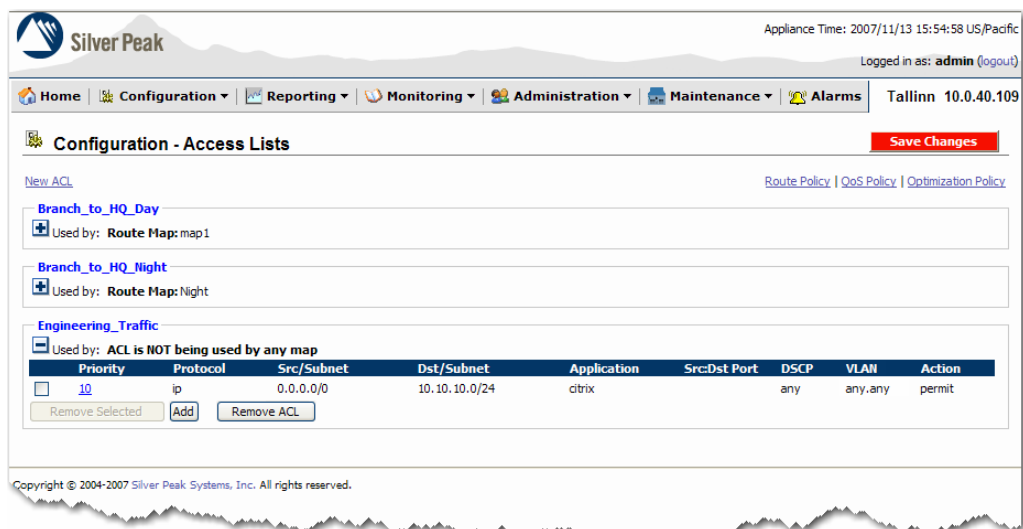


- 4 Click **OK**. The revised policy displays, with the ACL removed.




To return to the **Configuration - Access Lists** page, click **Access Lists**.


- 5 Return to the **Configuration - Access Lists** page, and expand the ACL.



- 6 Click **Remove ACL**. To confirm when prompted, click **OK**.



Appliance Manager removes the ACL.



- 7 Click **Save Changes**.

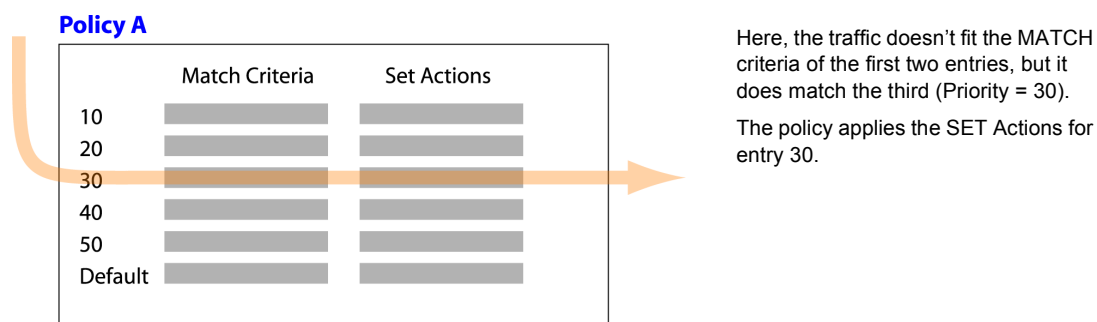
How Policies and ACLs Filter Traffic

The following three scenarios illustrate how policies and ACLs interact to isolate flows:

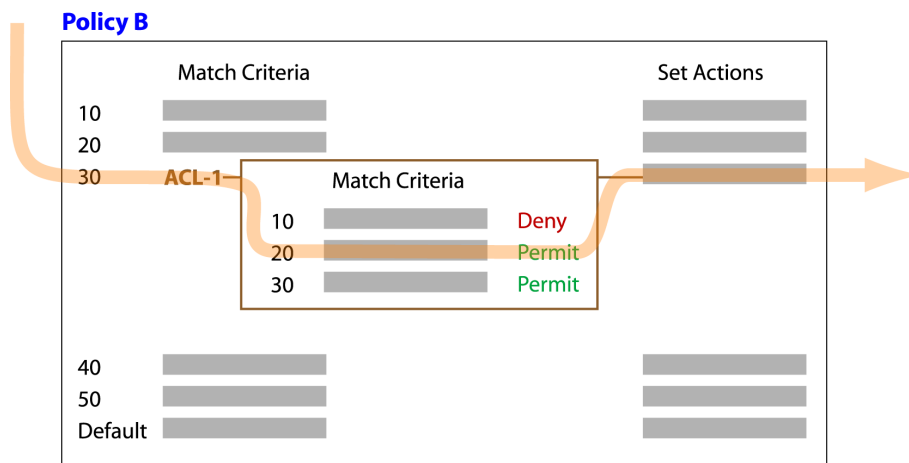
- **Scenario #1 — Policy with no ACLs in MATCH Criteria** See page 145.
- **Scenario #2 — Traffic matches ACL with Permit** See page 145.
- **Scenario #3 — Traffic matches ACL with Deny** See page 146.

It's important to remember that ACLs are *only* applied when called out for use in a policy's MATCH criteria.

♦ Scenario #1 — Policy with no ACLs in MATCH Criteria

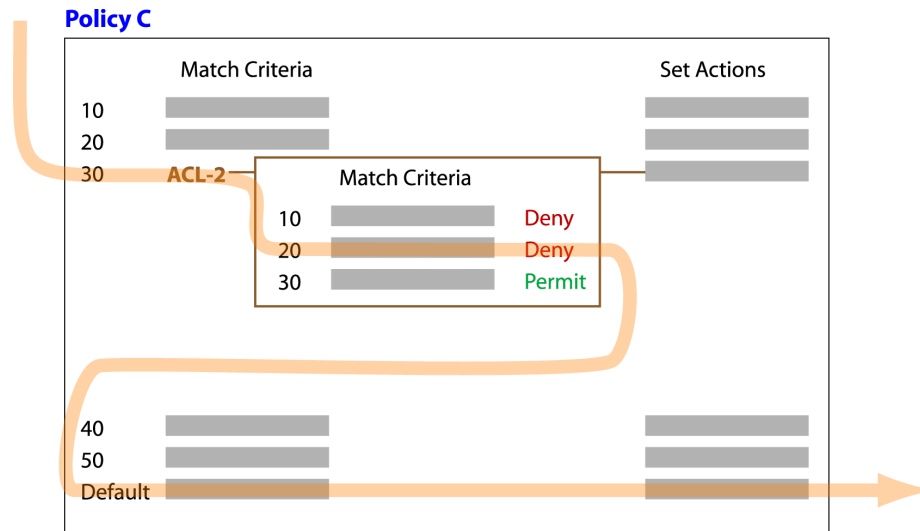


♦ Scenario #2 — Traffic matches ACL with Permit



- The traffic comes to entry **30** in the policy, where **ACL-1** defines the MATCH criteria. **ACL-1** has three rules.
- The traffic doesn't match ACL Rule 10, but it does match ACL Rule 20.
- ACL Rule 20 has a **Permit** action, so the appliance applies the SET actions for Policy entry 30.

◆ **Scenario #3 — Traffic matches ACL with Deny**



- a The traffic arrives at entry **30** in the policy, where **ACL-2** is the MATCH criteria. **ACL-2** has three rules.
- b The traffic doesn't match ACL Rule 10, but it does match ACL Rule 20.
- c ACL Rule 20 has a **Deny** action, so it prevents further processing of that ACL. Traffic looks for a match with the next policy entry.
- d No other user-configured policy entries fit, so the Default entry processes the traffic.

Managing Applications and Application Groups

The Appliance Manager provides you with many ways to define and organize the applications you use. These include the following:

- **Built-in Applications** See page 147.
- **Defining Custom Applications** See page 152.
- **Creating and Using Application Groups** See page 156.

Built-in Applications

Silver Peak appliances have a long list of *built-in applications*. For the latest information regarding default port numbers, see <http://www.iana.org/assignments/port-numbers>.

Name	Port Number(s)	Description and Inclusions
aol	5191-5193	America Online
aol_im	4443, 5190	AOL/ICQ Instant Messenger AOL/ICQ Image Transfer
backweb	370	Backweb is a generic, background downloading tool that software vendors can incorporate into their product to download data (for example, product updates) to the user's PC.
cifs_smb	139, 445	Microsoft's Common Internet File System/Server Message Block protocol
cisco_skinny	2000-2001	Cisco Skinny (SCCP) Control
citrix	1494, 1604	Citrix Citrix - ICA WinFrame Server
cuseeme	7648-7652, 24032	Cu-SeeMe Videoconferencing
cvs	--	Concurrent Versioning System
ddm	--	Distributed Data Management
ddm_ssl	--	Secure DDM (DDM over SSL)
dns	53	Domain Name Services Domain Name Service (DNS) over TCP (RFC 793)
doom	666	DOOM Game - Id Software
echo	7	Echo Protocol (RFC 863)
edonkey	4661-4662, 4665	eDonkey2000 Server
filenet	32768-32771	FileNet TMS Transfer Management System FileNet RPCRemote Procedure Call FileNet NCHNetwork Clearinghouse FileNet RMIRemote Method Invocation
ftp	20-21	File Transfer Protocol - Control Port (RFC 959) File Transfer Protocol - Data Port (RFC 959)
ftps	989-990	Secure FTP Data Port (FTP Data Port over SSL) Secure FTP Control Port (FTP Control Port over SSL)
gnutella	6346-6347	Gnutella Server Gnutella Router

Name	Port Number(s)	Description and Inclusions
h_323	1718-1720	H.323 Videoconferencing Call Signaling & Control
hostname	101	NIC Internet Hostname Server Protocol (RFC 953)
http	80, 591, 8008, 8080	WWW Hypertext Transfer Protocol (HTTP - RFC 1945, 2068, 2069, 2109, 2145) HTTP Alternate (see Port 80 for HTTP)
https	443	Secure HTTP (HTTP over SSL)
ibm_db2	523, 3700-3701	IBM DB2 Administration Server IBM-DB2 Connection Service IBM-DB2 Interrupt Connection Service
imap	143, 220	Internet Message Access Protocol (IMAP) Internet Message Access Protocol (IMAP) (v2 - RFC 1064, v4 - RFC 1730) Internet Message Access Protocol (IMAP) (v3 - RFC 1203)
imap4s	585, 993	Secure IMAPv4 (IMAPv4 over SSL)
ipsec	--	A collection of IP security measures that comprise an optional tunneling protocol for IPv6.
irc	194	Internet Relay Chat Protocol (RFC 1459)
irc_ssl	994	Secure IRC Chat (IRC Chat over SSL)
isakmp	500	Internet Security Association and Key Management Protocol (ISAKMP)
ivisit	9943, 9945, 56768	iVisit - Internet Video CHAT
kazaa	1214, 2002	Kazaa-Morpheus-Grokster P2P File Sharing Kazaa P2P File Sharing - File Download
kerberos	88	Kerberos
l2tp	1701	Layer 2 Tunneling Protocol
ldap	389	Lightweight Directory Access Protocol (LDAP over TCP - RFC 1777)
ldaps	636	Secure LDAP (LDAP over SSL)
lotus_cc_mail	3264	Lotus cc:Mail
lotus_notes	1352	Lotus NOTES
matip	350-351	MATIP (RFC 2351)
ms_exchange	--	Microsoft Exchange Server
ms_media	1755	Microsoft Media Player Microsoft Media Streaming Payload
ms_messenger	1863, 6891-6901, 7001	MSN Messenger MSN Messenger File Transfer MSN Messenger Voice
ms_odbc	--	Microsoft Open DataBase Connectivity
ms_ole	--	Microsoft Object Linking and Embedding
ms_rpc	135	Microsoft Remote Procedure Call

Name	Port Number(s)	Description and Inclusions
ms_sql	1433-1434	Microsoft-SQL Server Microsoft-SQL Monitor
ms_terminal_services	3389	Microsoft Terminal Services Microsoft Terminal Server
ms_zone	6073, 28000-29100, 47624	MSN Zone MSN Zone DirectX 7.0 Control MSN Zone DirectX 8.0 Control
nameserver	42	Name Server
netbios	137	Network Basic Input/Output System NetBIOS-over-TCP/UDP - Datagram Service (RFC 1001, 1002) NetBIOS-over-TCP/UDP - Name Service, WINS (RFC 1001, 1002) NetBIOS-over-TCP/UDP - Session Service (RFC 1001, 1002)
nfs	2049	Sun Network File System
nnntp	119	Network News Transfer Protocol (NNTP - RFC 977)
nntps	563	Secure NNTP (NNTP over SSL) or TLS [Transport Layer Security]
novell	--	Netware Core Protocol
ntp	123	Network Time synchronization Protocol -- protocol providing time across a network with precise clocks; implemented over TCP and UDP
openwindows	2000	Open Windows
oracle	1521, 1525, 1529, 1571, 1575, 1600, 1610, 1620, 1748, 1754, 1808-1809	Oracle Co-Author Database Oracle Enterprise Manager Oracle Names Database Oracle Remote Database Oracle Server Oracle TNS Server Oracle VP
pcanywhere	5631-5632	pcANYWHERE pcANYWHERE - Data
pcmail	158	PCMail PCMail Server (RFC 1056)
peoplesoft	--	PeopleSoft enterprise application software
pop	109-110	Post Office Protocol Post Office Protocol - Version 2 (RFC 937) Post Office Protocol - Version 3 (RFC 1725)
pop3s	995	Secure POP3 Mail (POP3 Mail over SSL)
pptp	1723	Microsoft Point-to-Point Tunneling Protocol (PPTP)
printer	515	Printer Spooler
quake	2600	Quake Quake-II
rlogin	513	BSD RLOGIN (remote login a la telnet)

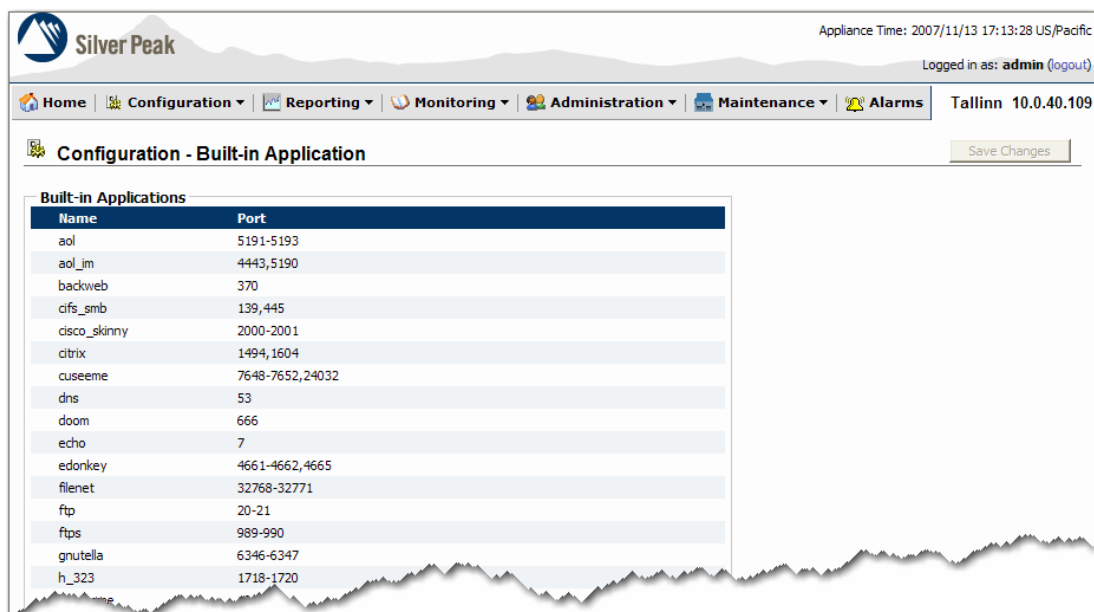
Name	Port Number(s)	Description and Inclusions
routing	179	BGP Border Gateway Protocol
	--	OSPF Open Shortest Path First
	--	IGP Interior Gateway Protocol
	--	RTM Routing Table Messaging Protocol
	--	EIRGP Enhanced Interior Gateway Routing Protocol
rtcp	5005	Real Time Transport Control Protocol
rtsp	554, 8554	Real Time Stream Control Protocol (RTSP - RFC 2326)
sap	3200, 3300-3388, 3390-3399, 3600-3699	Service Advertising Protocol (a NetWare protocol) SAP R/3
sgcp	440	Simple Gateway Control Protocol
sgmp	--	Signaling Gateway Monitoring Protocol
shell	514	RCMD, RSH (Remote execution; like exec, but automatic)
sip	5060	Session Initiated Protocol, or Session Initiation Protocol, an application-layer control protocol; a signaling protocol for Internet Telephony
sip_tls	--	SIP using Transport Layer Security
smtp	25	Simple Mail Transfer Protocol (SMTP - RFC 821)
smtps	465	Secure SMTP (SMTP over SSL)
snmp	161-162	Simple Network Management Protocol (RFC 1902, 1905) Simple Network Management Protocol - Traps (RFC 1902, 1905)
sql	156	SQL (Structured Query Language) SQL Services Oracle SQL*NET
ssh	22, 614	SSH (Secure Shell) Remote Login Protocol
sshell	--	Secure shell (shell over SSL)
sun_rpc	111	Sun Remote Procedure Call (RFC 1831)
sybase	1498, 2638	Sybase SQL Anywhere (v6.0) Sybase SQL Anywhere (v5.x & older)
syslog	514	Syslog
t_120	1503	T.120 Whiteboarding
tacacs	49, 65	Login Host Protocol (TACACS) TACACS - Default Server Port (RFC 1492)
telnet	23	Telnet (RFC 854)
telnets	992	Secure TELNET (TELNET over SSL)
tftp	69	Small, simple FTP used primarily in booting diskless systems
timbuktu	407, 1417-1420	Timbuktu
time	37	Time Protocol (RFC 868)
uucp	540	UUCP (Unix-to-Unix copy protocol) UUCP Path Service (RFC 915)

Name	Port Number(s)	Description and Inclusions
xwindows	6000-6063	X Window (x11) System
yahoo_im	5000-5010, 5050, 5055, 5100-5101	Yahoo Instant Messenger Yahoo Instant Messenger file transfer Yahoo Instant Messenger voice Yahoo Instant Messenger webcam
yahoo_games	11999	Yahoo Games

You can also define custom applications, by associating an application name with a protocol and a port number. These *user-defined applications* are accessible alongside built-in applications and application groups.

- ◆ **To view the list of built-in applications**

In the menu bar, click **Configuration > Application > Built-in** to access the **Configuration - Built-in Application** page.



When you create MATCH criteria in policies or ACLs, you have access to these applications via a drop-down list.

Defining Custom Applications

By default, the Appliance Manager already lists over 100 built-in applications.

When you want to differentiate traffic that is not already included in the built-in applications list, you can create custom applications, using the following parameters:

- priority
- protocol
- source IP address, subnet, and port
- destination IP address, subnet, and port
- DSCP
- VLAN

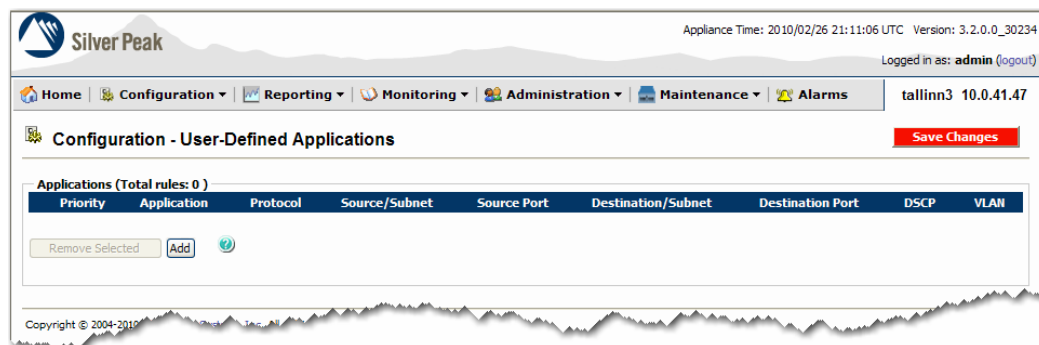
Once you've created and named an application, you can access it in the **Match Criteria** when configuring any of the traffic maps (Route, Optimization, QoS) or the Access [Control] Lists (ACLs).



Note When creating a custom application on one appliance, you must create the same application on each corresponding device so that there is reporting symmetry. Doing so ensures that if an application has a name on one appliance, it isn't listed as **unassigned application** on another, paired appliance.

♦ To create a user-defined application

- 1 In the menu bar, click **Configuration > Application > User-Defined** to access the **Configuration - User-Defined Applications** page.



- 2 Click **Add**. A set of editable fields appears.

The screenshot shows the Silver Peak web interface for configuring user-defined applications. The 'Add Application Rule' dialog is displayed over the 'Configuration - User-Defined Applications' page. The dialog contains the following fields:

- Priority:** 10
- Application:** (empty text field)
- Protocol:** tcp (dropdown menu)
- DSCP:** any (dropdown menu)
- VLAN:** any (dropdown menu)
- Source:**
 - IP Address/Subnet:** 0.0.0.0/0
 - Port:** 0
- Destination:**
 - IP Address/Subnet:** 0.0.0.0/0
 - Port:** 0

Buttons for 'Apply', 'Cancel', and a confirmation icon are at the bottom of the dialog.

- a Accept or modify the **Priority** number:
 - For the first entry, the **Priority** field defaults to **10**. By default, adding a rule/application increments the last **Priority** by **10**.
 - The range is from **1** to **50000** (no commas).
 - The numbers are unique across all applications.
 - If there is any overlap in the matching criteria within the custom applications, then the lower **Priority** number prevails for that criteria.
- b In the blank **Application** field, enter a name for your application.
 - **Application** names cannot be modified.
 - To rename an application, you must delete the existing one and create a new one.
- c From the **Protocol** drop-down menu, select the protocol.
 - You can create an application that uses the same port with **tcp** and with **udp**. In that case, use the option, **tcp/udp**.
 - If you select **tcp**, **udp**, or **tcp/udp**, then you can access the **Port** field. If you don't select one of those three specific protocols, then the **Port** field(s) are unavailable.
- d In the **DSCP** field, accept the default, **any**, or select another option from the drop-down menu.
- e If you want, specify a **VLAN** tag.
For more information about VLANs, see [Chapter 11, "Configuring and Managing VLANs."](#)

f Enter **Source** and/or **Destination** information:

- An IP address can specify a subnet. For example: **10.0.0.10/24**.
- An IP address can specify a range. For example: **10.0.0.20-30**.
- To specify *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**.

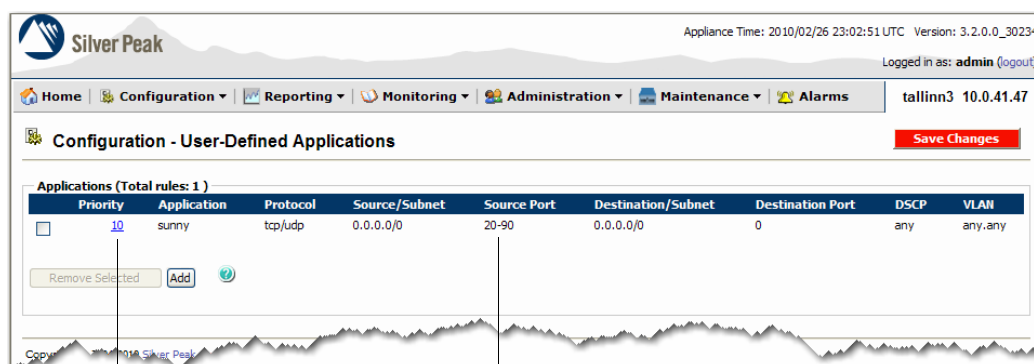


Note When it comes to flow and application statistics reports, user-defined applications are always checked before built-in applications.

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

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

- To specify *any port*, use **0**.
- Separate multiple items with any of the following: a line break, a comma, or a single space.

3 Click **Apply**.

If you want to edit an entry later, just click on the hyperlink in the **Priority** column.


In this example, the user-defined application's port range of **20–90** overrides claim to the following built-in ports:

- | | |
|--------------|--------|
| • dns | 53 |
| • ftp | 20–21 |
| • http | 80 |
| • kerberos | 88 |
| • nameserver | 42 |
| • smtp | 25 |
| • ssh | 22 |
| • tacacs | 49, 65 |
| • telnet | 23 |
| • time | 37 |

4 Click **Save Changes**.

♦ **To remove a user-defined application from the database**

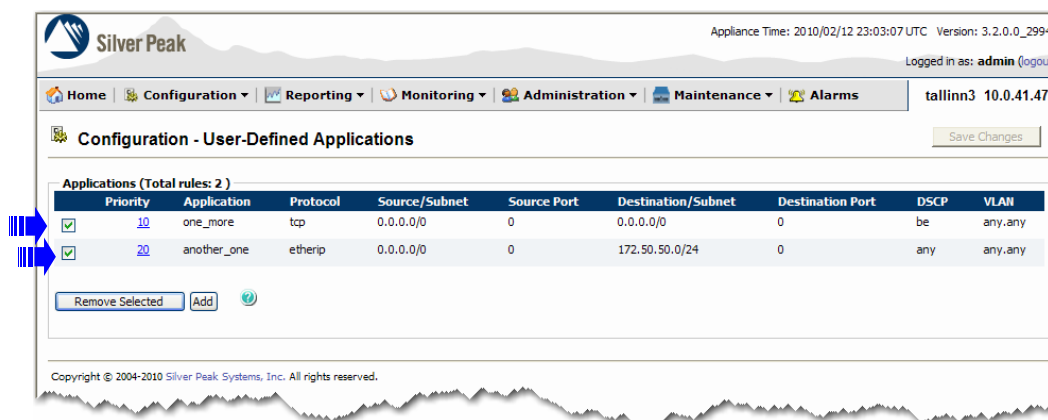
The **Configuration - User-Defined Application** page doesn't provide direct information about whether or not a custom application is in use. If you try to delete an application that an ACL and/or a policy *is* using, Appliance Manager provides the following message:



In that case, you'll need to either remove the ACL or edit the ACL to remove (or change) that application.


Once the application has no associations, you can delete it from the table on the **Configuration - User-Defined Application** page.

- 1 For the entry you want to remove, click its check box. If you want to remove more than one application, you can select all their check boxes. A green check mark displays.



WARNING: Clicking on the application name makes the line editable and displays different buttons from those shown here. To back out of that situation, just click **Cancel**.

- 2 Click **Remove Selected**. A confirmation message appears.



- 3 Click **OK**.

If the application were associated with a policy or ACL, Appliance Manager would display that message now.

- 4 Click **Save Changes**.

Creating and Using Application Groups

If your ACLs or policy maps contain MATCH conditions that involve multiple applications, you can simplify the MATCH criteria with **application groups**. Application groups are identifiers that you can create to represent a list of applications.

For example, an application group, *encrypted*, might include SSH, HTTPS, and SFTP.

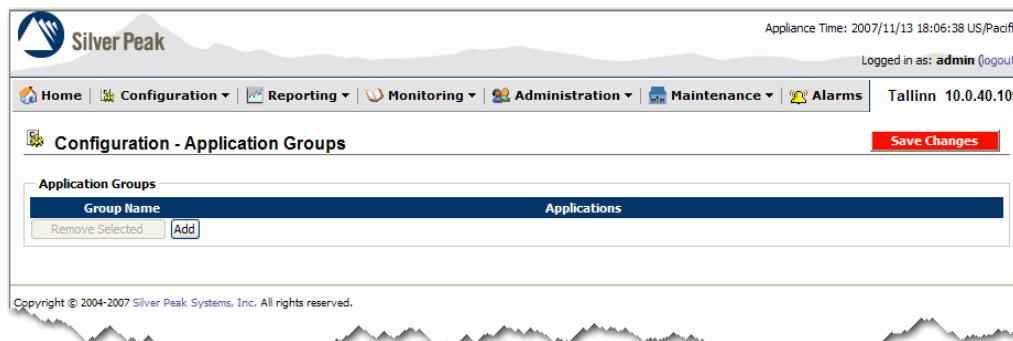
Application groups have the following properties:

- Any application can belong to multiple groups.
- You can modify an application group even when it's used by an ACL or policy map.
- You can only remove an application group if it is **not** associated with an ACL or map.

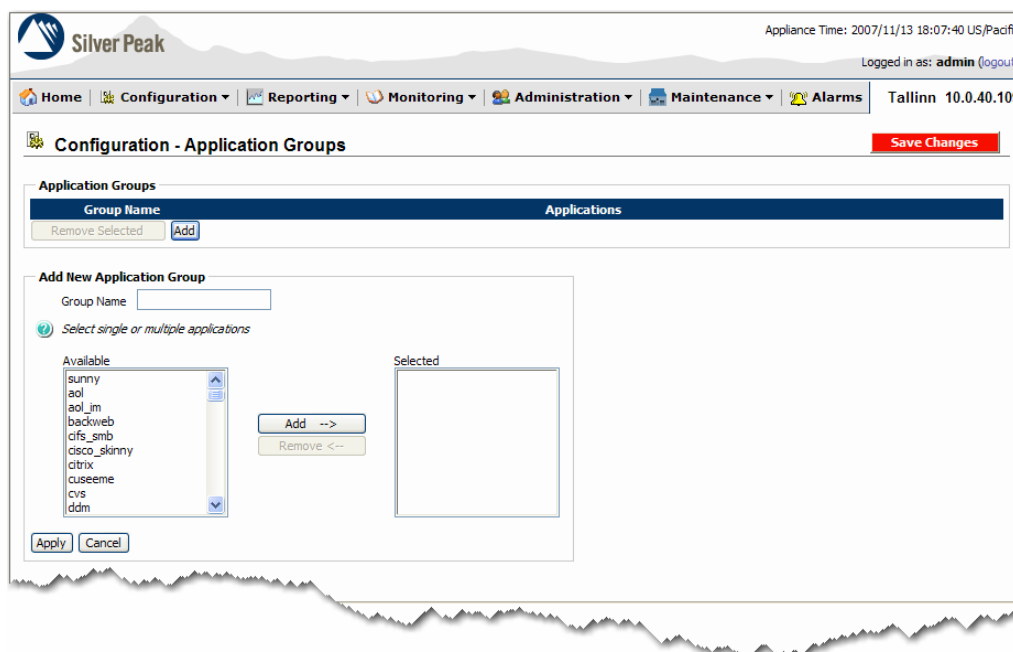
♦ To create an Application Group

When creating an application group on one appliance, you must create the same application group on each corresponding device so that there is reporting symmetry.

- 1 In the menu bar, click **Configuration > Application > Groups** to access the **Configuration - Application Groups** page.



- 2 Click **Add**. The **Add New Application Group** area appears.



- 3 In the **Group Name** field, enter a descriptive name.
- 4 From the **Available** list, select the built-in and/or user-defined applications you want and click **Add**.
 - To select multiple adjacent applications, drag across the names you want to select.
 - To select nonadjacent applications, select a single cell, and then hold down the **CTRL** key while you click other cells that you want to select.

The applications move to the **Selected** list.

To remove any or all applications from the group, select the applications and click **Remove**.

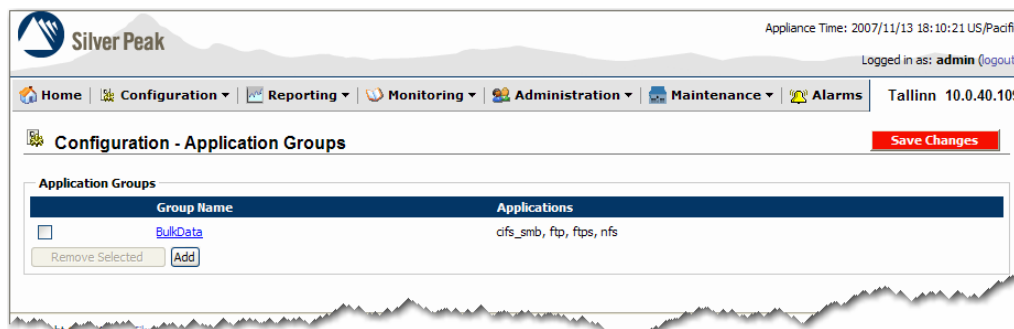
- 5 Click **Apply**. The **Add New Application Group** area disappears, and the new group displays in the table.

- 6 Click **Save Changes**.

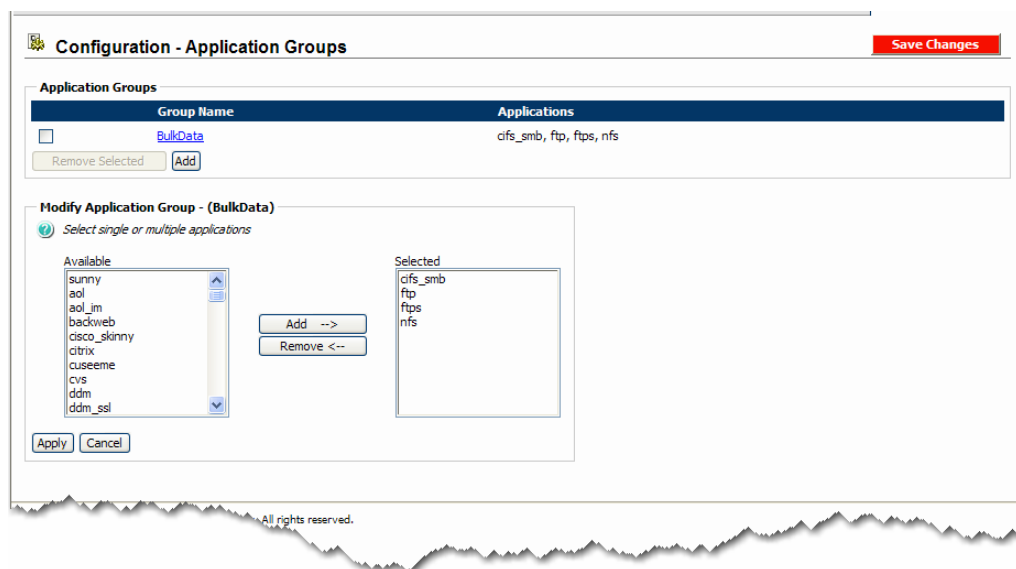
♦ To modify an Application Group

For any group you create, you can add applications or remove them — as needed — even if the group is specified in an ACL or policy.

- 1 In the menu bar, click **Configuration > Application > Groups** to access the **Configuration - Application Groups** page.



- 2 In the **Group Name** column, click the name of the group you want to modify. The **Modify Application Group** area displays.



- 3 Do one of the following:
 - To add an application to the group, go to the **Available** list, select the built-in and/or user-defined applications you want, and click **Add**.
 - To remove an application from the group, go to the **Selected** list, select the built-in and/or user-defined applications you want, and click **Remove**.


To select multiple adjacent applications, drag across the names you want to select.

To select nonadjacent applications, select a single cell, and then hold down the **CTRL** key while you click other cells that you want to select.


- 4 Click **Apply**. The **Modify Application Group** area disappears, and the modified results display in the table.
- 5 Click **Save Changes**.

- ◆ **To remove an Application Group from the database**

The **Configuration - Application Groups** page doesn't tell you whether or not an application group is in use. When you try to remove any application, the Appliance Manager first asks you to confirm:



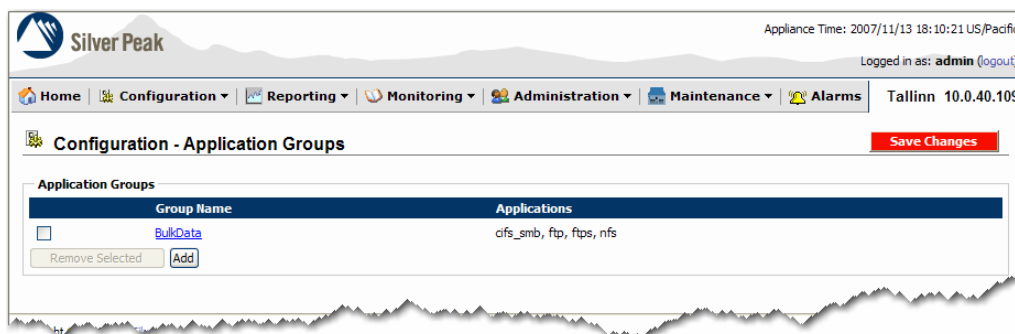
If you try to delete a group that an ACL and/or a policy *is* using, the Appliance Manager tells you which policy is using the application group and tells you that you must change that policy entry:



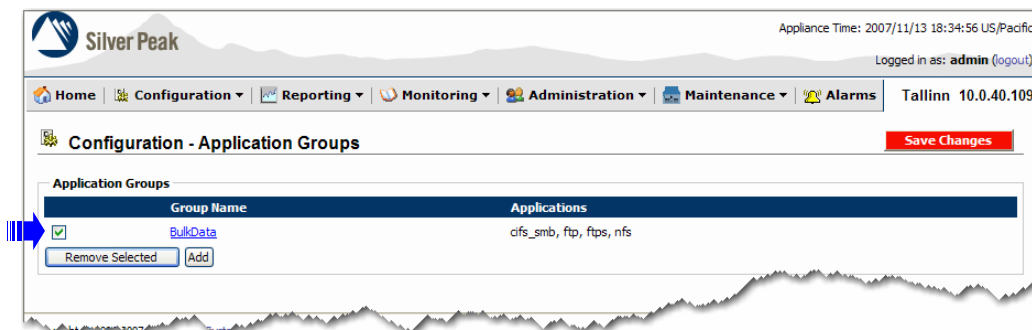
In that case, you'll need to either remove the ACL or edit the ACL to replace the group with another group or application.

Once the application group has no associations, you can delete it from the table on the **Configuration - Application Groups** page.

- 1 In the menu bar, click **Configuration > Application > Groups** to access the **Configuration - Application Groups** page.




- For the entry you want to remove, select its check box. If you want to remove more than one group, select their check boxes. A selected check box displays a green check mark.



WARNING: Clicking on the group name makes the line editable and displays different buttons from those shown here. To back out of that situation, just click **Cancel**.

- Click **Remove Selected**. A confirmation message appears.



- Click **OK**.
If the group were associated with a policy or ACL, Appliance Manager would display that message now.
- Click **Save Changes**.

The next three chapters describe each of the policies — Route, QoS, and Optimization — in terms of their SET actions and how to best utilize them.



Route Policy

This chapter describes the Route Policy.

Because MATCH criteria work the same way across all policies, the discussions focus on the SET actions that are specific to the Route policy. Where applicable, they also provide context relative to the Optimization and QoS policies.

You can manually configure explicit Route Policy entries, or you can choose to use auto-optimization for all your TCP traffic.

In This Chapter

- **Introduction** See page 162.
- **How Auto Optimization Works** See page 163.
- **Where the Route Policy Can Direct Flows** See page 165.
- **Route Policy Page Organization** See page 171.
- **Managing the Route Policy** See page 172.

Introduction

The Route Policy determines where to direct traffic and flows. A **Route Policy** asks:

- What traffic do I want to optimize?
- What traffic do I want to send to the WAN without optimization (pass-through traffic)?
- Do I want to shape all pass-through traffic or not?
- What traffic do I want to drop?
- How do I route traffic if a tunnel goes down?

Auto Optimization simplifies configurations in complex or simple networks by reducing or removing the need to create Route Policy entries. By default, the **Route Policy** auto-optimizes all traffic. Specifically, this means:

- TCP flows are automatically directed to the appropriate tunnel, shaped, and optimized.
- Non-TCP traffic passes through as unoptimized, but shaped to the maximum system bandwidth.
- To optimize non-TCP traffic, you simply add an entry to the Route map for the traffic you want to optimize.

You can, if you choose, modify the default entry's SET action.

As in all policies, the Route Policy consists of entries that pair MATCH criteria with SET actions. Each MATCH criteria delineates a flow. In the Route Policy, the SET actions determine whether an individual flow is ultimately:

- auto-optimized
- directed to a specific tunnel, shaped, and optimized
- processed as shaped, pass-through (unoptimized) traffic
- processed as unshaped, pass-through (unoptimized) traffic, or
- dropped.

How Auto Optimization Works

Auto Optimization begins with the sending of TCP control packets which—in the process of handshaking—determine which tunnel to use as they open the connection.

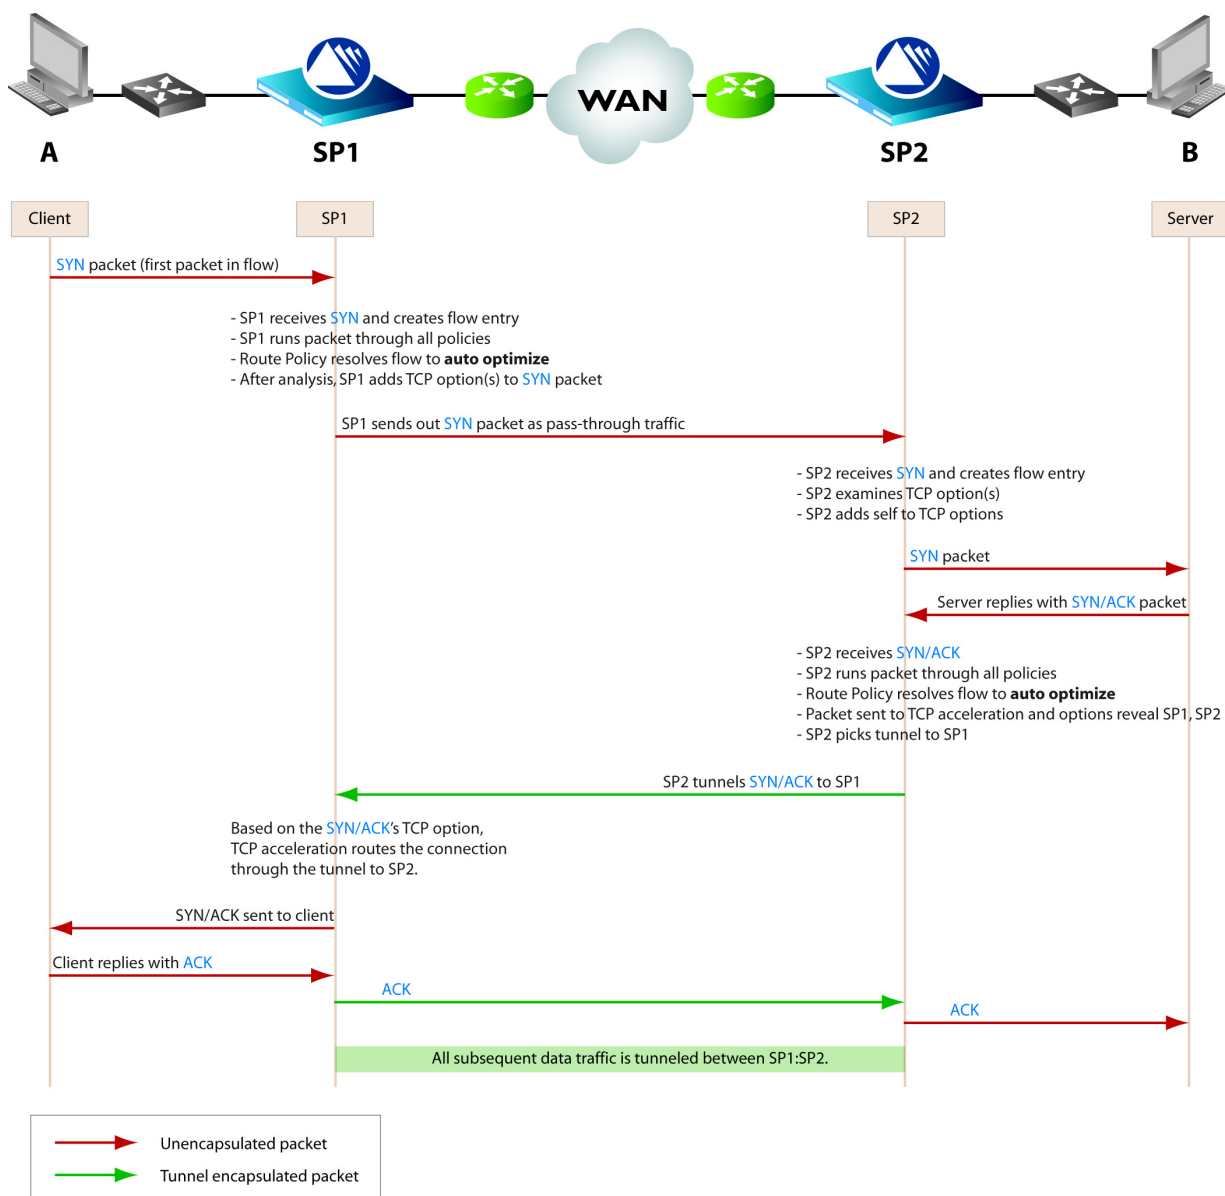
Basic TCP handshaking consists of three ordered steps:

- 1 The client sends a SYN packet to the server, as “hello”.
- 2 The server receives the SYN packet and acknowledges it by sending a SYN/ACK packet.
- 3 The client receives the SYN/ACK packet. The connection is established. The client then sends an ACK packet, along with the data, known as a **TCP flow**.

During this process, the appliances interact with the control packets to set up auto optimization.

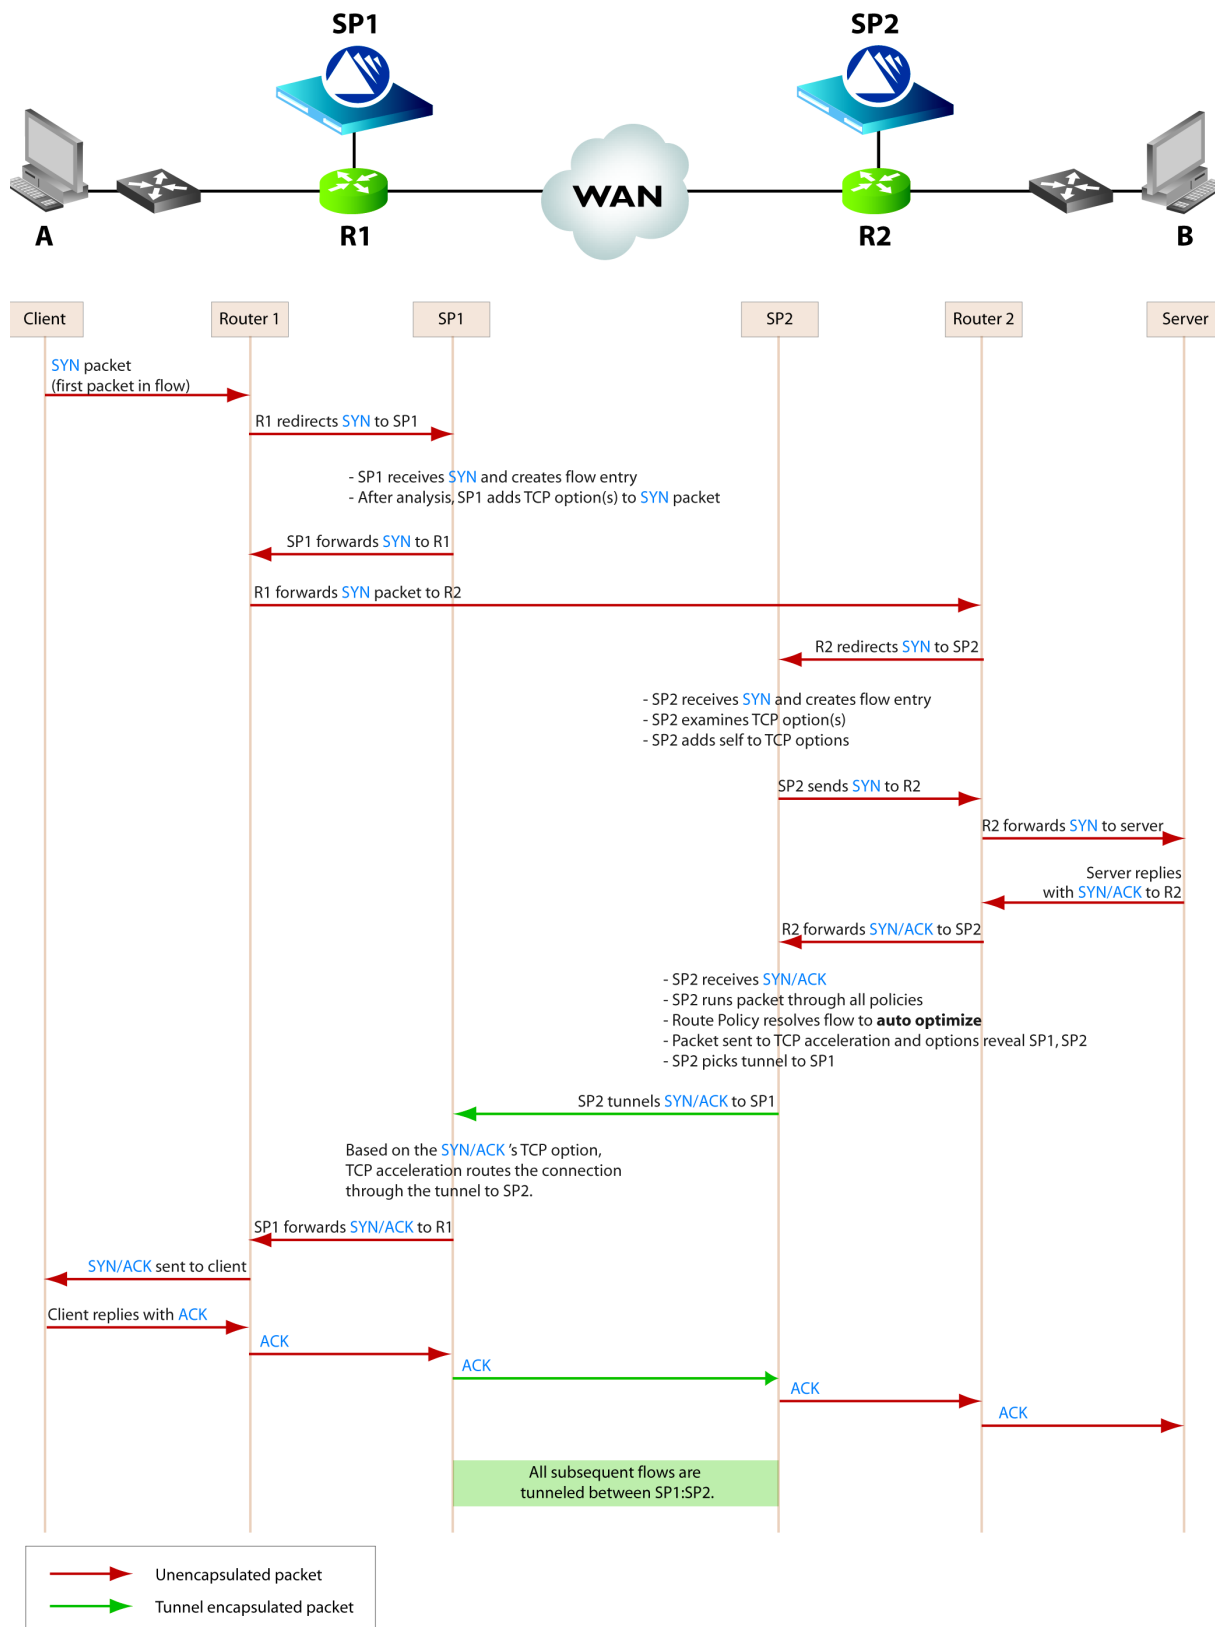
Handshaking for Auto Optimization in Bridge Mode

Beginning at the top and progressing to the bottom, this diagram summarizes the sequence of activities during handshaking in Bridge mode.



Handshaking for Auto Optimization in Router Mode

Beginning at the top and progressing to the bottom, this diagram summarizes the sequence of activities during handshaking in Router mode.



Where the Route Policy Can Direct Flows

The Route Policy's SET actions determine:

- where the appliance directs the traffic, and
- how traffic is treated if a tunnel is down.

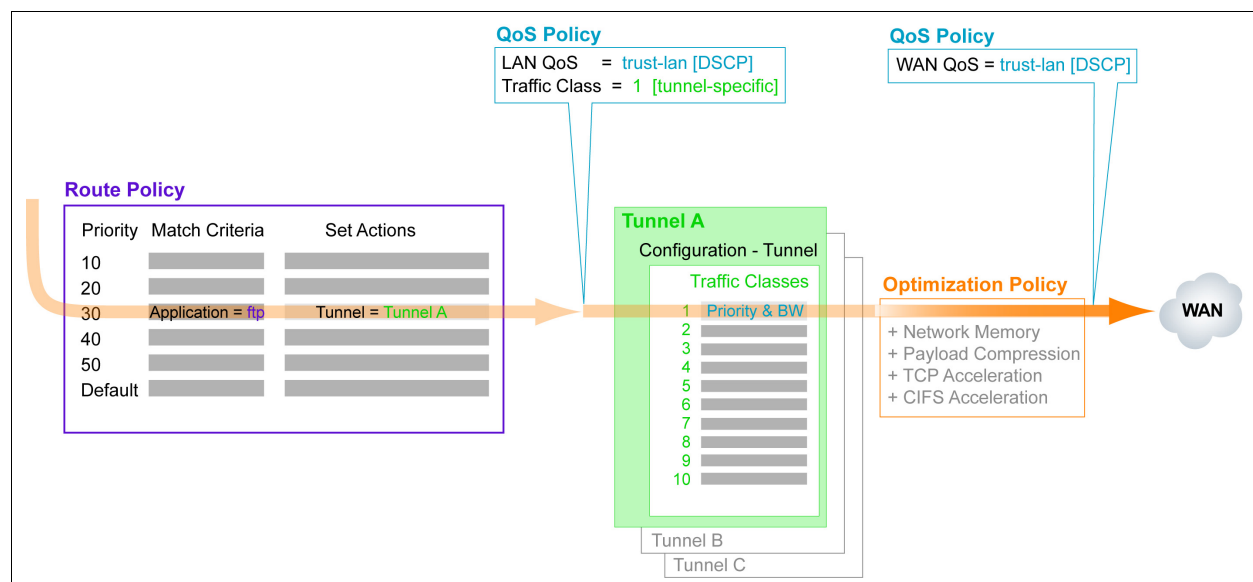
These actions correlate with what you choose for the options in **Tunnel** and **Tunnel Down Action**. The following diagrams illustrate the consequences for each:

- **Flow directed to a tunnel** See page 165.
- **Flow designated as auto-optimized** See page 166.
- **Flow designated as shaped pass-through traffic** See page 167.
- **Flow designated as unshaped pass-through traffic** See page 168.
- **Flow dropped** See page 169.
- **Continue option used in Tunnel Down Action** See page 170.

Flow directed to a tunnel

The most important thing to remember is that **the only way to optimize traffic is to direct flows to tunnels**, either by specifying the tunnel or selecting auto-optimization.

This diagram shows how the appliance processes a flow assigned to a tunnel by the Route Policy. The QoS and Optimization policies are shown only in the interest of providing a broader context for interested users.

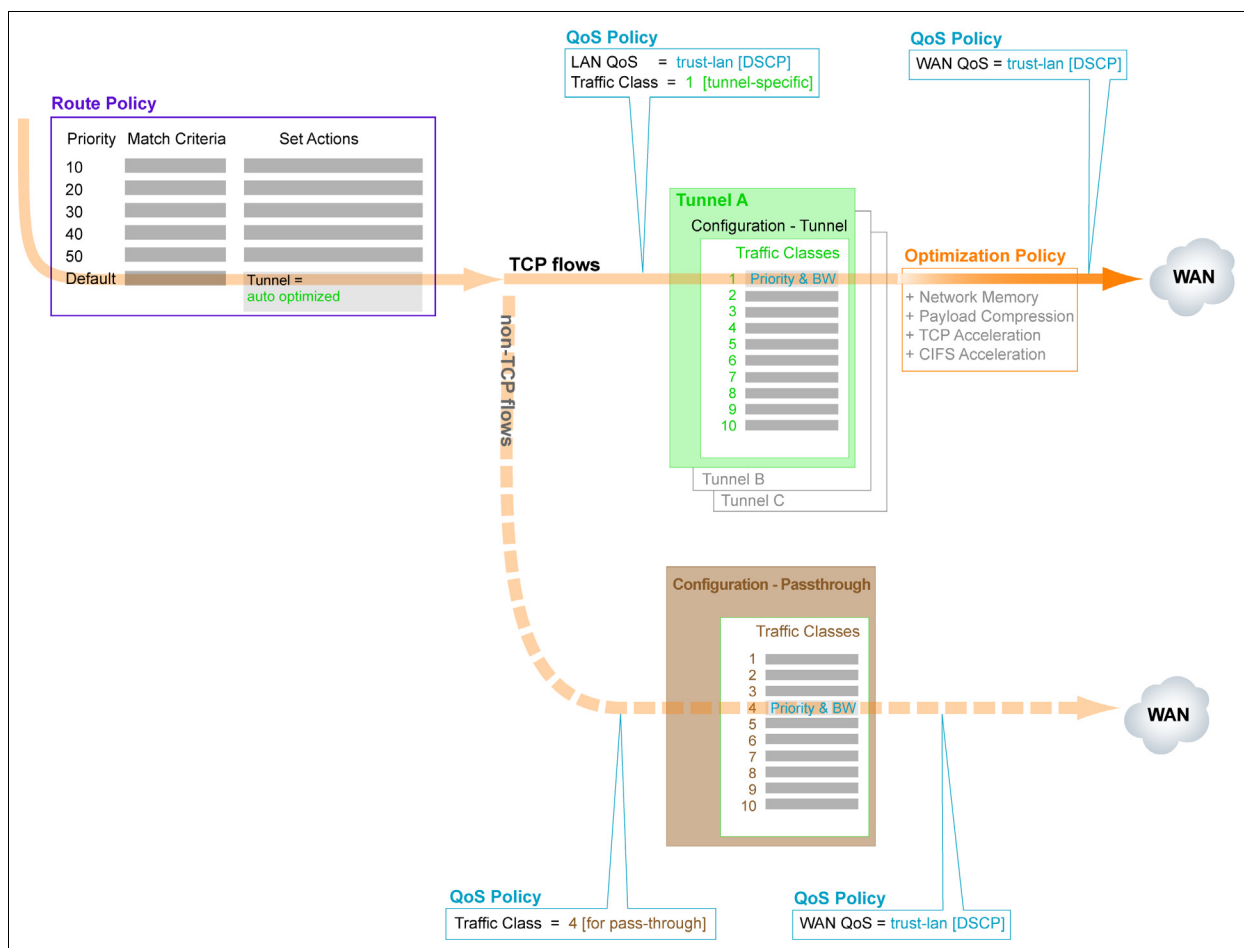


- 1 First, the Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. Entries 10 and 20 don't match the traffic, but Entry 30 does.
- 2 The policy applies the entry's SET actions to the identified flow. In this case, it sends the flow to **Tunnel A**. Once any traffic matches an entry, no subsequent entries are examined.

- 3 Before the flow reaches **Tunnel A**, the QoS Policy checks against its entries and
 - applies the DSCP marking specified for LAN QoS, and
 - tells the flow which of **Tunnel A**'s traffic classes to use. All traffic classes for optimized flows are tunnel-specific. That is, they're part of the tunnel's configuration.
- 4 The appliance places the flow in Traffic Class #1 and passes the flow to the Optimization Policy.
Only flows directed to tunnels are subject to the Optimization Policy.
- 5 After optimization, the QoS Policy applies the rest of the SET action, which is to apply the DSCP markings for the WAN QoS.
- 6 The appliance queues the optimized flow to exit the physical WAN interface.

Flow designated as auto-optimized

When a Route Policy entry has a SET action of *auto optimized* — as is the case with the default entry — the appliance first employs a TCP handshaking procedure to determine the appropriate tunnel (if there's more than one configured) and establish a connection. The appliance then sends TCP flows through the appropriate tunnel, and any non-TCP traffic is sent as pass-through-shaped.

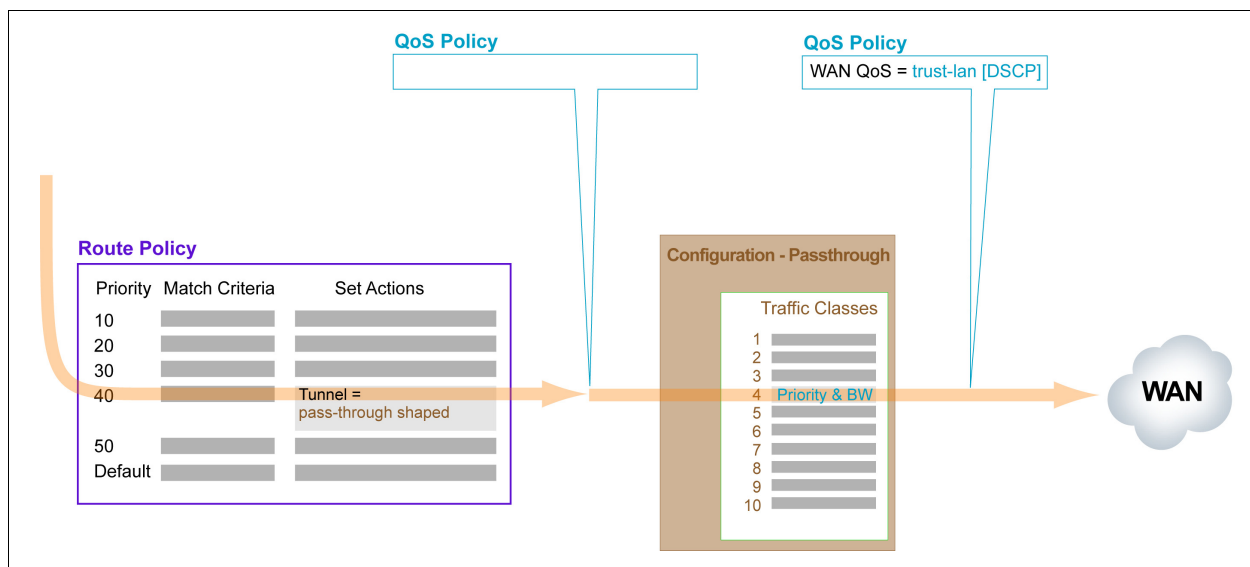


- 1 The Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. No user-configured entries match the traffic, so the Default entry applies.

- 2 The Default entry tells the appliance to process the flow as auto-optimized traffic. In this case, it means it sends TCP flows to **Tunnel A**, and processes non-TCP flows as pass-through-shaped:
 - a Before the TCP flow reaches **Tunnel A**, the QoS Policy checks against its entries and
 - applies the DSCP marking specified for LAN QoS, and
 - tells the flow which of **Tunnel A**'s traffic classes to use. All traffic classes for optimized flows are tunnel-specific. That is, they're part of the tunnel's configuration.
 - b The appliance places the flow in Traffic Class #1 and passes the flow to the Optimization Policy.
Only flows directed to tunnels are subject to the Optimization Policy.
 - c After optimization, the QoS Policy applies the rest of the SET action, which is to apply the DSCP markings for the WAN QoS.
- At the same time, it processes non-TCP flows as shaped, pass-through traffic:
- d Based on the configuration for pass-through traffic, the QoS Policy checks against its entries and
 - ignores the DSCP marking specified for LAN QoS, and
 - tells the flow which traffic class to use. For *shaped*, pass-through traffic, the appliance references the **Configuration – Pass-through** page.
 - e After shaping, the QoS Policy applies the rest of the SET action, which is to apply the DSCP markings for the WAN QoS.
- 3 The appliance queues the flows to exit the physical WAN interface.

Flow designated as shaped pass-through traffic

Flows tagged by the Route Policy as shaped, pass-through traffic follow this path:

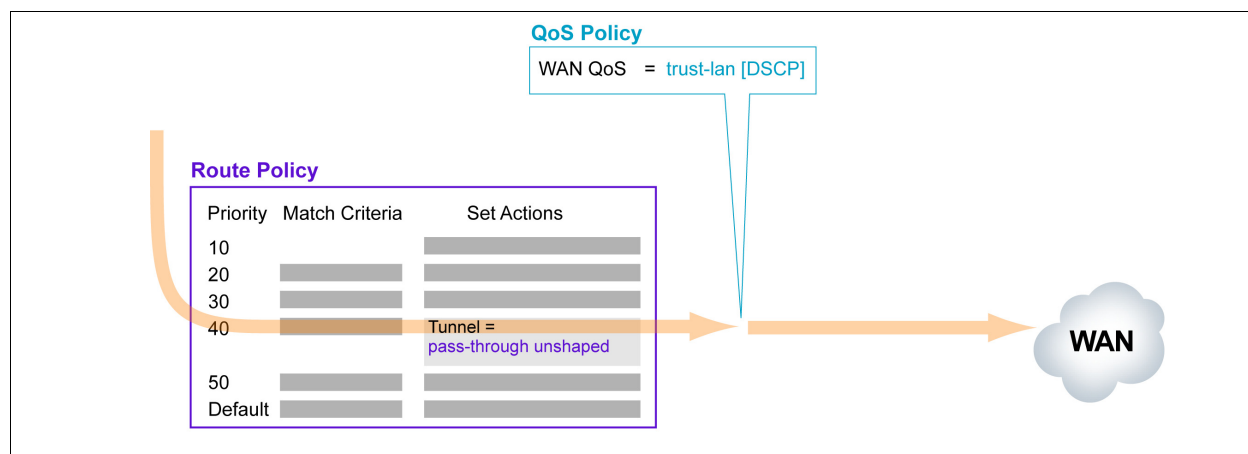


- 1 The Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. No user-configured entries match the traffic, so the Default entry applies.
- 2 The Default entry tells the appliance to process the flow as shaped, pass-through traffic.

- 3 Then, the QoS Policy checks against its entries and
 - ignores the DSCP marking specified for LAN QoS, and
 - tells the flow which traffic class to use. For *shaped*, pass-through traffic, the appliance references the **Configuration – Pass-through** page.
- 4 After shaping, the QoS Policy applies the rest of the SET action, which is to apply the DSCP markings for the WAN QoS.
- 5 The appliance queues the flow to exit the physical WAN interface.

Flow designated as unshaped pass-through traffic

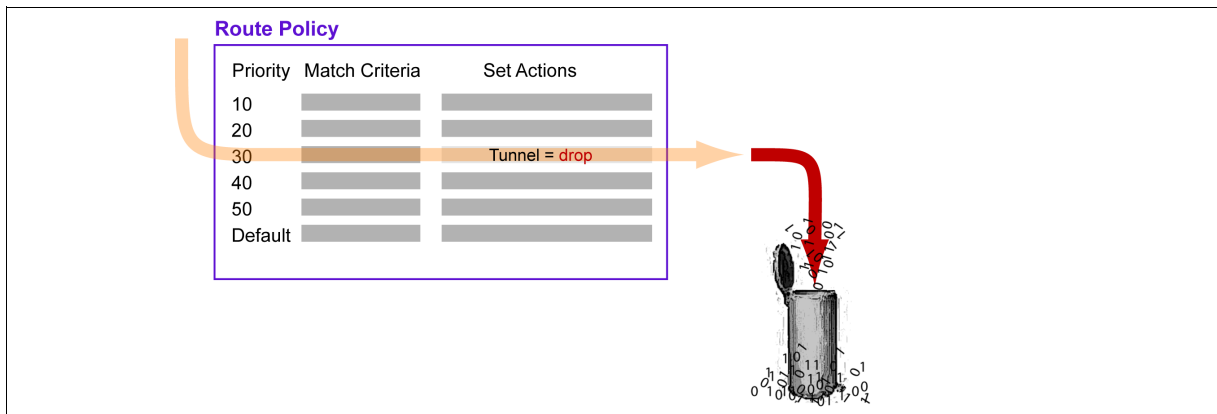
Flows marked by the Route Policy as unshaped, pass-through traffic follow this path:



- 1 The Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. The first three entries don't match the traffic, but Entry 40 does.
- 2 The policy applies the entry's SET actions to the identified flow. In this case, the flow is to be processed as unshaped, pass-through traffic. Once any traffic matches an entry, no subsequent entries are examined.
- 3 Then, the QoS Policy checks against its entries and only applies the DSCP marking specified for WAN QoS.
- 4 The appliance queues the flow to exit the physical WAN interface.

Flow dropped

Flows that have a SET action of **drop** follow this path:



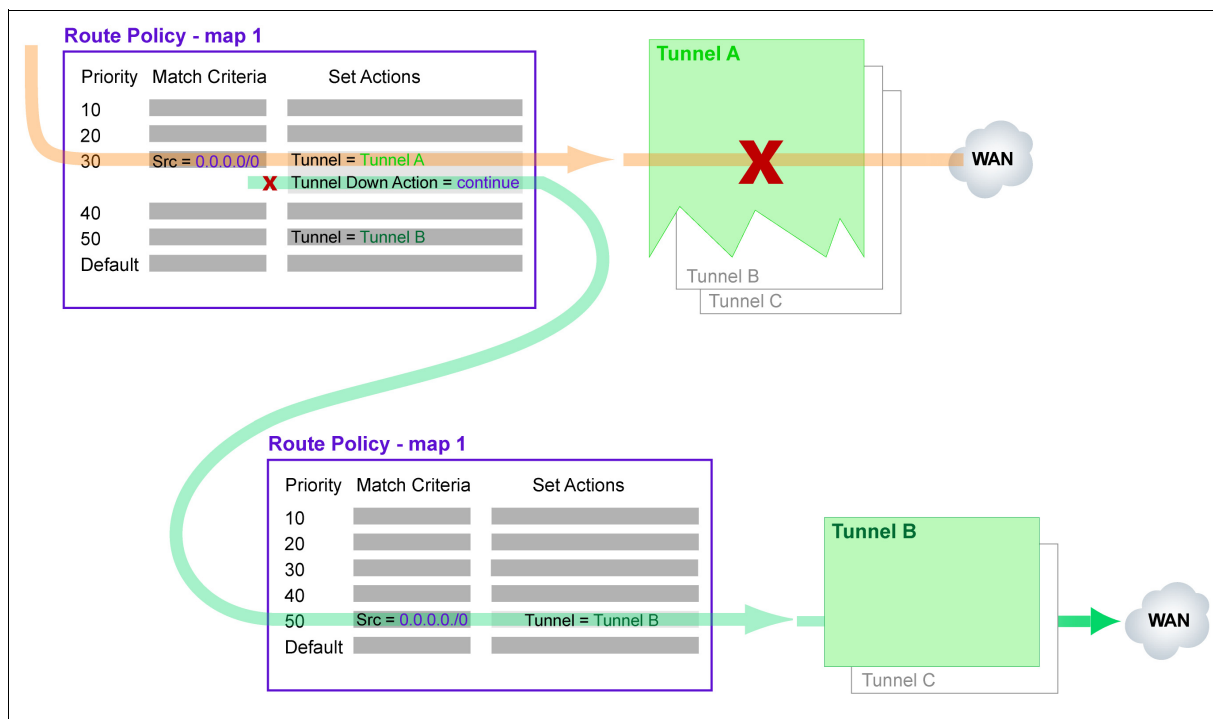
- 1 First, the Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. Entries 10 and 20 don't match the traffic, but Entry 30 does.
- 2 With a SET action of **drop**, the appliance stops all processing on the flow.

Continue option used in Tunnel Down Action

The **Continue** option in the **Tunnel Down Action** field enables the appliance to read ensuing entries in the Route Policy in the event that the tunnel used in a previous entry goes down.

Flows that have a **Tunnel Down** SET action of **Continue** follow this path:

(We've simplified this last diagram, skipping over the sequenced application of Optimization and QoS Policies. To refresh your memory, see *"Flow directed to a tunnel"* on page 165.)



- 1 First, the Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. Entries 10 and 20 don't match the traffic, but Entry 30 does.
- 2 The policy applies the entry's SET actions to the identified flow. In this case, it sends the flow to **Tunnel A**. Once any traffic matches an entry, no subsequent entries are examined.
- 3 **Tunnel A** goes down, and the Route Policy refers back to the policy entry's **Tunnel Down Action**. The action prescribed is to **continue** to the next applicable MATCH criteria, which is Entry 50, putting all traffic into **Tunnel B**.

This configuration provides redundancy for high availability environments:

- If **Tunnel A** is subsequently restored, the Route Policy directs new flows matching Entry 30 to **Tunnel A**.
- Flows that were continued from Entry 30 to Entry 50 (and **Tunnel B**) persist until complete.

Route Policy Page Organization

The following shows the SET actions for the appliance, **SP1**:

This Route Policy has three maps — **gms_RouteMap**, **gms_RouteMap_alt**, and **map1**.

By definition, the active map is the **policy** and always displays at the top of the list when there is more than one map. Here, it's **gms_RouteMap**.

To switch to another Route Policy, select from the drop-down menu and click **Activate**. Any change governs all new flows.

Configuration - Route Policy

Map Name: **gms_RouteMap** **Activate**

Total entries: 4 Max entries: 600

gms_RouteMap (active) - Number of entries: 4

Match Criteria									Set Actions	
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Tunnel	Tunnel Down Action
803		ip	7.7.7.0/24	7.7.7.0/24	any		any	any.any	[pass-through-unshaped]	
804		ip	8.8.8.0/24	8.8.8.0/24	any		any	any.any	[pass-through-unshaped]	
805		ip	0.0.0.0/0	2.2.0.0/16	any		any	any.any	gms_SP1_SP4	continue
806		ip	0.0.0.0/0	2.2.0.0/16	any		any	any.any	gms_SP1_SP2	continue
default									[pass-through-unshaped]	

Remove Selected Add Remove Route Map

gms_RouteMap_alt - Number of entries: 0

rm - Number of entries: 0

The last column is only accessible if the **Tunnel** entry is a **specific tunnel**.

The following options are available when configuring the **Tunnel**:

- **auto optimized**
- the name of any tunnel from the **Configuration - Tunnels** page
- **pass-through [shaped]**
- **pass-through-unshaped**, and
- **drop**

Set Actions	
Tunnel	Tunnel Down Action
[pass-through-unshaped]	
[pass-through-unshaped]	
gms_SP1_SP4	continue
gms_SP1_SP2	continue
[pass-through-unshaped]	

No matter how many entries you add, the default entry in any policy (or map) is *always* the last line. In the case of a Route Policy, the default is **auto optimized**.

Tunnel Down Action has the following options:

- **pass-through [shaped]**
- **pass-through-unshaped**
- **drop**, and
- **continue**

Managing the Route Policy

This section discusses procedures for managing the Route Policy. It includes:

- **Adding an Entry to a Map** See page 172.
- **Editing an Entry** See page 174.
- **Deleting an Entry** See page 175.
- **Adding a New Route Map** See page 175.
- **Deleting a Map** See page 178.
- **Activating a New Policy** See page 178.

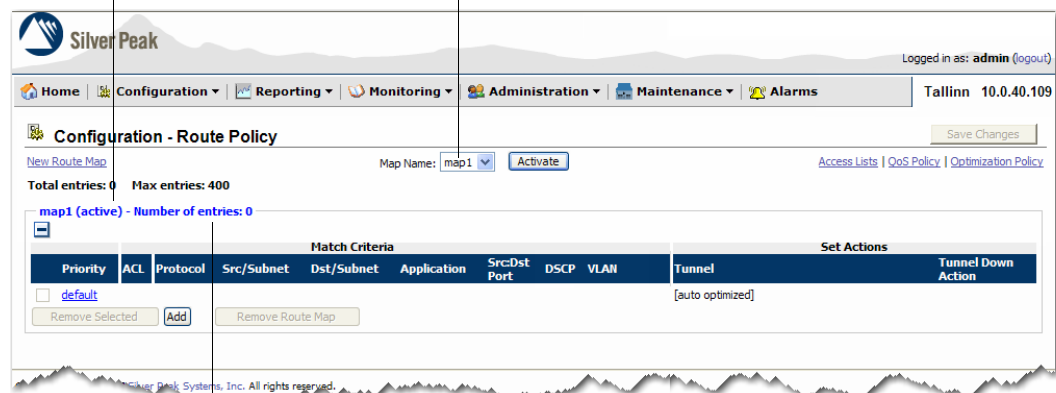
Adding an Entry to a Map

This section describes how to add an entry to a route map.

♦ To add an entry to a map

- 1 From the **Configuration** menu, select **Route Policy**. The **Configuration – Route Policy** page appears. If this is the first time you're accessing the page, only the default entry appears.

The default Route Policy is always **map1**...
...and it's active.



This refers only to entries made by a user.
The default entry isn't part of the count.

The Route Policy has a default entry that isn't numbered. It always occupies the last priority position. In reality, its assigned value is **65535**.

You can modify the default entry's Action. The options are:

- auto-optimization
- pass-through [shaped]
- pass-through-unshaped
- drop

If traffic matches no user-configured entries, then the default entry:

- auto-optimizes all TCP traffic, and
- sends non-TCP traffic to the WAN as shaped pass-through traffic.

- Click **Add**. A new entry appears, with editable fields.

If you need to review the list of existing ACLs, click **Access Lists**. The **Configuration – Access Lists** page displays.

After you've reviewed the entries, you can return to this page by clicking **Route Policy** in the same location on that page.

To use an existing ACL, select its name from the **ACL** column's drop-down menu. Once you do, the other fields in the **Match Criteria** area are inaccessible.

To complete the tuple elements individually, complete these fields.

As described in Chapter 6, “Theory of Operations,” do one of the following:

- Complete the individual fields belonging to the **Match Criteria**, or
- From the **ACL** column, select the name of an existing Access Control List.

- Make your selection from the **Tunnel** field.

This column is only accessible if the **Tunnel** entry is a **specific tunnel**.

Tunnel names display without square brackets

- If you select a tunnel, then you must select an option from the **Tunnel Down Action** field. The available options include: **pass-through** [shaped], **pass-through-unshaped**, and **drop**.
- Click **Apply**.
- Click **Save Changes**.

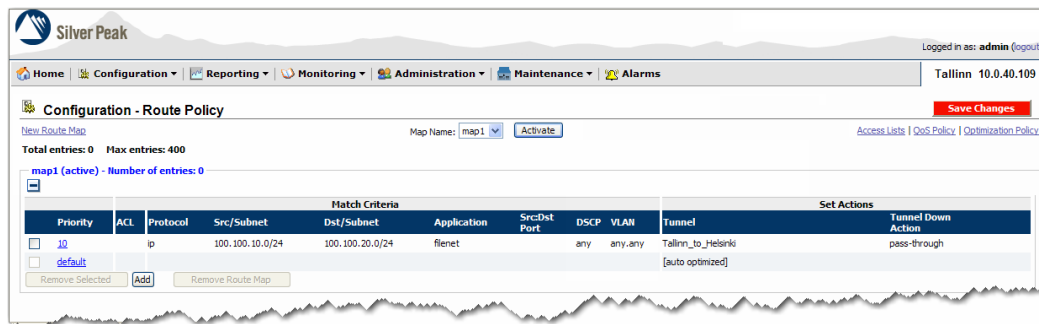
Editing an Entry

You can edit an active entry at any time. Changes to a route map entry affect new connections only.

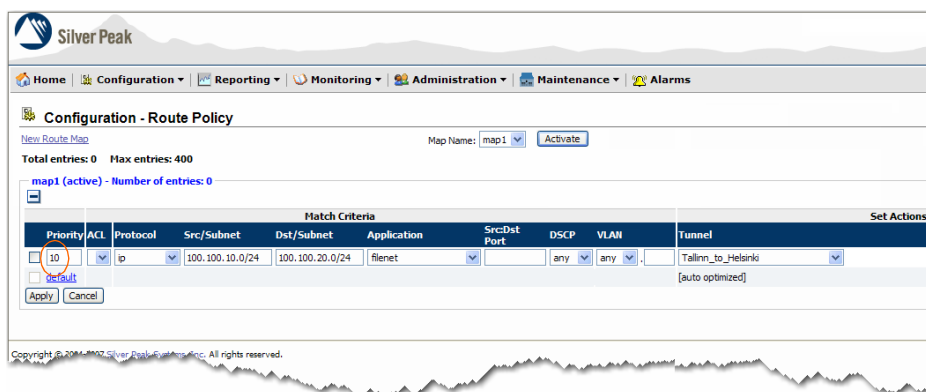
If you've added additional route maps, you can also edit their entries.

♦ To edit an entry

- 1 From the **Configuration** menu, select **Route Policy**. The **Configuration – Route Policy** page appears.



- 2 To edit an entry, click that entry's **Priority** number. The fields become editable.



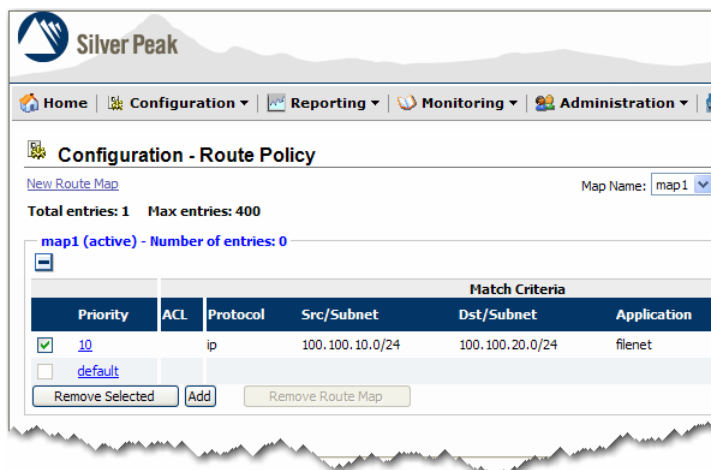
- 3 Make the desired changes.
- 4 Click **Apply**.
- 5 Click **Save Changes**.

Deleting an Entry

The procedure for deleting an entry is the same across all maps, lists, and policies.

♦ To delete an entry

- 1 In the leftmost column of the policy, click the check box(es) for the entries you want to delete.



A green check displays inside, and the **Remove Selected** button becomes accessible.

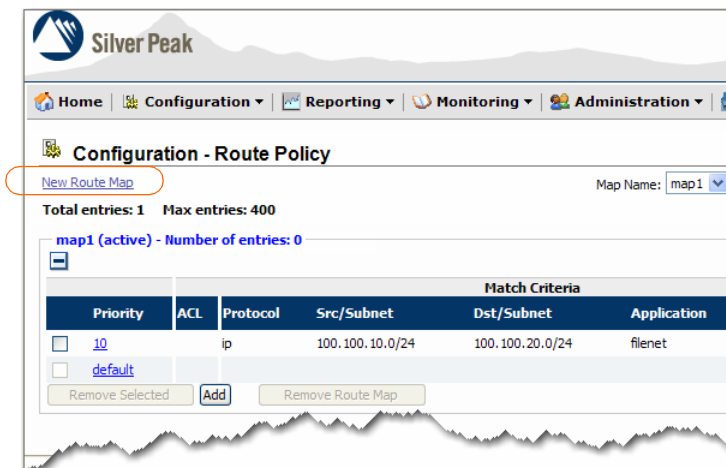
- 2 Click **Remove Selected**.
- 3 Click **Save Changes**.

Adding a New Route Map

You can always add another Route map, whether or not you choose to activate it immediately.

♦ To add a new map

- 1 In the **Configuration – Route Policy** page, click **New Route Map**.



The **new** section displays above the active map.

The screenshot shows the 'Configuration - Route Policy' interface. At the top, there is a 'New Route Map' link and a 'Map Name' dropdown set to 'map1' with an 'Activate' button. Below this, it states 'Total entries: 0' and 'Max entries: 400'. The 'new' section is active, showing a 'Route Map Name' field with an 'Apply' and 'Cancel' button. Below the 'new' section, the 'map1 (active) - Number of entries: 0' section is shown. It contains a table with the following data:

Match Criteria					
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application
<input type="checkbox"/> 10		ip	100.100.10.0/24	100.100.20.0/24	filenet
<input type="checkbox"/> default					

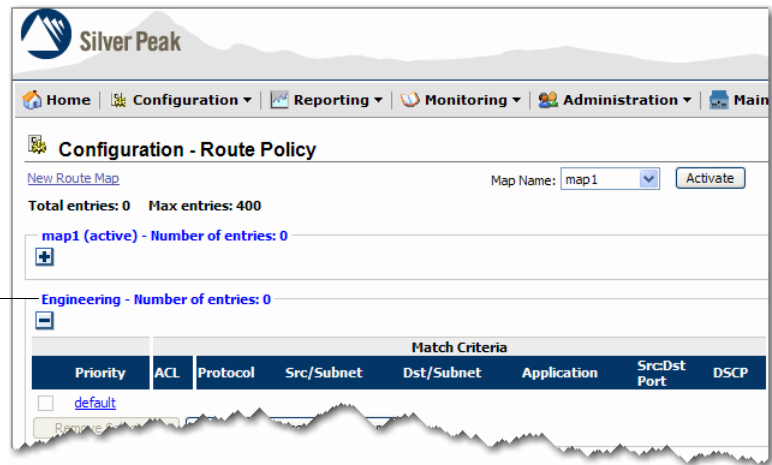
Below the table are buttons for 'Remove Selected', 'Add', and 'Remove Route Map'.

- 2 In the **Route Map Name** field, enter a new map name.

The screenshot shows the 'Configuration - Route Policy' interface. The 'Route Map Name' field in the 'new' section is now filled with the text 'Engineering'. The 'map1 (active) - Number of entries: 0' section remains the same as in the previous screenshot.

3 Click **Apply**.

The new map, **Engineering**, displays in expanded view below the active map (which is the policy). Meanwhile, all other maps are collapsed.




Each time you create a new Route map, the **default** entry is always the same. It auto-optimizes all TCP traffic and then sends the rest through as pass-through shaped.

- 4 Add new entries, as necessary.
- 5 Click **Save Changes**.

Deleting a Map


You can **only delete** an **inactive** map.

♦ To delete a map

- 1 Under the map's name, click the  to expand the inactive map you want to delete.



- 2 Click **Remove Route Map**. A window appears, requesting confirmation.




- 3 Click **OK**.
- 4 Click **Save Changes**.

Activating a New Policy

Activating a new Route map deactivates the old one. All new traffic matches against the new Route Policy.

♦ To activate a new policy

- 1 From the **Map Names** field, select the map you want to activate.



- 2 Click **Activate**. The old map deactivates and the new one activates, beginning with all new flows.
- 3 Click **Save Changes**.



Bandwidth Management & QoS Policy

This chapter describes the QoS Policy's SET actions. It also explains how to configure traffic classes for optimized and pass-through traffic, along with providing best practices guidelines for effectively managing bandwidth.

In This Chapter

- **Overview** See page 180.
- **How the QoS Policy Affects Flows** See page 180.
- **Best Practices for Bandwidth Management** See page 184.
- **Configuring Maximum System Bandwidth** See page 188.
- **How Tunnel Auto BW Works** See page 189.
- **Configuring Pass-Through Traffic Bandwidths** See page 190.
- **Configuring Traffic Classes** See page 191.
- **Handling and Marking Packets** See page 194.
- **QoS Policy Page Organization** See page 201.
- **Managing the QoS Policy** See page 202.

Overview

Effective use of QoS ensures that bandwidth is adequately distributed when a mix of traffic, with varying levels of priority, must be delivered over a shared limited resource.

In a well-designed network, QoS will be managed at every potential bottleneck point. It is particularly important to implement QoS in the WAN acceleration appliance for the following reasons:

- It can offload the router.
- It's the only element that collects real-time metrics — such as packet loss and delay — for both a pre- and post-optimization view of the traffic.

The main purpose of the Silver Peak QoS feature portfolio is to optimize performance in the presence of network impairments and, in the event that demand exceeds available bandwidth, to give user-defined preferential treatment to selected flows, as defined by MATCH criteria in the **QoS Policy**.

A **QoS policy** asks:

- How do I want to prioritize traffic via the DSCP markings?
- How are the traffic flows to be shaped via traffic classes?
- For a flow directed to Tunnel_A, what Tunnel_A traffic class should it be assigned?
- For a flow designated as shaped, pass-through traffic, what pass-through traffic classes should it be assigned?
- How should the DSCP markings be treated? Trust the incoming LAN or re-mark for the WAN?

The default QoS Policy honors incoming DSCP tags and sets all flows to Traffic Class 1. For the majority of users, the need to adjust this will be a “corner case”.

How the QoS Policy Affects Flows

As in all policies, the QoS policy consists of entries that pair MATCH criteria with SET actions. Each MATCH criteria delineates a flow. In the **QoS Policy**, the SET actions determine how flows are queued and marked.

- If the Route Policy directs the flow to a **tunnel** for optimization, then this number refers to the specific traffic class in that tunnel, found on the **Configuration - Tunnels** page.
- If the Route Policy designates the flow as **pass-through shaped** traffic, then this number refers to the traffic class found on the **Configuration - Pass-through** page.

The default is to use **Traffic Class 1**.

This parameter has no bearing on unshaped pass-through traffic.

Traffic Class	LAN QoS	WAN QoS
1	trust-lan	trust-lan

Below the table are three dropdown menus, each currently set to '1'.

LAN QoS and **WAN QoS** refer to the DSCP markings.

For more information, see [“Handling and Marking Packets” on page 194](#).

The QoS Policy's SET actions determine two things:

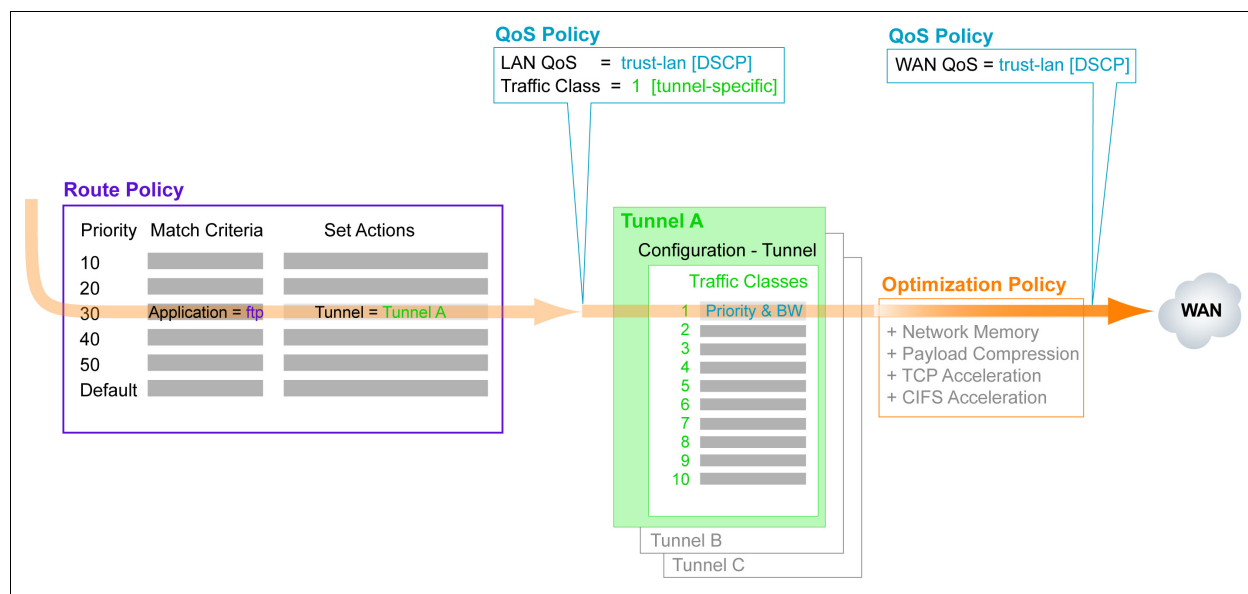
- what traffic class a shaped flow — whether optimized or pass-through — is assigned
- how to handle DSCP markings for all flows leaving the appliance's WAN interface, whether the marking is for over-the-WAN or for the LAN on the remote side.

The following diagrams illustrate the consequences for each:

- **Flow directed to a tunnel** See page 181.
- **Flow designated as pass-through shaped traffic** See page 182.
- **Flow designated as unshaped pass-through traffic** See page 183.

Flow directed to a tunnel

This diagram shows how the appliance applies QoS to a flow that's been directed to a tunnel.



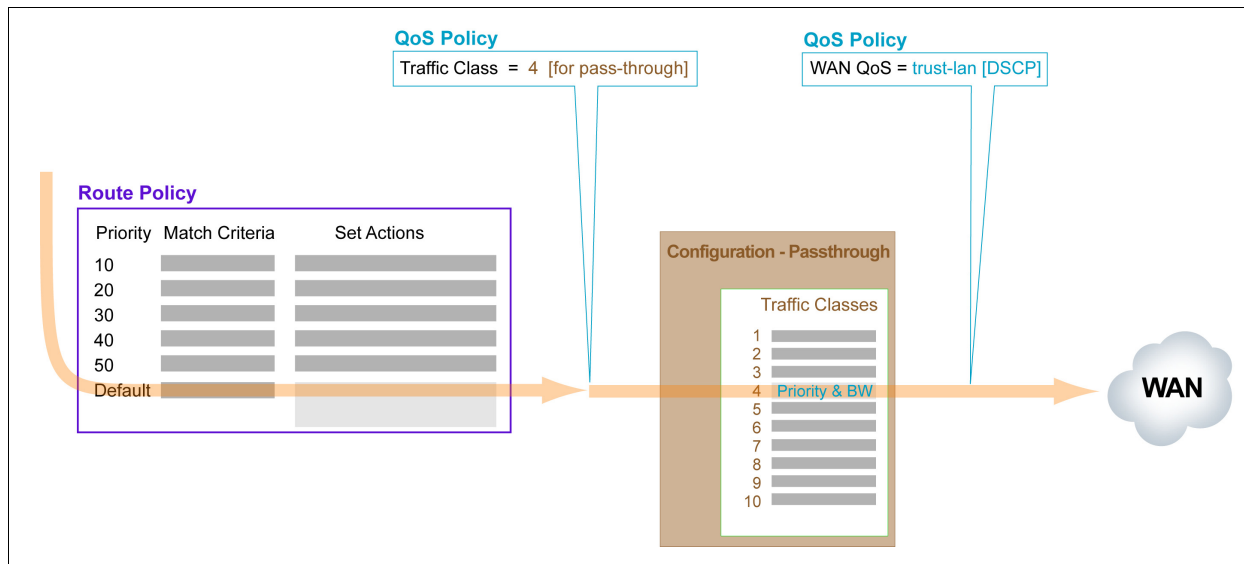
- 1 First, the Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. Entries 10 and 20 don't match the traffic, but Entry 30 does.
- 2 The policy applies the entry's SET actions to the identified flow. In this case, it sends the flow to **Tunnel A**. Once any traffic matches an entry, no subsequent entries are examined.
- 3 Before the flow reaches **Tunnel A**, the QoS Policy checks against its entries and
 - applies the DSCP marking specified for LAN QoS, and
 - tells the flow which of **Tunnel A**'s traffic classes to use. All traffic classes for optimized flows are tunnel-specific. That is, they're part of the tunnel's configuration.
- 4 The appliance places the flow in Traffic Class #1 and passes the flow to the Optimization Policy.
- 5 After optimization, the QoS Policy applies the DSCP markings for the WAN QoS.
- 6 The optimized flow exits the physical WAN interface.



Handling of DSCP markings is further explained in *"Applying DSCP Markings to Optimized Traffic"* on page 194.

Flow designated as pass-through shaped traffic

Flows tagged by the Route Policy as pass-through shaped traffic follow this path:



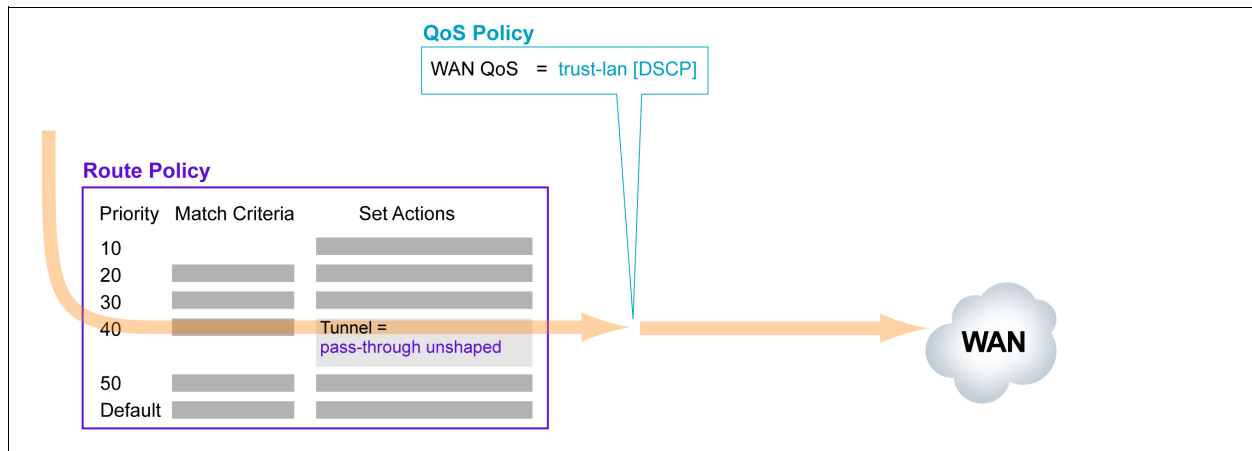
- 1 The Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. No user-configured entries match the traffic, so the Default entry applies.
- 2 The Default entry tells the appliance to process the flow as pass-through shaped traffic.
- 3 Then, the QoS Policy checks against its entries and
 - ignores the DSCP marking specified for LAN QoS (because packets are not encapsulated), and
 - tells the flow which traffic class to use. For *shaped*, pass-through traffic, the appliance references the **Configuration – Pass-through** page.
- 4 After shaping, the QoS Policy applies the rest of the SET action, which is to apply the DSCP markings for the WAN QoS.
- 5 The flow exits the physical WAN interface.



Handling of DSCP markings is further explained in *“Applying DSCP Markings to Shaped and Unshaped Pass-through Traffic”* on page 197.

Flow designated as unshaped pass-through traffic

Flows marked by the Route Policy as unshaped, pass-through traffic follow this path:



- 1 The Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. The first three entries don't match the traffic, but Entry 40 does.
- 2 The policy applies the entry's SET actions to the identified flow. In this case, the flow is to be processed as unshaped, pass-through traffic. Once any traffic matches an entry, no subsequent entries are examined.
- 3 Because the traffic is set to pass-through unshaped, it is **not** encapsulated. The QoS Policy checks against its entries and only applies the DSCP marking specified for WAN QoS.
- 4 The flow exits the physical WAN interface.



Handling of DSCP markings is further explained in *"Applying DSCP Markings to Shaped and Unshaped Pass-through Traffic"* on page 197.

Best Practices for Bandwidth Management

Congestion is unlikely on either of the LAN segments to which the Silver Peak device connects directly, since these are typically operating at 100Mbps or 1000Mbps.

In a typical deployment, congestion is most likely to arise at the near-end WAN interface.

With wise bandwidth management and QoS, the Silver Peak appliance can guarantee shaping and prioritization for all traffic. For smooth network operation, it's wisest to consider your overall bandwidth allocation in advance and then to revisit it each time you add, edit, or remove a tunnel.

This section describes the following:

- **Summary of Bandwidth Assessment and Management Tasks** See page 184.
- **Guidelines for Configuring Minimum and Maximum Bandwidth Values** See page 185.
- **Which Appliance Manager Pages to Use** See page 187.

Summary of Bandwidth Assessment and Management Tasks

The following table summarizes the tasks when configuring multiple tunnels for an appliance and/or more than one traffic class per entity.

	Task	Notes	For detailed instructions, see...
1	Configure the maximum system bandwidth, based on the bandwidth of the WAN link.	Because of where network congestion typically occurs, you want to ensure that the appliance doesn't deliver more than the WAN can manage.	"Configuring Maximum System Bandwidth" on page 188.
2	Let the appliance negotiate tunnel maximum bandwidth(s)	When Auto BW is active (as it is by default), the appliance negotiates maximum bandwidth for each tunnel.	"How Tunnel Auto BW Works" on page 189.
3	Configure the tunnel minimum bandwidth		"Which Appliance Manager Pages to Use" on page 187.
4	Configure traffic classes for tunnels.	Each tunnel's default traffic class is number 1. You can configure up to 10 traffic classes per tunnel.	"Configuring Traffic Classes" on page 191.
5	Set the configuration for pass-through shaped traffic	Configure: <ul style="list-style-type: none"> • minimum and maximum bandwidth values for the traffic as a whole • traffic classes 	"Configuring Pass-Through Traffic Bandwidths" on page 190. "Configuring Traffic Classes" on page 191.
6	Review your configuration	Make sure that you haven't over- or undersubscribed the link.	"Guidelines for Configuring Minimum and Maximum Bandwidth Values" on page 185.

Guidelines for Configuring Minimum and Maximum Bandwidth Values

This sample worksheet provides an overall picture of the items and relationships you need to consider.

System Bandwidth	1500	
Also known as the WAN Bandwidth . Set this to the value at router's WAN interface.		
	MIN	MAX
Tunnel #1	500 *	1500
Tunnel #2	500 *	1500
Traffic Class #1	300	1500
Traffic Class #2	150	1500
Traffic Class #3	50	1500
Pass-through - shaped	500 *	1500
Traffic Class #1	300	1500
Traffic Class #2	200	1500

Appropriately configuring the **Minimum Bandwidth** values is the **most critical aspect**.
See "Rules of Thumb" on page 186.

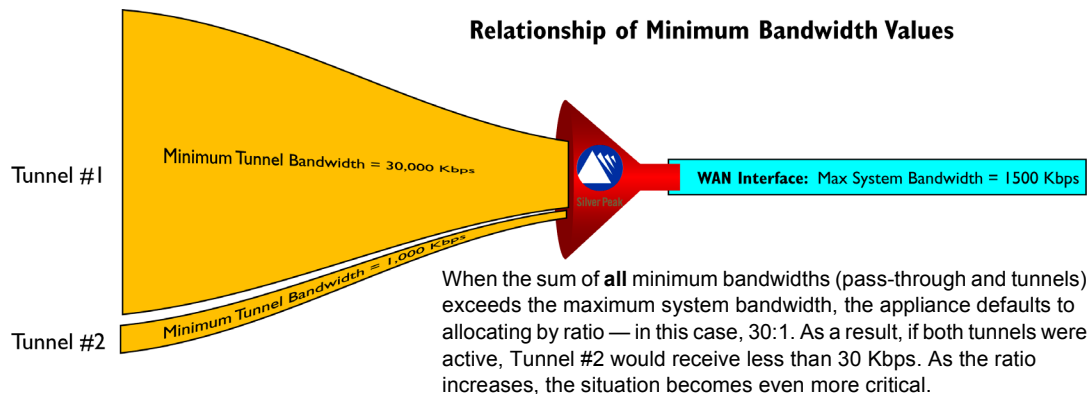
* For minimums, the sum of the component traffic classes should never exceed its parent's configured value.

Guidelines for Maximum Bandwidth Values

- All tunnels and pass-through traffic can have a Max(imum) Bandwidth equal to the WAN Bandwidth.
- For all tunnels and pass-through traffic, each of their traffic classes can have a max(imum) bandwidth equal to the WAN Bandwidth.

Guidelines for Minimum Bandwidth Values

If you oversubscribe the system bandwidth, the appliance defaults to allotting minimum bandwidths based on ratio. In that case, you won't get your expected minimum. In fact, if the ratio is too large, the smaller tunnel may not get any bandwidth at all.



Rules of Thumb

- The minimum pass-through bandwidth plus the sum of all minimum tunnel bandwidths should not exceed the value of the **WAN Bandwidth**.
- Within a tunnel, the sum of traffic class minimum bandwidths should not exceed the tunnel's minimum bandwidth.
- For (shaped) pass-through traffic, the sum of traffic class minimum bandwidths should not exceed the pass-through minimum bandwidth.
- For traffic classes, minimum bandwidth allocation is based on Priority. This becomes critical when you oversubscribe.

Following this sequence prevents bottlenecks and helps keep traffic flowing smoothly to/from the WAN.



Tip When you set a traffic class Minimum Bandwidth to zero, you are explicitly not guaranteeing any bandwidth for that class.

By default, a tunnel's Minimum Bandwidth is set to 32 kbps.

Which Appliance Manager Pages to Use

The worksheet information originates from three **Configuration** menu pages. Highlighted areas show where the necessary data is, positionally.

Go to the **Configuration - System** page to enter the bandwidth at the router's WAN interface in the **WAN Bandwidth** field.

As a best practice, this ensures that the appliance doesn't contribute to congestion.

A System Bandwidth	1500
Also known as the WAN Bandwidth . Set this to the value at router's WAN interface.	
B Tunnel #1	500 *
B Tunnel #2	500 *
Traffic Class #1	300
Traffic Class #2	150
Traffic Class #3	50
C Pass-through - shaped	500 *
Traffic Class #1	300
Traffic Class #2	200

* For minimums, the sum of the component traffic classes should never exceed its parent's configured value.

On the **Configuration - Tunnels** page, click on a tunnel name to see its details:

- The **Setting** tab displays the tunnel Min and Max Bandwidths.
- The **Traffic Class** tab [shown here] has hyperlinks to individual traffic classes for their min/max bandwidth values.

ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	10	1,000,000	0	500

ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	10	1,000,000	0	500

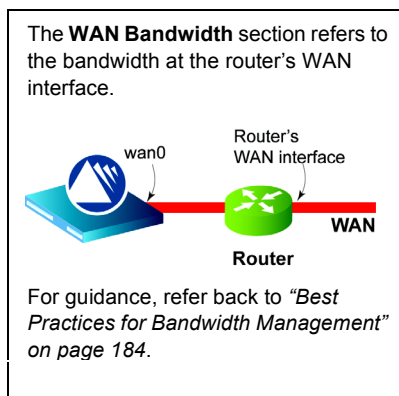
The **Configuration - Pass-through** [shaped] page is home to the overall Min and Max Pass-through values, as well as the traffic class values.

Configuring Maximum System Bandwidth

To make sure that the appliance doesn't send more traffic than the router's WAN interface can manage, enter the router's **WAN Bandwidth** in the **Configuration - System** page.

2-Port Configurations

When using **wan0** and **lan0**, make sure that this value is **not** greater than the WAN interface. In other words, don't oversubscribe.



The default value varies with appliance model. To maintain the best traffic flow, refer to the section, "How Tunnel Auto BW Works" on page 189.

4-Port Configurations

When using a 4-port configuration (that is, with **lan0**, **lan1**, **wan0**, and **wan1**) with two WAN next-hops, use these rules of thumb:

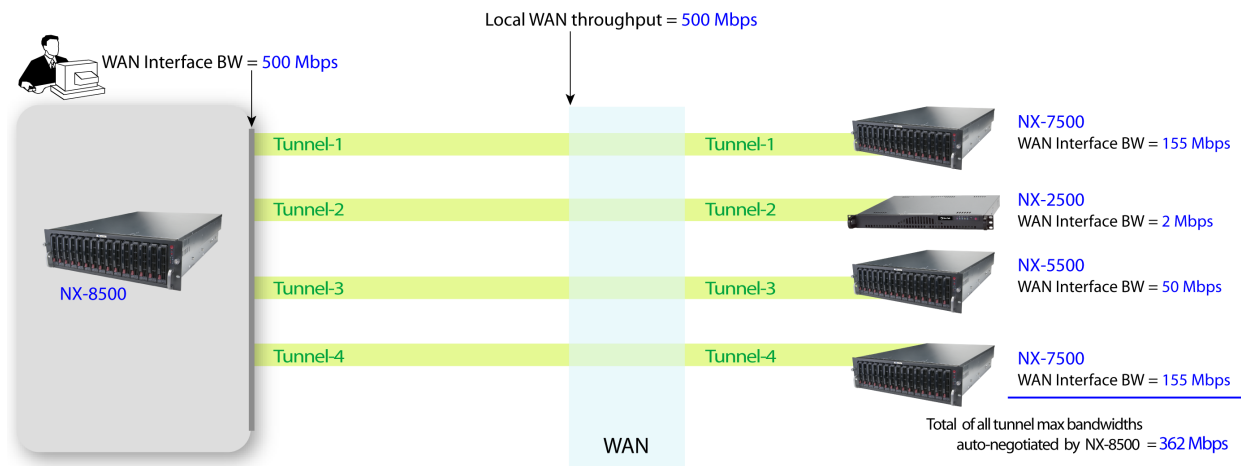
- If the ISPs are configured **Active/Active**, then use the **sum** of the two routers' WAN bandwidths.
- If the ISPs are configured **Active/Standby**, then use the **larger** of the two routers' WAN bandwidth values.

How Tunnel Auto BW Works

Each model of NX Series appliance has a specific maximum system bandwidth. That is, the amount of bandwidth it can support for optimized traffic at the WAN interface. With all features enabled, the amount of WAN transmit traffic that each appliance optimally manages is as follows:

NX Series Appliance	WAN Capacity (all features) (all options enabled)
NX-2500	2 Mbps
NX-1700 / NX-2600	4 Mbps
NX-2610	8 Mbps
NX-2700 / NX-3500	10 Mbps
NX-3600 / NX-3690 / NX-3700	20 Mbps
NX-5500 / NX-5504 / NX-5600 / NX-5700	50 Mbps
NX-7500 / NX-7504 / NX-7600 / NX-7700	155 Mbps
NX-8504 / NX-8600	500 Mbps
NX-8700	622 Mbps
NX-9610 / NX-9700	1 Gbps

By default, all tunnels are set to automatically negotiate tunnel bandwidth to the lowest common value. The following illustrations show this negotiation from the perspective of an NX-8500 with multiple tunnels. The maximum values assume that all options are enabled.



After negotiating bandwidth for all four tunnels, **138 Mbps** (500 minus 362) are left over for shaped pass-through traffic.

Configuring Pass-Through Traffic Bandwidths

On the **Configuration - Pass-through** page, you can configure QoS **Bandwidth** and **Queuing**.

♦ **To configure bandwidth limits for pass-through traffic**

- 1 Access the **Configuration - Pass-through** page.
- 2 Under **Traffic Bandwidth**, configure the **Min(imum) Bandwidth** and **Max(imum) Bandwidth**.

Specify **Minimum** and **Maximum Bandwidths** here.

For more information, see:

- [“Guidelines for Configuring Minimum and Maximum Bandwidth Values” on page 185.](#)
- [“Configuring Pass-Through Traffic Bandwidths” on page 190.](#)

Configuration - Pass-through

Save Changes

Traffic Bandwidth

Max Bandwidth: (0..10000) Kbps

Min Bandwidth: (0..10000) Kbps

Apply Cancel

Bandwidth				Queuing
ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	10	1,000,000	0	500

Maximum value (in **milliseconds**) for how long packets of this traffic class can queue up before exiting the appliance at the WAN interface. Packets that are queued longer than this maximum value will be dropped.

These are absolute values for the upper and lower **bandwidth** limits, in **kilobits per second**.

Priority determines which traffic class gets serviced first. Multiple classes can have the same priority number. The lower the number, the higher the priority.

Configuring Traffic Classes

Each tunnel — as well as all pass-through shaped traffic — comes with 10 (ten) configurable traffic classes. As the default, the **QoS Policy** points to Traffic Class #1 for any shaped flow.

To give preferential treatment to selected flows, you can configure additional traffic classes to provide the level of granularity you need. Then it's just a matter of pairing **MATCH** criteria with specific traffic classes within the **QoS Policy**.

- If the **Route Policy** directs a flow to a tunnel, then the appliance uses that *specific tunnel's* traffic class.
- If the **Route Policy** designates a flow as pass-through shaped, then the appliance uses the traffic class configured for pass-through shaped traffic.

This can be simplified with ACLs and application groups as part of the **MATCH** criteria. For example, you may want to group VoIP-based traffic for special handling.

Initially, all traffic class pages start with the same default configuration values. In each instance below, Traffic Class #10 has been changed from the default:

On the lower half of the **Configuration - Tunnels** page, this is the **Traffic Class** tab for the selected tunnel, **Tallinn_to_Helsinki**. In the original default, Traffic Classes 9 and 10 had identical values.

Modify Tunnel - (Tallinn_to_Helsinki)

[\[Setting\]](#) [\[Traffic Class\]](#)

ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	9	1,000,000	10	400

On the lower half of the **Configuration - Pass-through** page, this is the **Traffic Class** section. It only applies to flows that the Route Policy designates as *pass-through shaped*.

Traffic Class

ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	7	1,000,000	10	600

Practically speaking, you'll likely define the same 10 traffic classes across all traffic — optimized and pass-through. What will vary are the **Max Wait Time**, and the bandwidth settings, based on the size of each tunnel.

If your appliance only has one tunnel, then you don't need to make any changes to either the tunnel or (shaped) pass-through traffic classes. However, once you configure an additional traffic class or add another tunnel, or both, best practice dictates that you review the values to ensure that you don't run the risk of squeezing out any traffic by over- or under-subscribing bandwidth. To review a model for doing this, see "[Summary of Bandwidth Assessment and Management Tasks](#)" on page 184.

The **Max Wait Time** feature makes it easy to precisely configure QoS queue depths by using a maximum allowable queue wait time, in milliseconds.



Tip For real-time applications that are most sensitive to delay and jitter (for example, voice), we recommend that you set **Max Wait Time** = **100 ms**.

Traffic Class Components

Traffic classes are sets of parameters for *bandwidth* and *queuing*.

These items are **Bandwidth**-related:

- Priority
- Maximum bandwidth (Kbps)
- Minimum bandwidth (Kbps)

To set **Queue** limits, configure:

- Maximum Wait Time (milliseconds)

It's important to understand that this QoS mechanism only has an effect at congestion points.

Modify Tunnel - (Tallinn_to_Helsinki)

[Setting] [Traffic Class]				
ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	9	1,000,000	10	400

The parameter, **Max Wait Time**, allows precise setting of allowable queue latency.

♦ **To edit a traffic class**

- 1 To access the traffic classes, do one of the following:
 - a **For a tunnel**, go to the **Configuration - Tunnels** page and in the **Tunnels** table, click on a tunnel's **Name**. When the details display, click **Traffic Class**.
 - b **For pass-through shaped traffic**, go to the **Configuration - Pass-through** page.

These 10 traffic classes are specific to the tunnel, **Tallinn_to_Helsinki**.

For precise queue length, use **Max Wait Time** to control latency.

The screenshot shows the Silver Peak web interface. At the top, the navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The current page is 'Configuration - Tunnels'. Below the navigation bar, there's a 'Total Tunnels: 1' section. A table lists the tunnel 'Tallinn_to_Helsinki' with status 'down - idle', admin 'up', local IP '192.168.1.2', remote IP '172.50.6.11', max BW 'AUTO/10000', min BW '32/32', IPsec 'no', and up time '0s'. Below this is a 'Modify Tunnel - (Tallinn_to_Helsinki)' section. It has two tabs: 'Setting' and 'Traffic Class'. The 'Traffic Class' tab is active, showing a table with 10 traffic classes. The table has columns: ID, Priority, Max BW (Kbps), Min BW (Kbps), and Max Wait Time (ms). The traffic classes are numbered 1 through 10. Class 1 has priority 5, max BW 1,000,000, min BW 500,000, and max wait time 500. Classes 2 through 9 have priority 10, max BW 1,000,000, min BW 0, and max wait time 500. Class 10 has priority 9, max BW 1,000,000, min BW 10, and max wait time 400. A blue circle highlights the '1' in the ID column of the first row.

ID	Priority	Max BW (Kbps)	Min BW (Kbps)	Max Wait Time (ms)
1	5	1,000,000	500,000	500
2	10	1,000,000	0	500
3	10	1,000,000	0	500
4	10	1,000,000	0	500
5	10	1,000,000	0	500
6	10	1,000,000	0	500
7	10	1,000,000	0	500
8	10	1,000,000	0	500
9	10	1,000,000	0	500
10	9	1,000,000	10	400

These are absolute values for the upper and lower **bandwidth** limits, in **kilobits per second**. The sum of these ten **Min BW** values should *never* exceed the **Min BW** value configured for the tunnel in the **Setting** tab.

Priority determines which traffic class gets serviced first. Multiple classes can have the same priority number. The lower the number, the higher the priority.

To access and edit the **Bandwidth** and **Queuing**, click the desired **ID**.

The screenshot shows the Silver Peak web interface. At the top, the navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The current page is 'Configuration - Traffic Class'. Below the navigation bar, there's a 'Modify Traffic Class (Tallinn_to_Helsinki - class 2)' section. It has a 'Priority' dropdown set to '10' (with a note '1 is the highest priority'). Below this are input fields for 'Max BW' (1000000), 'Min BW' (0), and 'Max Wait Time' (500). Each input field has a range in parentheses: (0..1000000) Kbps for Max BW, (0..1000000) Kbps for Min BW, and (0..2000) ms for Max Wait Time. There are 'Apply', 'Cancel', and 'OK' buttons at the bottom.

- 2 To permanently save edits, first click **Apply** and then click **Save Changes**.

Handling and Marking Packets

All flows that are not explicitly dropped by the Route Policy are subject to DSCP marking by the QoS Policy. DSCP markings enforce end-to-end QoS policies throughout a network.

As with all policies, the appliance searches sequentially through the policy for the first MATCH criteria that applies. If no user-created entries match, then ultimately the default entry applies. For the QoS Policy, the default DSCP values for LAN QoS and WAN QoS are **trust-lan**.

The appliance encapsulates optimized traffic. This process adds an IP outer header to packets for travel across the WAN. However, because pass-through traffic doesn't receive this additional header, its handling is different. The following two sections provide illustrated examples:

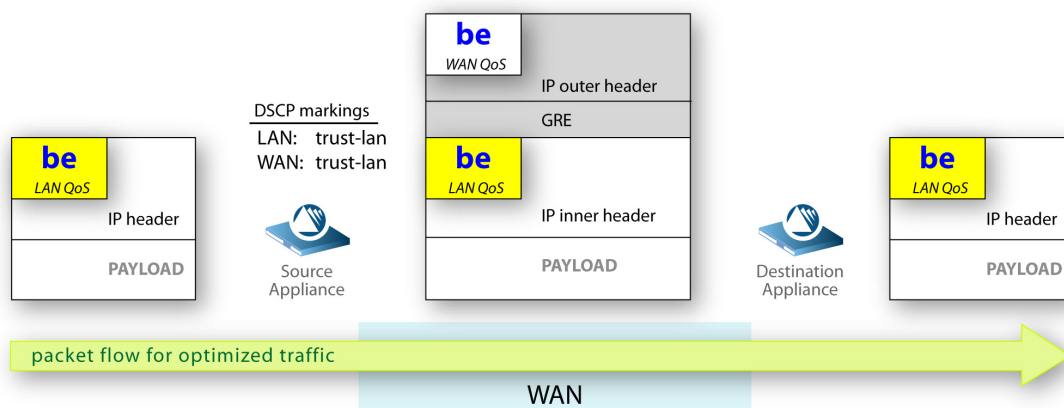
- **Applying DSCP Markings to Optimized Traffic** See page 194.
- **Applying DSCP Markings to Shaped and Unshaped Pass-through Traffic** See page 197.
- **Definitions of DSCP Markings** See page 199.

Applying DSCP Markings to Optimized Traffic

This section illustrates and explains how the appliance applies the QoS Policy to optimized traffic in the following scenarios:

- **LAN and WAN set to trust-lan** See page 194.
- **LAN setting changed, WAN is trust-lan** See page 195.
- **LAN is trust-lan, WAN setting changed** See page 195.
- **LAN setting changed, WAN setting changed** See page 196.

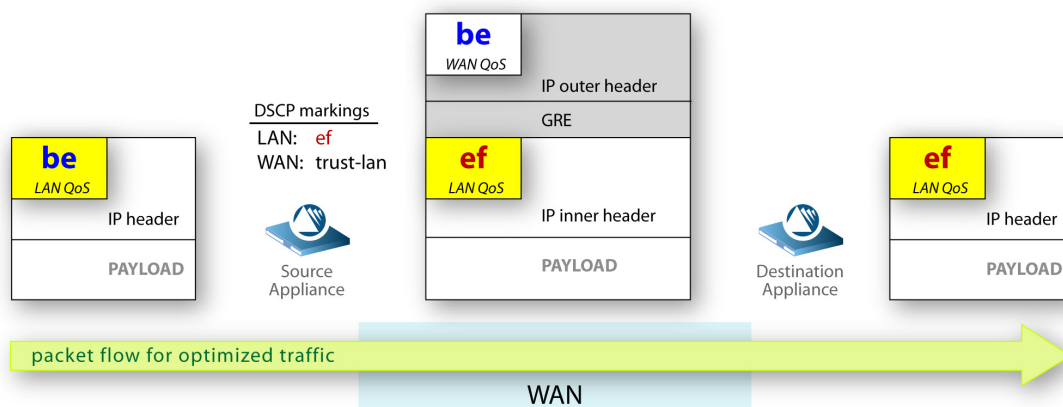
LAN and WAN set to trust-lan



- 1 The source appliance receives the packet from the LAN with a DSCP marking of **be** (best effort).
- 2 Based on MATCH criteria, the QoS Policy applies the LAN QoS setting of **trust-lan**, leaving the LAN DSCP markings as **be** (best effort). As the packet is encapsulated, this is now part of the IP inner header.
- 3 Since the WAN QoS is **trust-lan**, the appliance also sets the WAN QoS bits to **be** in the encapsulating IP outer header.

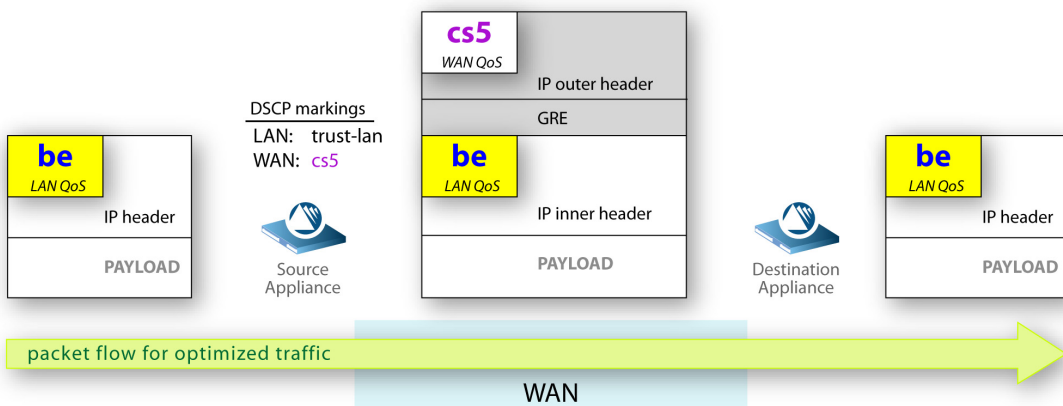
- When the packet reaches the destination appliance, the appliance de-encapsulates the packet, and the packet traverses the LAN with the DSCP markings set to **be**.

LAN setting changed, WAN is trust-lan



- The source appliance receives the packet from the LAN. It has a DSCP marking of **be** (best effort).
- Based on MATCH criteria, the QoS Policy changes the LAN QoS setting to **ef** (express forwarding). As the packet is encapsulated, this is now part of the IP inner header.
- Since the policy's WAN QoS is **trust-lan**, the appliance refers back to the original DSCP markings and sets the WAN QoS bits to **be** in the encapsulating IP outer header.
- When the packet reaches the destination appliance, the appliance de-encapsulates the packet, and the packet traverses the LAN with the DSCP markings set to **ef**.

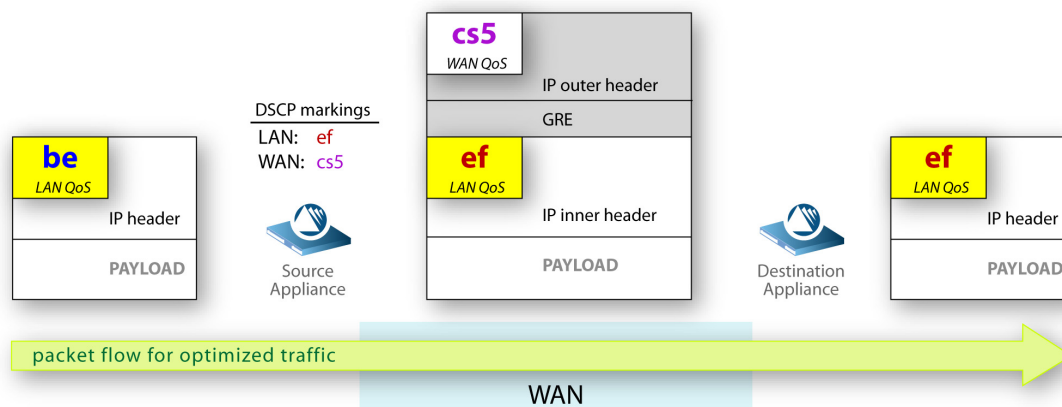
LAN is trust-lan, WAN setting changed



- The source appliance receives the packet from the LAN.
- Based on MATCH criteria, the QoS Policy applies the LAN QoS setting of **trust-lan**, leaving the LAN DSCP markings as **be** (best effort). As the packet is encapsulated, this is now part of the IP inner header.
- Since the policy's WAN QoS action is **cs5** (class selector 5), the appliance sets the bits to **cs5** in the encapsulating IP outer header.

- 4 When the packet reaches the destination appliance, the appliance de-encapsulates the packet, and the packet traverses the LAN with the DSCP markings set to **be**.

LAN setting changed, WAN setting changed



- 1 The source appliance receives the packet from the LAN. It has a DSCP marking of **be** (best effort).
- 2 Based on MATCH criteria, the QoS Policy changes the LAN QoS setting to **ef**. As the packet is encapsulated, this is now part of the IP inner header.
- 3 Since the policy's WAN QoS action is **cs5**, the appliance sets the bits to **cs5** in the encapsulating IP outer header.
- 4 When the packet reaches the destination appliance, the appliance de-encapsulates the packet, and the packet traverses the LAN with the DSCP markings set to **ef**.

Applying DSCP Markings to Shaped and Unshaped Pass-through Traffic

The appliance applies the QoS Policy's DSCP markings to all pass-through flows — whether shaped or unshaped — in the same way:

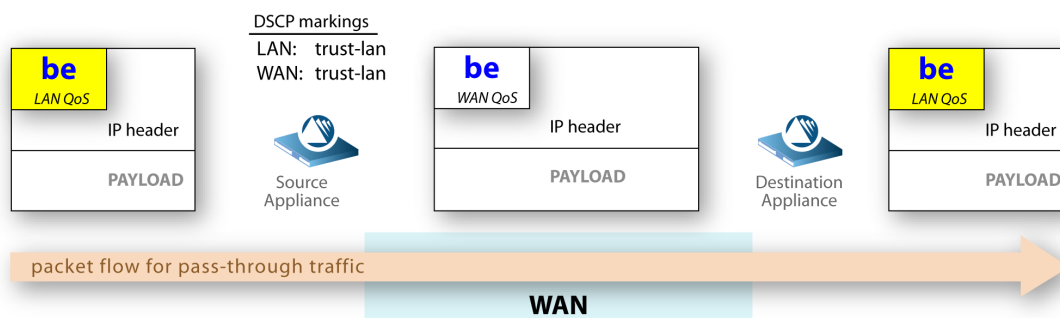
- If there is a match, the appliance applies the WAN QoS setting to the packet (in the IP ToS/DSCP field).
- If there is a LAN QoS setting in the policy match, it is ignored.
- If there is a **trust-lan** setting in the policy match, it is ignored.

To summarize, **all** pass-through traffic is **trust-lan** unless it's modified by the WAN QoS setting. When that's the case, the packet retains the modified QoS setting as it travels through the WAN to the destination appliance.

The following three examples illustrate how the QoS Policy's LAN QoS and WAN QoS settings affect a matched flow's DSCP markings:

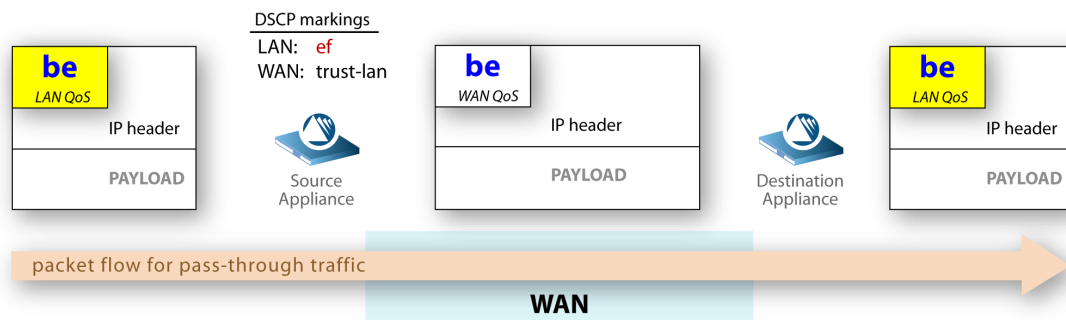
- **LAN and WAN set to trust-lan** See page 197.
- **LAN setting changed, WAN is trust-lan** See page 198.
- **LAN is trust-lan, WAN setting changed** See page 198.

LAN and WAN set to trust-lan



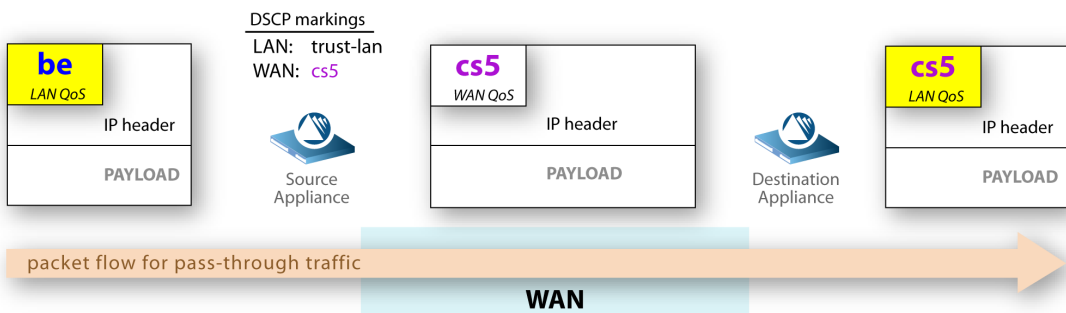
- 1 Because it's pass-through traffic, the appliance ignores the LAN QoS setting.
- 2 Since the WAN QoS is **trust-lan**, the appliance sets the WAN QoS bits to **be** (best effort).
- 3 When the packet reaches the destination appliance, it retains the **be** setting as the LAN receives it.

LAN setting changed, WAN is trust-lan



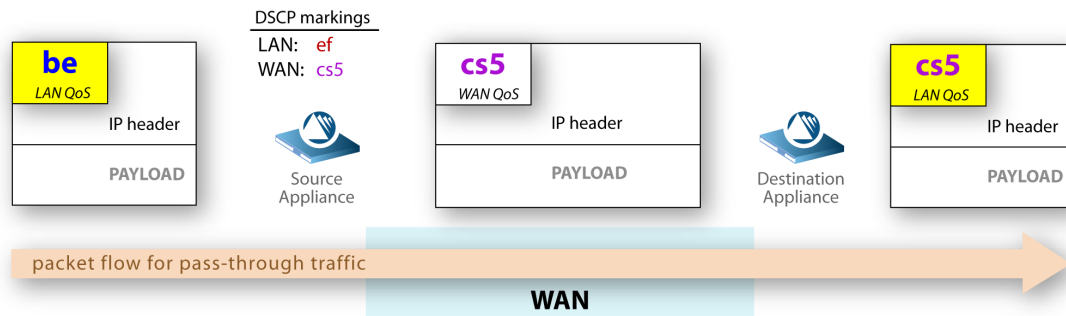
- 1 Because it's pass-through traffic, the appliance ignores the new LAN QoS setting.
- 2 Since the WAN QoS is **trust-lan**, the appliance sets the WAN QoS bits to **be** (best effort).
- 3 When the packet reaches the destination appliance, it retains the **be** setting as the LAN receives it.

LAN is trust-lan, WAN setting changed



- 1 Because it's pass-through traffic, the appliance ignores the LAN QoS setting.
- 2 The appliance sets the WAN QoS bits to **cs5**.
- 3 When the packet reaches the destination appliance, it retains the **cs5** setting as the LAN receives it.

LAN setting changed, WAN setting changed



- 1 Because it's pass-through traffic, the appliance ignores the LAN QoS setting.
- 2 The appliance sets the WAN QoS bits to **cs5**.
- 3 When the packet reaches the destination appliance, it retains the **cs5** setting as the LAN receives it.

Definitions of DSCP Markings

Following is a list of definitions for the available Differentiated Services Code Point (DSCP) markings, which use a 6-bit value to indicate Per-Hop Behavior (PHB):

DSCP Marking	Per-Hop Behavior Group	Codepoint	Number
be	Best Effort	000000	DSCP 0
af11	Assured Forwarding 11	001010	DSCP 10
af12	Assured Forwarding 12	001100	DSCP 12
af13	Assured Forwarding 13	001110	DSCP 14
af21	Assured Forwarding 21	010010	DSCP 18
af22	Assured Forwarding 22	010100	DSCP 20
af23	Assured Forwarding 23	010110	DSCP 22
af31	Assured Forwarding 31	011010	DSCP 26
af32	Assured Forwarding 32	011100	DSCP 28
af33	Assured Forwarding 33	011110	DSCP 30
af41	Assured Forwarding 41	100010	DSCP 34
af42	Assured Forwarding 42	100100	DSCP 36
af43	Assured Forwarding 43	100110	DSCP 38
cs1	Class Selector 1 (precedence 1)	001000	CS1
cs2	Class Selector 2 (precedence 2)	010000	CS2
cs3	Class Selector 3 (precedence 3)	011000	CS3
cs4	Class Selector 4 (precedence 4)	100000	CS4
cs5	Class Selector 5 (precedence 5)	101000	CS5

DSCP Marking	Per-Hop Behavior Group	Codepoint	Number
cs6	Class Selector 6 (precedence 6)	110000	CS6
cs7	Class Selector 7 (precedence 7)	111000	CS7
ef	Expedited Forwarding	101110	DSCP 46

QoS Policy Page Organization

The following shows the SET actions for the appliance, **spdcnx01**.

This QoS Policy has two maps — **gms_QoSMap** and **map1**. Here, **gms_QoSMap** is active and therefore it's the policy.

To activate another QoS Policy, select a map from the drop-down menu and click **Activate**. Any change governs all new flows.

When there are multiple maps, they display below the active map, or policy. You can **only** change a map's name when it's inactive.

The screenshot shows the Silver Peak Configuration - QoS Policy page. The page has a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The current page is Configuration - QoS Policy. The page shows two QoS maps: **gms_QoSMap (active)** and **map1**. Each map has a table of Match Criteria and Set Actions. The **gms_QoSMap** table has 5 rows of match criteria and 3 rows of set actions. The **map1** table has 5 rows of match criteria and 3 rows of set actions. The **map1** map is currently inactive.

Match Criteria										Set Actions		
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class	LAN QoS	WAN QoS	
<input type="checkbox"/> 20		ip	0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	5	trust-lan	trust-lan	
<input type="checkbox"/> 30		ip	0.0.0.0/0	0.0.0.0/0	Transactional		any	any.any	4	trust-lan	trust-lan	
<input type="checkbox"/> 40		ip	0.0.0.0/0	0.0.0.0/0	Interactive		any	any.any	3	trust-lan	trust-lan	
<input type="checkbox"/> 50		ip	0.0.0.0/0	0.0.0.0/0	VOIP		any	any.any	2	trust-lan	trust-lan	
<input type="checkbox"/> 60		ip	0.0.0.0/0	0.0.0.0/0	any		any	any.any	10	trust-lan	trust-lan	
<input type="checkbox"/> default									1	trust-lan	trust-lan	

Match Criteria										Set Actions		
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class	LAN QoS	WAN QoS	
<input type="checkbox"/> 1		ip	172.20.38.0/24	0.0.0.0/0	VOIP		any	any.any	2	trust-lan	trust-lan	
<input type="checkbox"/> 20		ip	0.0.0.0/0	0.0.0.0/0	Interactive		any	any.any	3	trust-lan	trust-lan	
<input type="checkbox"/> 30		ip	0.0.0.0/0	0.0.0.0/0	Transactional		any	any.any	4	trust-lan	trust-lan	
<input type="checkbox"/> 40		ip	0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	5	trust-lan	trust-lan	
<input type="checkbox"/> 50		ip	0.0.0.0/0	0.0.0.0/0	any		any	any.any	10	trust-lan	trust-lan	
<input type="checkbox"/> default									1	trust-lan	trust-lan	

Every policy — Route, Optimization, or QoS — starts with an active default map named, **map1**. Here, **map1** is inactive.

However, since they're different **types** of policies, the fact that they have the same name **does not** indicate (or cause) a relationship among them.

Set Actions		
Traffic Class	LAN QoS	WAN QoS
1	trust-lan	trust-lan
1	trust-lan	trust-lan

The QoS Policy doesn't apply DSCP markings for **LAN QoS** if the flows are pass-through shaped or unshaped.

Managing the QoS Policy

This section discusses procedures for managing the QoS Policy. It includes:

- **Adding an Entry to a Map** See page 202.
- **Editing an Entry** See page 204.
- **Deleting an Entry** See page 205.
- **Adding a New QoS Map** See page 206.
- **Deleting a Map** See page 208.
- **Activating a New Policy** See page 209.

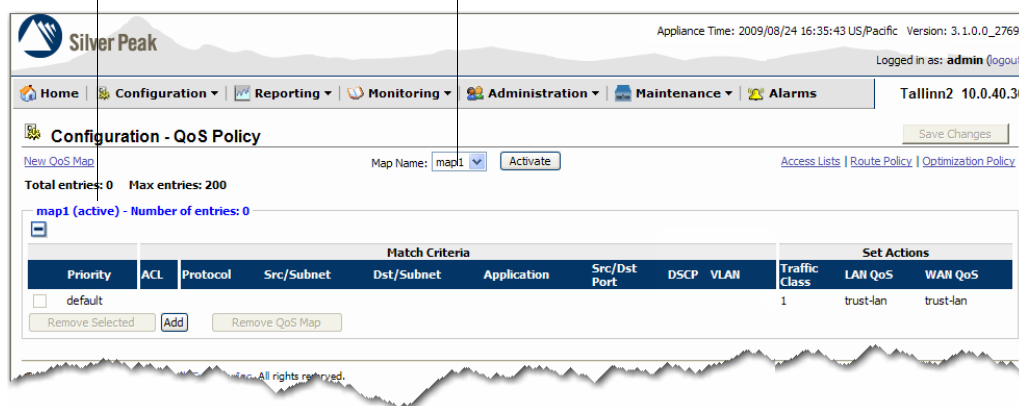
Adding an Entry to a Map

When you create a new entry in the **QoS Policy**, the Appliance Manager defaults to directing all flows to Traffic Class #1 and trusting the DSCP markings the packets had when arriving from the LAN.

♦ To add an entry to a map

- 1 From the **Configuration** menu, select **QoS Policy**. The **Configuration – QoS Policy** page appears. If this is the first time you're accessing the page, only the default entry appears.

The default QoS Policy is always **map1**...
...and it is active.



The QoS Policy has a default entry that isn't numbered. It always occupies the last priority position. In reality, its assigned value is **65535**.

If traffic doesn't match any user-configured entries, then the policy applies the default entry.

- Click **Add**. A new entry appears, with editable fields.

If you need to review the list of existing ACLs, click **Access Lists**. The **Configuration – Access Lists** page displays.

After you've reviewed the entries, you can return to this page by clicking **QoS Policy** in the same location on that page.

To use an existing ACL, select its name from the **ACL** column's drop-down menu. Once you do, the other fields in the **Match Criteria** area are inaccessible.

To complete the tuple elements individually, complete these fields.

As described in Chapter 6, “Theory of Operations,” do one of the following:

- Complete the individual fields belonging to the **Match Criteria**, or
- From the **ACL** column, select the name of an existing Access Control List.

- Edit the SET actions as desired.

By default, the appliance:

- trusts and propagates the DSCP markings that the packets had on arrival from the LAN.
- directs flows going to a tunnel to the tunnel's Traffic Class #1.
- directs pass-through shaped flows to Traffic Class #1, as defined on the **Configuration - Pass-through** page.

- Click **Apply**.
- Click **Save Changes**.

Editing an Entry

You can edit an active entry without having to “deactivate” it first. Changes to an QoS map entry affect new connections only.

If you’ve added additional QoS maps, you can also edit their entries.

The **default** entry (always the last one in a policy) cannot be edited.

♦ To edit an entry

- 1 From the **Configuration** menu, select **QoS Policy**. The **Configuration – QoS Policy** page appears.

Configuration - QoS Policy

Map Name: map1 [Access Lists](#) [Route Policy](#) [Optimization Policy](#)

Total entries: 1 Max entries: 200

map1 (active) - Number of entries: 1

Priority	ACL	Match Criteria						Set Actions			
		Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class	LAN QoS	WAN QoS
<input type="checkbox"/> 10		ip	0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	1	be	be
<input type="checkbox"/> default									1	trust-lan	trust-lan

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

BulkData is a user-defined application group.

- 2 To edit an entry, click that entry's **Priority** number. The fields become editable.

Configuration - QoS Policy

Map Name: map1 [Access Lists](#) [Route Policy](#) [Optimization Policy](#)

Total entries: 1 Max entries: 200

map1 (active) - Number of entries: 1

Priority	ACL	Match Criteria						Set Actions			
		Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class	LAN QoS	WAN QoS
<input type="text" value="10"/>		ip	0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	1	be	be
<input type="checkbox"/> default									1	trust-lan	trust-lan

- 3 Make the desired changes.
- 4 Click **Apply**.
- 5 Click **Save Changes**.

Deleting an Entry

The procedure for deleting an entry is the same across all maps, lists, and policies.

♦ To delete an entry

- 1 In the leftmost column of the policy, click the check box(es) for the entries you want to delete.

The screenshot shows the Silver Peak Configuration - QoS Policy interface. The top navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The main section is titled 'Configuration - QoS Policy' and shows 'map1 (active) - Number of entries: 1'. Below this is a table with columns: Priority, ACL, Protocol, Src/Subnet, Dst/Subnet, Application, Src/Dst Port, DSCP, VLAN, Traffic Class, and LAN QoS. The first entry has a green checkmark in the leftmost column. Below the table are buttons for 'Remove Selected', 'Add', and 'Remove QoS Map'.

Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class	LAN QoS
<input checked="" type="checkbox"/> 10		ip	0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	1	be
<input type="checkbox"/> default									1	trust-lan

Buttons: Remove Selected, Add, Remove QoS Map

A green check displays inside, and the **Remove Selected** button becomes accessible.

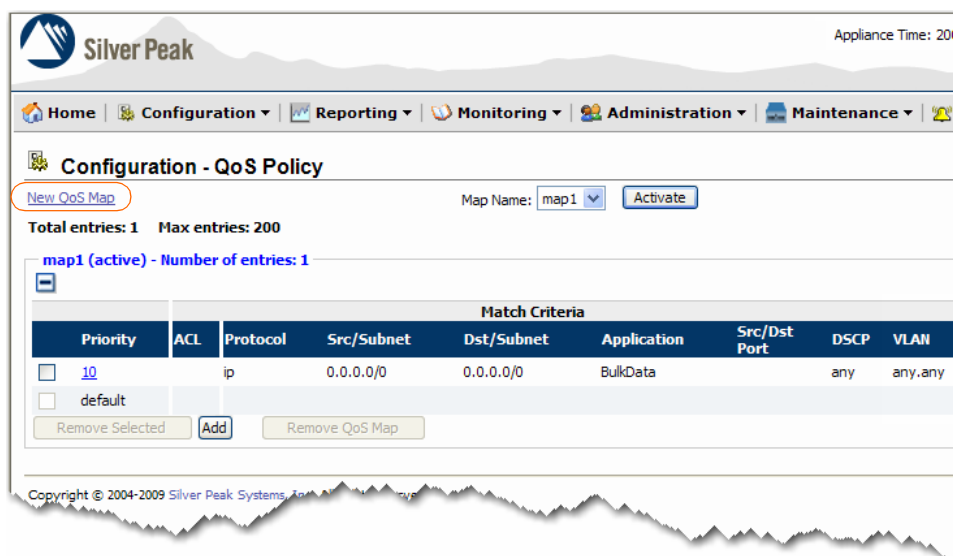
- 2 Click **Remove Selected**.
- 3 Click **Save Changes**.

Adding a New QoS Map

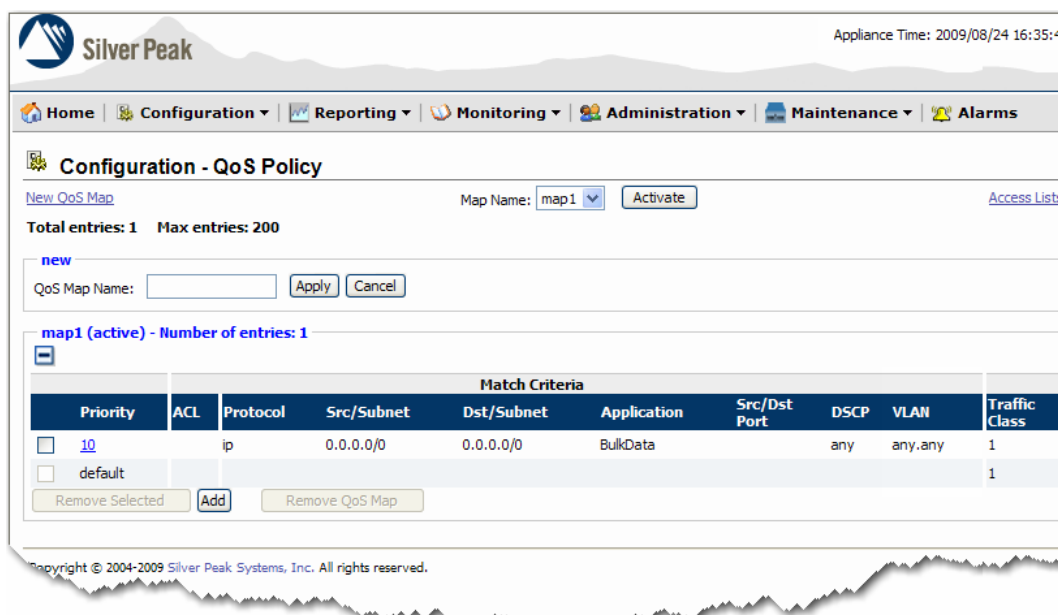
You can always add another QoS map, whether or not you choose to activate it immediately.

◆ To add a new map

- 1 In the **Configuration – QoS Policy** page, click **New QoS Map**.



The **new** section displays above the active map.



- 2 In the **QoS Map Name** field, enter a new map name.

Silver Peak

Appliance Time: 2009/08/24 16:35:4

Home Configuration Reporting Monitoring Administration Maintenance Alarms

Configuration - QoS Policy

[New QoS Map](#) Map Name: [Access Lists](#)

Total entries: 1 Max entries: 200

new

QoS Map Name:

map1 (active) - Number of entries: 1

Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Traffic Class
<input type="checkbox"/> 10	ip		0.0.0.0/0	0.0.0.0/0	BulkData		any	any.any	1
<input type="checkbox"/> default									1

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

- 3 Click **Apply**.

The new map, **Financial**, displays in expanded view below the active map (which is the policy). Meanwhile, all other maps are collapsed.

Silver Peak

Appliance Time: 2009/0

Home Configuration Reporting Monitoring Administration Maintenance Alarms

Configuration - QoS Policy

[New QoS Map](#) Map Name:

Total entries: 1 Max entries: 200

map1 (active) - Number of entries: 1

Financial - Number of entries: 0

Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN
<input type="checkbox"/> default								

© 2004-2009 Silver Peak Systems, Inc. All rights reserved.

Each time you create a new QoS map, the **default** entry is always the same.


- 4 Add new entries, as necessary.
- 5 Click **Save Changes**.

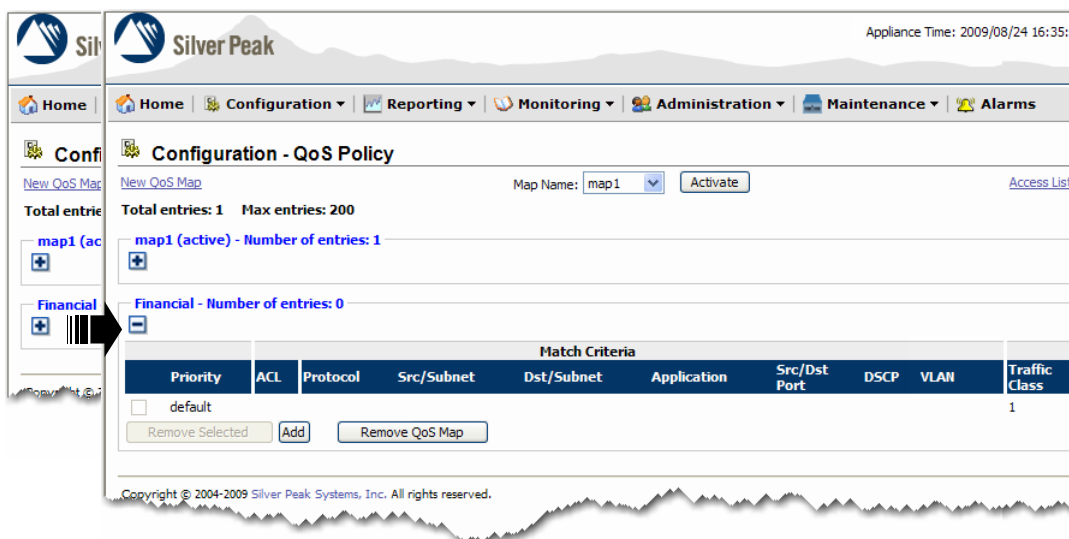
Deleting a Map

You can **only delete** an **inactive** map. That is, a map that's not in use as the policy.


If the map you want to delete is active, first see *"Activating a New Policy" on page 209*.

♦ To delete a map

- 1 Under the map's name, click the  to expand the inactive map you want to delete.



- 2 Click **Remove QoS Map**. A window appears, requesting confirmation.




- 3 Click **OK**.
- 4 Click **Save Changes**.

Activating a New Policy

Activating a new QoS map deactivates the old one. All new traffic matches against the new QoS Policy.

♦ To activate a new policy

- 1 From the **Map Names** field, select the map you want to activate.



- 2 Click **Activate**. The old map deactivates and the new one activates, beginning with all new flows.
- 3 Click **Save Changes**.



Optimization Policy

This chapter describes how the appliance optimizes tunnelized traffic — improving the performance of applications across the WAN.

In This Chapter

- **Introduction** See page 212.
- **Making the Best Use of Optimizations** See page 214.
- **When the Appliance Can Apply the Optimization Policy** See page 215.
- **Optimization Policy Page Organization** See page 216.
- **Managing the Optimization Policy** See page 217.

Introduction

The **Optimization Policy** uses optimization techniques to improve the performance of applications across the WAN. It asks:

- What optimization techniques do I want to apply to a given flow?
- Are there any flows that don't need optimization?



Note If a flow is **not** directed to a tunnel, it's **not** subject to the Optimization Policy.

The Optimization Policy's SET actions include:

- **Network Memory** See page 212.
- **Payload Compression** See page 213.
- **TCP Acceleration** See page 213.
- **CIFS Acceleration** See page 213.

Network Memory

All Silver Peak NX Series appliances are equipped with Network Memory™ technology. Network Memory is used to inspect all inbound and outbound WAN traffic in real-time, storing a single local instance of data on each appliance.

Before sending information across the WAN, NX Series appliances compare real-time traffic streams to patterns stored using Network Memory. If a match exists, a short reference pointer is sent to the remote Silver Peak appliance, instructing it to deliver the traffic pattern from its local instance. Repetitive data is never sent across the WAN.

If content is modified, the Silver Peak appliance detects the change at the byte level and updates the network's "memory". Only the modifications are sent across the WAN. These are combined with original content by NX Series appliances at the destination location.

Benefit scenarios

The following scenarios exemplify the benefits of Network Memory.

File Server Even when the file is not identical to the version that was previously downloaded, significant performance improvements are realized by transporting only the incremental changes across the WAN.

Web If a web application is generating dynamic pages (for example, using HTTP), only delta information is transferred. For example, a SharePoint table with many rows updates by just transmitting the delta for the row, rather than the whole page.

Video streaming and Video On Demand If several employees in an office chose to watch the same video (for example, a distance learning module or a taped CEO address), Network Memory eliminates the need to send multiple copies across the WAN. This has the same advantage if they all are watching the video simultaneously, or at different times.

Software patch distribution and upgrades If several employees in an office need to download the same software patch, Network Memory eliminates the need to send multiple copies across the WAN.

Remote backups Once the first backup is completed, future "full" backups are effectively reduced to "incremental backups" as far as WAN traffic is concerned.

Payload Compression

Compression reduces the bandwidth consumed by traffic traversing the WAN. ***Payload compression*** uses algorithms to identify relatively short byte sequences that are repeated frequently over time. These sequences 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.

Header compression provides additional bandwidth gains by reducing packet header information using specialized compression algorithms. It's always enabled and not user-configurable.

Silver Peak NX Series appliances include state of the art, cross-flow data compression and header compression as part of a broader Local Instance Networking solution. Information gleaned from the compression of one flow can be applied to other flows.

Payload compression is used in conjunction with Network Memory to provide compression on “first pass” data.

TCP Acceleration

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

CIFS Acceleration

Silver Peak appliances provide enhancements to accelerate CIFS—including read-aheads, write-behinds, and metadata caching. This reduces the impact of latency on data transfers using this protocol.

Making the Best Use of Optimizations

By default, the Optimization Policy's default entry applies all four user-selectable optimizations — Network Memory, payload compression, TCP acceleration, and CIFS acceleration.

In general, leaving all optimizations defaulted to ON is the best practice. However, there are scenarios where it makes sense to change these defaults. For example:

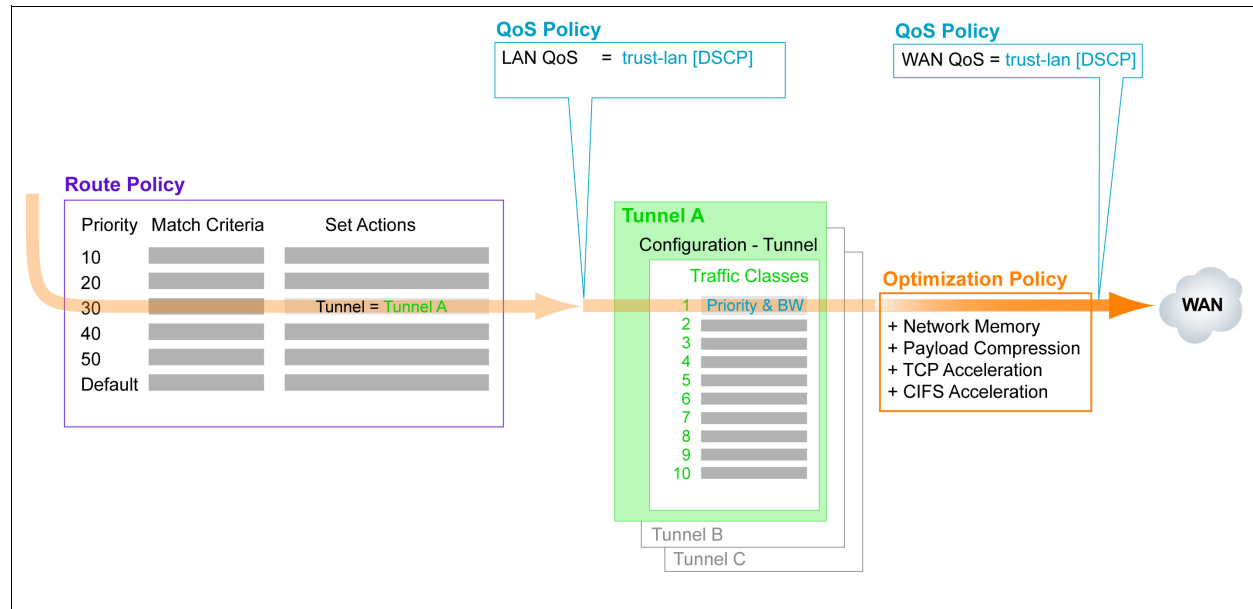
- 1 For real-time calls using VoIP, configure the optimizations as follows:
 - Network Memory – **OFF**
 - Payload Compression – **OFF**
 - TCP Acceleration – ON
 - CIFS Acceleration – ON
- 2 For end-to-end encrypted traffic (that's encrypted on a per-session basis), configure the optimizations as follows:
 - Network Memory – **OFF**
 - Payload Compression – **OFF**
 - TCP Acceleration – ON
 - CIFS Acceleration – ON
- 3 If there's a VoIP element in a videocast (from one source to multiple, individual users), then configure the optimizations as follows:
 - Network Memory – ON
 - Payload Compression – ON
 - TCP Acceleration – ON
 - CIFS Acceleration – ON

Usually, the only reasons to turn off TCP Acceleration and/or CIFS Acceleration are for:

- debugging
- testing the impacts of various acceleration techniques.

When the Appliance Can Apply the Optimization Policy

This diagram shows how the appliance processes a flow assigned to a tunnel.



- 1 First, the Route Policy checks traffic incoming from the LAN against the MATCH criteria in its prioritized entries. Entries 10 and 20 don't match the traffic, but Entry 30 does.
- 2 The policy applies the entry's SET actions to the identified flow. In this case, it sends the flow to **Tunnel A**. Once any traffic matches an entry, no subsequent entries are examined.
- 3 Before the flow reaches **Tunnel A**, the QoS Policy checks against its entries and
 - applies the DSCP marking specified for LAN QoS, and
 - tells the flow which of **Tunnel A**'s traffic classes to use. All traffic classes for optimized flows are tunnel-specific. That is, they're part of the tunnel's configuration.
- 4 The appliance places the flow in Traffic Class #1 and passes the flow to the Optimization Policy.
Only flows directed to tunnels are subject to the Optimization Policy.
- 5 After optimization, the QoS Policy applies the rest of the SET action, which is to apply the DSCP markings for the WAN QoS.
- 6 The appliance queues the optimized flow to exit the physical WAN interface.

Optimization Policy Page Organization

The following shows the SET actions within an Optimization Policy on a Silver Peak NX appliance.

This Optimization Policy has one map — **map1**. It's active and therefore it is **the** Optimization Policy that will be applied to all tunneled data flows.

Every policy — Route, Optimization, or QoS — has an active default map with this name. However, since they're different **types** of policies, the fact that they have the same name **does not** indicate (or cause) a relationship among them.

To activate another Optimization Policy, select a map from the drop-down menu and click **Activate**. Any change governs all new flows.

The screenshot shows the 'Configuration - Optimization Policy' page. At the top, there is a navigation bar with links like Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The main content area shows the 'map1' configuration. It includes a 'Map Name' dropdown set to 'map1' and an 'Activate' button. Below this, there is a table with match criteria and a table with set actions. The match criteria table has columns: Priority, ACL, Protocol, Src/Subnet, Dst/Subnet, Application, Src/Dst Port, DSCP, and VLAN. The set actions table has columns: Network Memory, Payload Compression, TCP Accel, and CIFS Accel. The 'VoIP' application is selected in the match criteria table. The 'Network Memory' set action is unchecked for the 'VoIP' application.

When there are multiple maps, they list under the active map, or policy. Here, there are none.

You can **only** change a map's name when it's inactive.

Here, **VoIP** is a user-created application group that includes **h_323** and **cisco_skinny**.

Those details aren't apparent on *this* screen, but you could see them by selecting **Configuration > Application > Groups** from the main menu bar.

Notice that for the application group, **VoIP**, **Network Memory** is unchecked, that is, turned **OFF**.

Set Actions			
Network Memory	Payload Compression	TCP Accel	CIFS Accel
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

The Optimization Policy is applied to any flow that matches the Route Policy entries..

Managing the Optimization Policy

This section discusses procedures for managing the Optimization Policy. It includes:

- **Adding an Entry to a Map** See page 217.
- **Editing an Entry** See page 219.
- **Deleting an Entry** See page 220.
- **Adding a New Optimization Map** See page 220.
- **Deleting a Map** See page 223.
- **Activating a New Policy** See page 224.

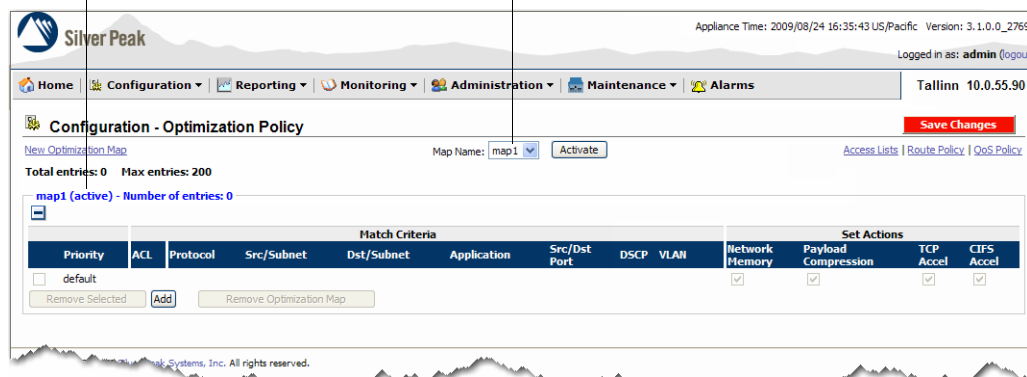
Adding an Entry to a Map

When you create a new entry in the Optimization Policy, the Appliance Manager defaults to leaving all four optimizations ON, until you change it.

♦ To add an entry to a map

- 1 From the **Configuration** menu, select **Optimization Policy**. The **Configuration – Optimization Policy** page appears. If this is the first time you’re accessing the page, only the default entry appears.

The default Optimization Policy is always map1...
...and it is active.



The Optimization Policy has a default entry that isn't numbered. It always occupies the last priority position. In reality, its assigned value is 65535.

If traffic doesn't match any user-configured entries, then the default entry applies all optimizations to all flows that the Route Policy is directing to a tunnel.

- 2 Click **Add**. A new entry appears, with editable fields.

If you need to review the list of existing ACLs, click **Access Lists**. The **Configuration – Access Lists** page displays.

After you've reviewed the entries, you can return to this page by clicking **Optimization Policy** in the same location on that page.

To use an existing ACL, select its name from the **ACL** column's drop-down menu. Once you do, the other fields in the **Match Criteria** area are inaccessible.

To complete the tuple elements individually, complete these fields.

Configuration - Optimization Policy

Map Name: map1 [Activate]

Total entries: 0 Max entries: 200

map1 (active) - Number of entries: 0

Priority	ACL	Match Criteria							Set Actions			
		Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Network Memory	Payload Compression	TCP Accel	CIFS Accel
new 10	[dropdown]	ip	0.0.0.0/0	0.0.0.0/0	any	0:0	any	any	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Apply Cancel

As described in [Chapter 6, “Theory of Operations,”](#) do one of the following:

- Complete the individual fields belonging to the **Match Criteria**, or
- From the **ACL** column, select the name of an existing Access Control List.

- 3 By default, all optimizations are ON. If you want to turn any off, click the green check mark to uncheck the option.

To review when you might want to turn any of the optimizations off, see [“Making the Best Use of Optimizations” on page 214](#).

Set Actions			
Network Memory	Payload Compression	TCP Accel	CIFS Accel
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

- 4 Click **Apply**.
- 5 Click **Save Changes**.

Editing an Entry

You can edit an active entry without having to “deactivate” it first. Changes to an Optimization map entry affect new connections only.

If you’ve added additional Optimization maps, you can also edit their entries.

The **default** entry (always the last one in a policy) cannot be edited.

♦ To edit an entry

- 1 From the **Configuration** menu, select **Optimization Policy**. The **Configuration – Optimization Policy** page appears.

- 2 To edit an entry, click that entry’s **Priority** number. The fields become editable.

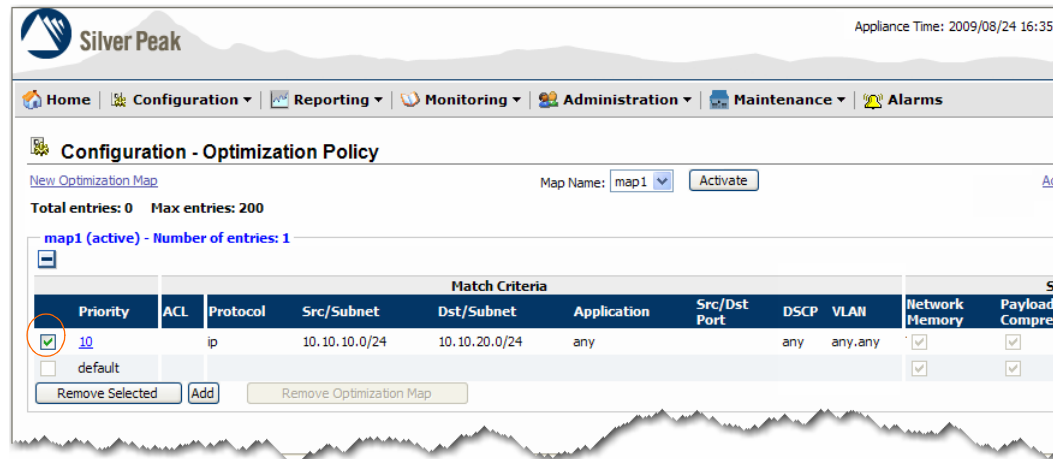
- 3 Make the desired changes.
- 4 Click **Apply**.
- 5 Click **Save Changes**.

Deleting an Entry

The procedure for deleting an entry is the same across all maps, lists, and policies.

♦ To delete an entry

- 1 In the leftmost column of the policy, click the check box(es) for the entries you want to delete.



A green check displays inside, and the **Remove Selected** button becomes accessible.

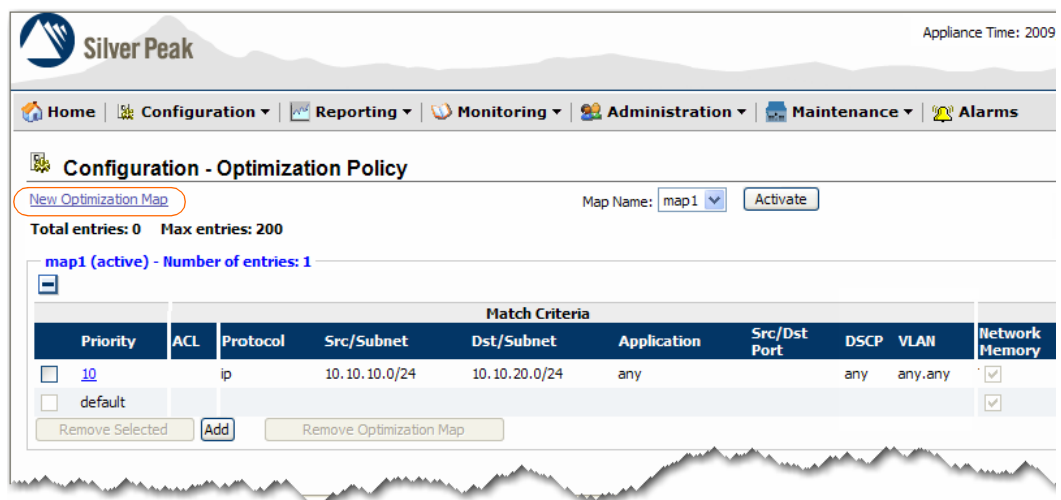
- 2 Click **Remove Selected**.
- 3 Click **Save Changes**.

Adding a New Optimization Map

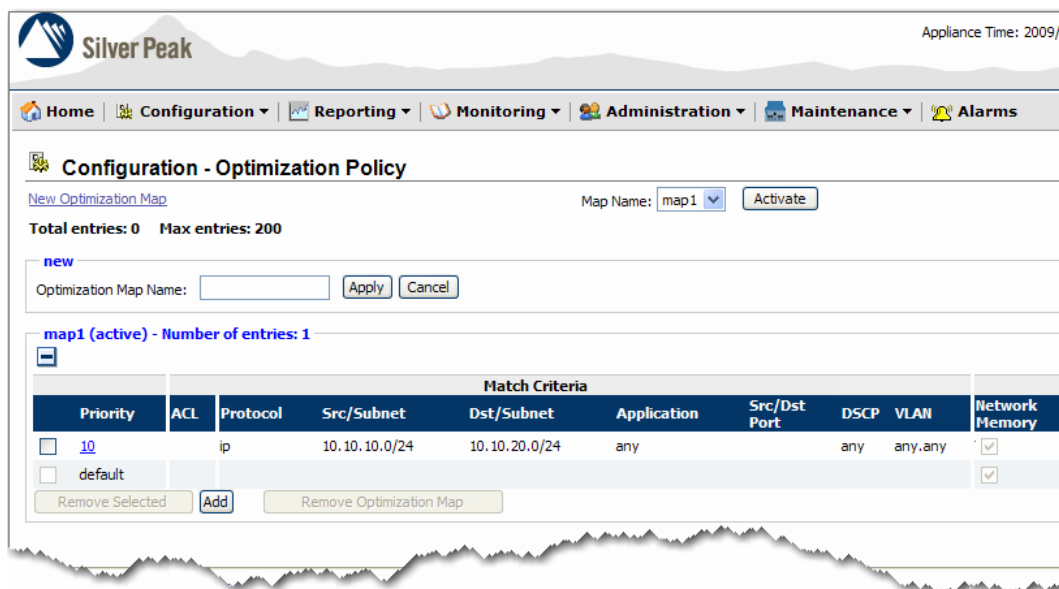
You can always add another Optimization map, whether or not you choose to activate it immediately.

♦ To add a new map

- 1 In the **Configuration – Optimization Policy** page, click **New Optimization Map**.



The **new** section displays above the active map.



Silver Peak Appliance Time: 2009/

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms

Configuration - Optimization Policy

[New Optimization Map](#) Map Name:

Total entries: 0 Max entries: 200

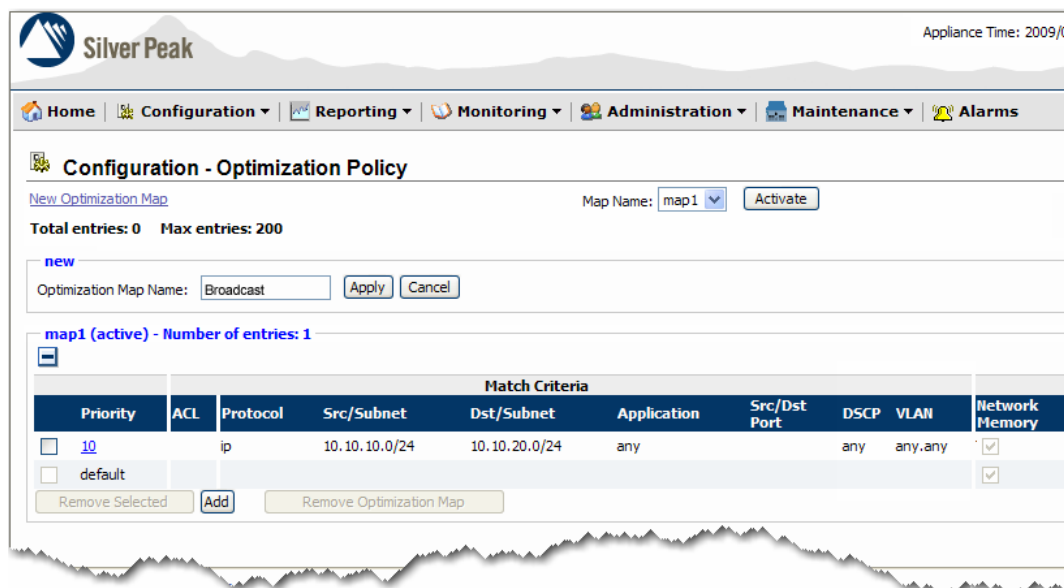
new

Optimization Map Name:

map1 (active) - Number of entries: 1

Match Criteria									
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Network Memory
<input type="checkbox"/> 10		ip	10.10.10.0/24	10.10.20.0/24	any		any	any.any	<input checked="" type="checkbox"/>
<input type="checkbox"/> default									<input checked="" type="checkbox"/>

- 2 In the **Optimization Map Name** field, enter a new map name.



Silver Peak Appliance Time: 2009/

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms

Configuration - Optimization Policy

[New Optimization Map](#) Map Name:

Total entries: 0 Max entries: 200

new

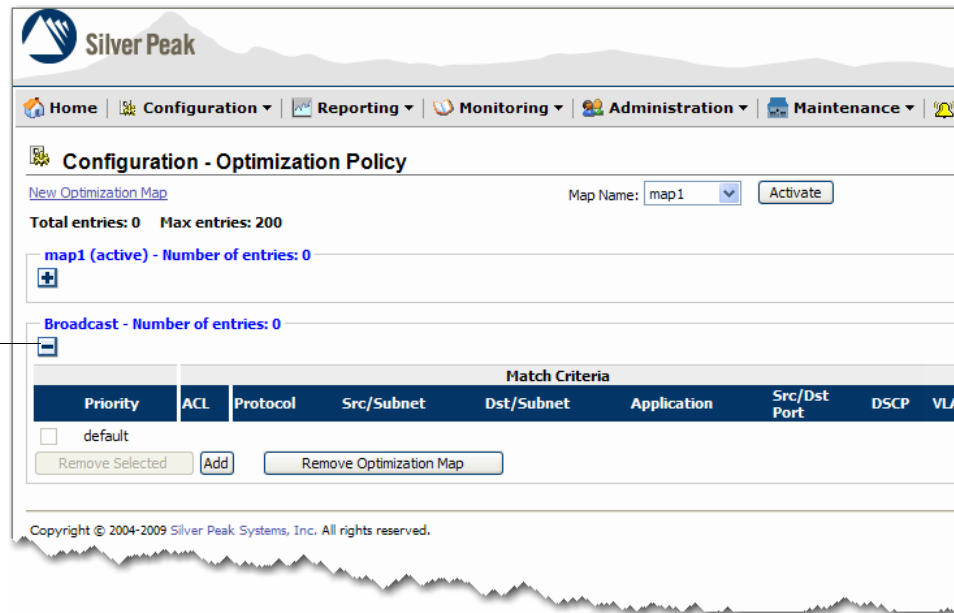
Optimization Map Name:

map1 (active) - Number of entries: 1

Match Criteria									
Priority	ACL	Protocol	Src/Subnet	Dst/Subnet	Application	Src/Dst Port	DSCP	VLAN	Network Memory
<input type="checkbox"/> 10		ip	10.10.10.0/24	10.10.20.0/24	any		any	any.any	<input checked="" type="checkbox"/>
<input type="checkbox"/> default									<input checked="" type="checkbox"/>

3 Click Apply.

The new map, **Broadcast**, displays in expanded view below the active map (which is the policy). Meanwhile, all other maps are collapsed.




Each time you create a new Optimization map, the **default** entry is always the same. It applies all optimizations to flows directed to tunnels.

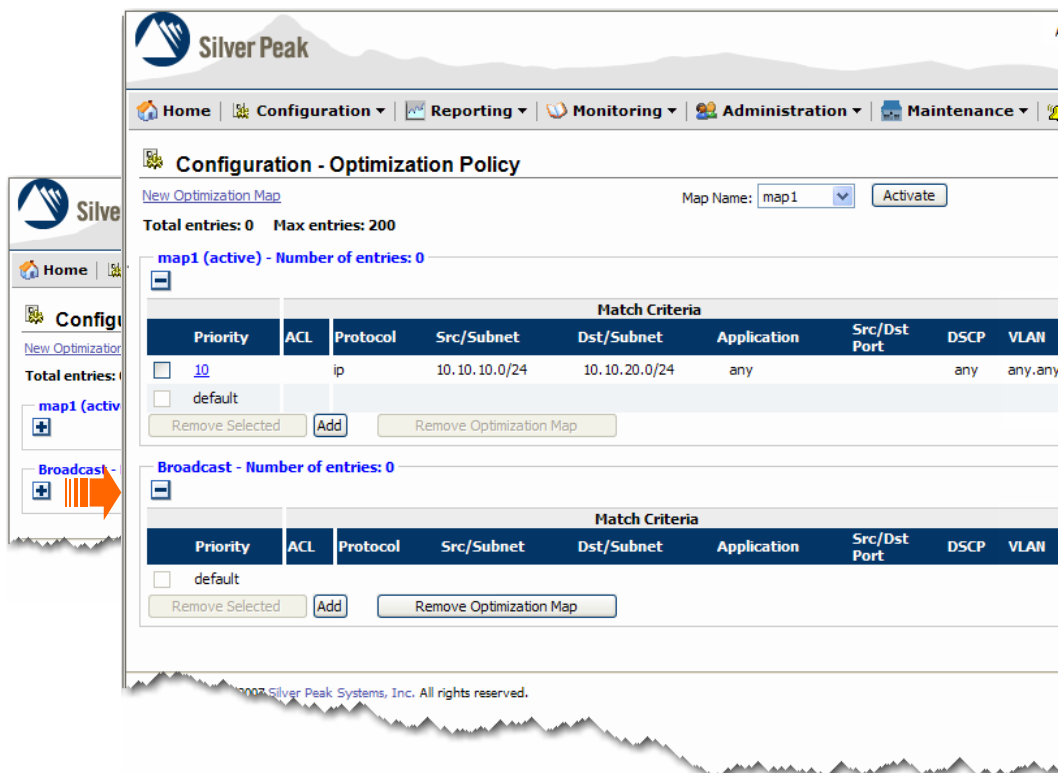
- 4** Add new entries, as necessary.
- 5** Click **Save Changes**.

Deleting a Map


You can **only delete** an **inactive** map. That is, a map that's not in use as the policy.

♦ To delete a map

- 1 Under the map's name, click the  to expand the inactive map you want to delete.



- 2 Click **Remove Optimization Map**. A window appears, requesting confirmation.




- 3 Click **OK**.
- 4 Click **Save Changes**.

Activating a New Policy

Activating a new Optimization map deactivates the old one. All new traffic matches against the new Optimization Policy.

♦ To activate a new policy

- 1 From the **Map Names** field, select the map you want to activate.



- 2 Click **Activate**. The old map deactivates and the new one activates, beginning with all new flows.
- 3 Click **Save Changes**.



Using Flow Redirection to Address TCP Asymmetry

This chapter describes how Flow Redirection allows Silver Peak appliances to optimize asymmetrically routed flows by redirecting packets between appliances.

The flow redirection feature is implemented solely in software, and is available in both bridge and router modes.

In This Chapter

- **Introduction** See page 226.
- **Configuring Flow Redirection** See page 229.
- **Flow Reporting** See page 234.


Introduction

A network is asymmetric when a client request and its server response don't use the same path through the network. This asymmetric network configuration is common for:

- Financial institutions, which virtualize geographically separate data centers for load balancing and redundancy.
- Businesses that have multiple ISP paths across a customer network.

Asymmetrical Networks and Flows

The following diagram shows a sample asymmetric network. In this example, each server appliance sees only one direction of the traffic flow.



For TCP flows to be optimized, both directions must travel through the same client and server appliances.

Removing Asymmetry with Flow Redirection

Flow redirection removes the asymmetry locally by merging the traffic of an asymmetric flow into a single appliance. An appliance that handles both directions of traffic for a flow can then optimize the flow properly. Specifically, this sets the stage for TCP acceleration and CIFS acceleration.

With flow redirection, the appliance that receives the first packet — that is, the TCP SYN packet — owns the flow and eventually receives all of that flow's traffic. To be able to redirect, appliances are configured into **clusters**, whereby they communicate with each other and keep track of flows. Any given appliance can own multiple flows and redirect others, depending on whether or not the appliance received the initial TCP SYN packet.

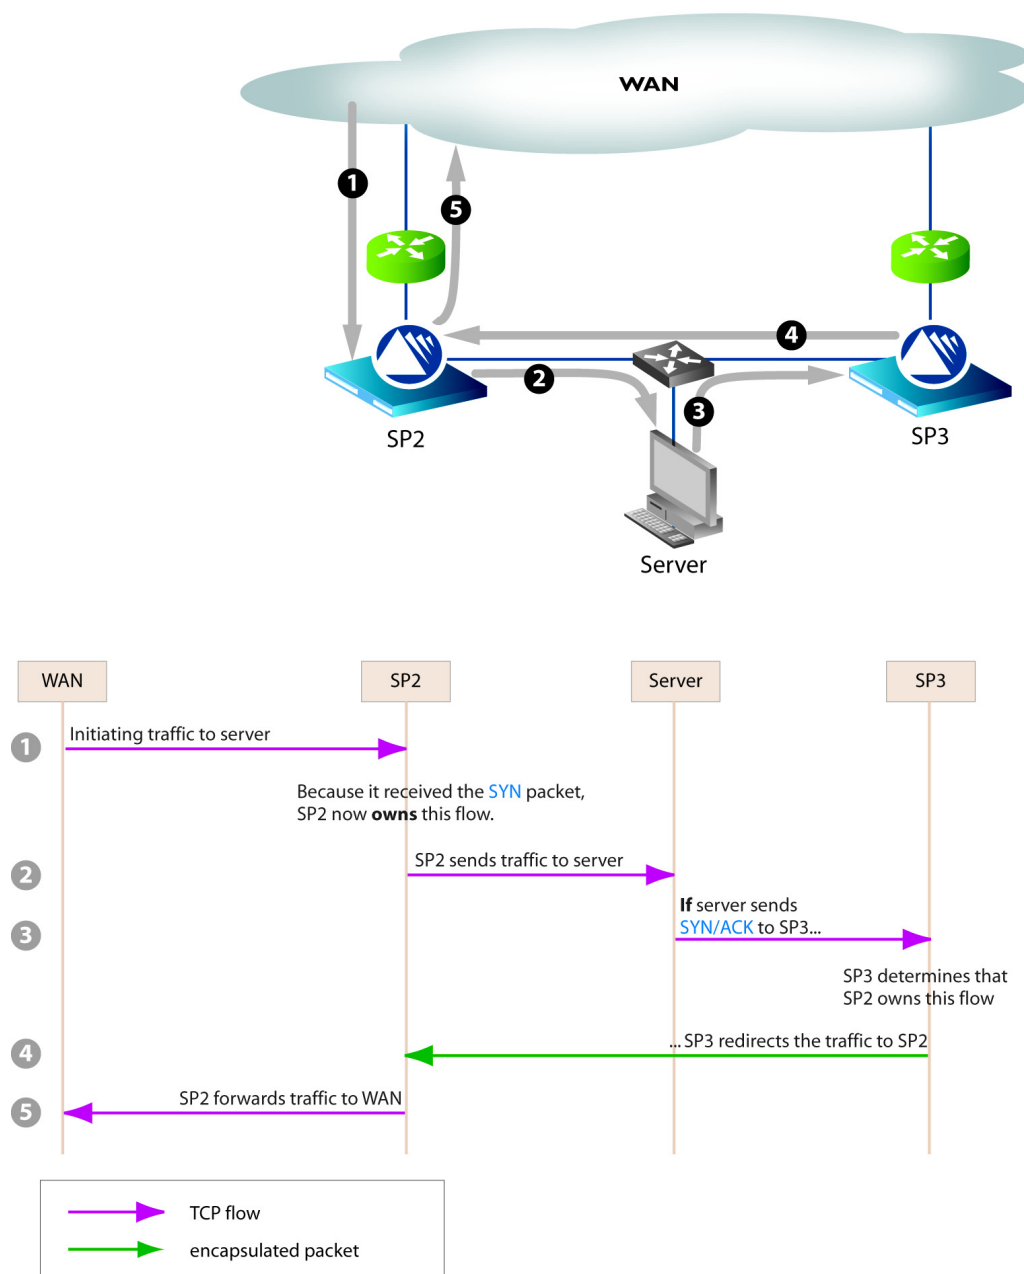
The client request — in the form of an initiating SYN packet — may be received from the WAN side or the LAN side. This results in two possible redirection scenarios:

- **Redirection for WAN-initiated Traffic** See page 227.
- **Redirection for LAN-initiated Traffic** See page 228.

The assumption is that flow redirection happens across a LAN environment. Redirection across a WAN is not supported.

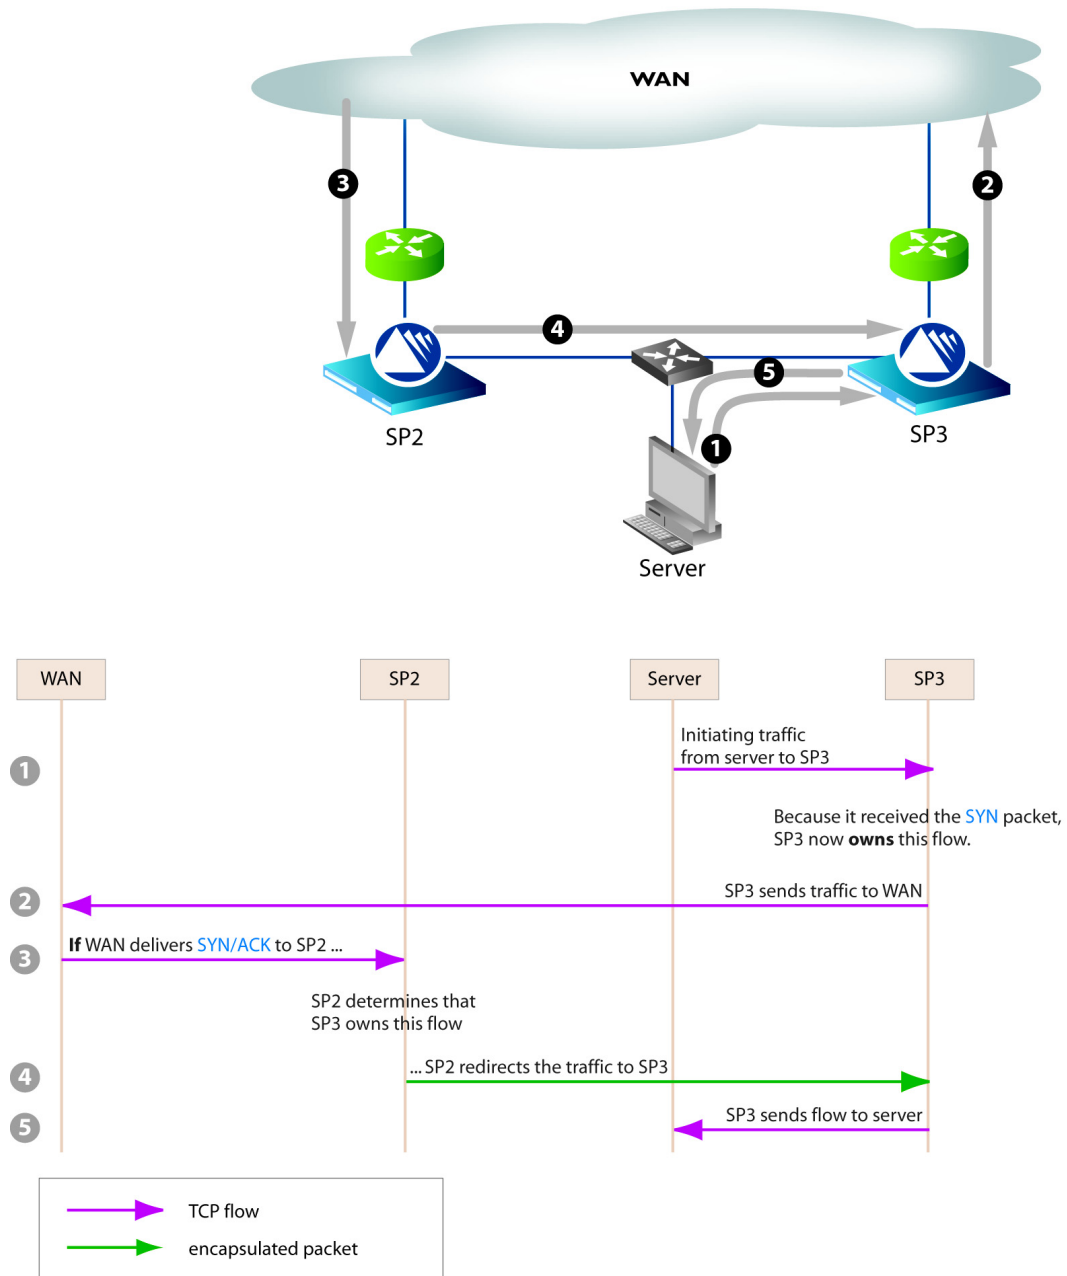
Redirection for WAN-initiated Traffic

In this scenario, the WAN initiates the flow. All traffic returned from the server is redirected to the appliance that first received traffic from the WAN.



Redirection for LAN-initiated Traffic

In this scenario, the LAN initiates the flow. All traffic returned from the WAN is redirected to the appliance that first received the traffic from the LAN.



Configuring Flow Redirection

If you have one path from the client to the server and a different path from the server to the client, you need to enable flow redirection and configure the appliances to communicate with each other.

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

- A cluster may 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 the 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, the process of configuring flow redirection uses three of the Appliance Manager's pages:

- The **Configuration - Interfaces** page, for configuring the **mgmt1** IP address.
- The **Configuration - IP Route** page, for configuring the necessary static route(s).
- The **Configuration - Flow Redirection** page, 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 need to be in a separate subnet from the **mgmt0** interfaces.



Tip Typically, you'll use the **mgmt1** interface. However, when the LAN-side is greater than 1 Gbps and your Silver Peak appliance has a 10-Gbps interface, then you may consider using a 10-Gbps interface for flow direction. In that case, contact Support for assistance.


Following is the complete example:

- **Example #1: Simple Cluster with Two Physically Connected Peers** See page 230.

Because of their physical proximity, a crossover cable connects two peers' **mgmt1** interfaces.

Example #1: Simple Cluster with Two Physically Connected Peers

When you want to cluster two appliances that are in the same subnet (with Layer 2 connectivity), and they're located in the same room, you can physically cable the two **mgmt1** interfaces together, in lieu of setting up an IP static route.



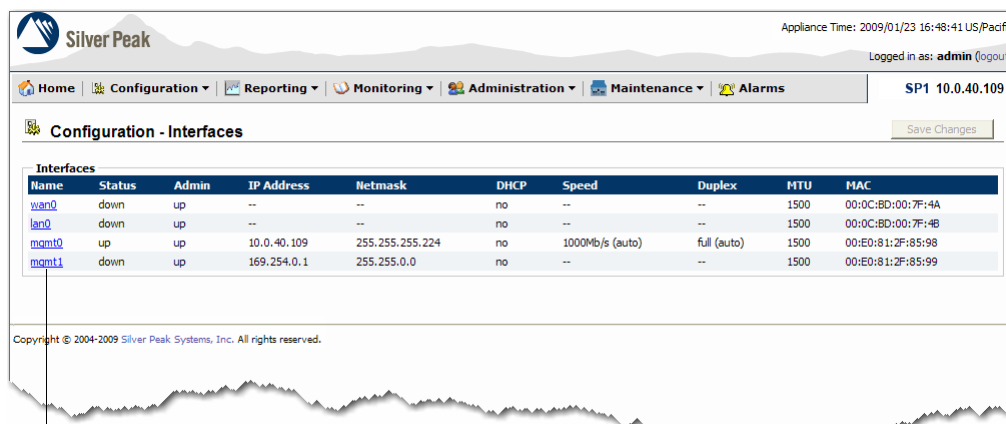
Instead of physically cabling the appliances, you also have the option of connecting the **mgmt1** interfaces via the local area network.



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

◆ To configure this scenario

- 1 Using a crossover cable, physically connect one appliance's **mgmt1** port to the other appliance's **mgmt1** port.
- 2 From **SP1**'s **Configuration** menu, select **Networking** and then **Interfaces**. The **Configuration - Interfaces** page appears.



The **mgmt1** interface shipped with a default IP address, to make initial configuration easy. You don't need this any longer, so we'll reconfigure it to use as a cluster interface for flow redirection.

- a Click **mgmt1**. The **Configure interface mgmt1** area appears.

Configuration - Interfaces

Name	Status	Admin	IP Address	Netmask	DHCP	Speed	Duplex	MTU	MAC
wan0	down	up	--	--	no	--	--	1500	00:0C:BD:00:7F:4A
lan0	down	up	--	--	no	--	--	1500	00:0C:BD:00:7F:4B
mgmt0	up	up	10.0.40.109	255.255.255.224	no	1000Mb/s (auto)	full (auto)	1500	00:E0:81:2F:85:98
mgmt1	down	up	169.254.0.1	255.255.0.0	no	--	--	1500	00:E0:81:2F:85:99

Configure interface mgmt1

Admin: Speed: Duplex:

☐ Automatically obtain IP Address settings with DHCP

☒ Statically configure IP address settings

IP address:

- b Configure **SP1** to have the **mgmt1** IP address, **10.10.10.1/24**.

Configure interface mgmt1

Admin: Speed: Duplex:

☐ Automatically obtain IP Address settings with DHCP

☒ Statically configure IP address settings

IP address:

- c Click **Apply**. The new **mgmt1** IP address appears in the table.

Configuration - Interfaces

Name	Status	Admin	IP Address	Netmask	DHCP	Speed	Duplex	MTU	MAC
wan0	down	up	--	--	no	--	--	1500	00:0C:BD:00:7F:4A
lan0	down	up	--	--	no	--	--	1500	00:0C:BD:00:7F:4B
mgmt0	up	up	10.0.40.109	255.255.255.224	no	1000Mb/s (auto)	full (auto)	1500	00:E0:81:2F:85:98
mgmt1	down	up	10.10.10.1	255.255.255.0	no	--	--	1500	00:E0:81:2F:85:99

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

To complete the appliance's configuration, you'll enable flow redirection and specify the other peer in the cluster.

- 3 From the **Configuration** menu, select **Networking > Flow Redirection**. The **Configuration - Flow Redirection** page appears.

Configuration - Flow Redirection

Settings

Enable ☐

Interface: mgmt1

Apply Cancel

Peers

Peer IP	State	Flow Redirection
---------	-------	------------------

Remove Selected Add

In the **Settings** area:

- a Click **Enable**.
 - b Verify that the default **Interface** is **mgmt1**.
 - c Click **Apply**.
- 4 In the **Peers** area, click **Add**. The **Add Peer** area appears.

Configuration - Flow Redirection

Settings

Enable ☒

Interface: mgmt1

Apply Cancel

Peers

Peer IP	State	Flow Redirection
---------	-------	------------------

Remove Selected Add

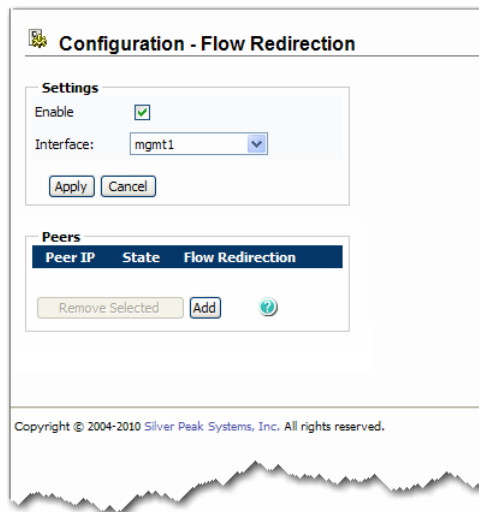
Add Peer

Peer IP: 10.10.10.2

Apply Cancel

Copyright © 2004-2010 Silver Peak Systems, Inc. All rights reserved.

- a Enter the **mgmt1** IP address for the appliance, **SP2**. Here, it's **10.10.10.2**.
- b Click **Apply**. The **Peers** table appears, displaying the **SP2**'s **mgmt1** interface IP address.



Configuration - Flow Redirection

Settings

Enable ☒

Interface: mgmt1

Apply Cancel

Peers

Peer IP	State	Flow Redirection
---------	-------	------------------

Remove Selected Add

Copyright © 2004-2010 Silver Peak Systems, Inc. All rights reserved.

- 5 Click **Save Changes**.

Now, repeat the entire procedure for the other appliance. When that's complete, both cluster interfaces are able to communicate with each other.

Flow Reporting

For flow redirection, the appliance handles flow reporting as follows:

- The reporting mechanism of locally owned flows is unchanged.
- When a peer redirects a flow, it does no per-flow reporting. Rather, that traffic's statistics are maintained by the owner of the flow. That is, the peer appliance to which the flow is redirected.
- On the **Monitoring - Current Flows** page, a flow's **Details** link has a field named, **Flow Redirected From** which displays which peer appliance IP is redirecting the flow to this appliance. This field only has an entry if the appliance owns the flow.

For more information, see “[Viewing Current Flows](#)” on page 273.

- The **Monitoring - Flow Redirection** page lists the **mgmt1** IP address for each peer in the cluster. It also displays statistics for the control packet, number of current flows redirected to/from the appliance, and a cumulative tally of the packets and bytes redirected to/from peers.

For more information, see “[Viewing Flow Redirection Statistics](#)” on page 302.



Configuring and Managing VLANs

This chapter describes how to configure and manage VLANs. This applies only when the appliance is in Bridge mode.

In This Chapter

- **Introduction** See page 236.
- **Configuring a VLAN IP Interface** See page 237.
- **Setting VLAN Tags in Outgoing WAN-side Packets** See page 238.

Introduction

The NX appliances provide full support for deployments on 802.1q trunked links.

By default, in a tagged environment (that is, an environment with VLANs), the NX appliance behaves as follows:

- optimizes all traffic according to existing Route policies, independent of VLAN
- delivers LAN-side packets untagged

You need to create a VLAN IP interface(s) on the appliance if you want to do any of the following:

- deliver tagged LAN-side packets. Any exploded LAN-side packets destined to VLANs for which the NX appliance has an interface will be tagged.

For more information, see “Configuring a VLAN IP Interface” on page 237.

- (**optionally**) set VLAN tags in outgoing WAN-side packets

For more information, see “Setting VLAN Tags in Outgoing WAN-side Packets” on page 238.

- (**optionally**) use VLAN tags as discriminators in MATCH criteria for Route/QoS/Optimization maps or Access Lists

For more information, see “Configuring MATCH Criteria in a Map or Policy” on page 131.

- view VLAN information for **Monitoring - Current Flows**

For more information, see “Current Flow Details” on page 281.

Configuring a VLAN IP Interface

You can only create a VLAN IP interface if the appliance is in Bridge mode.

◆ To configure a VLAN IP interface

- 1 On the **Configuration - System** page, verify that the appliance is in Bridge mode. There are 4 possible deployments.
- 2 From the **Configuration** menu, select **Networking > VLAN**. The **Configuration - VLAN** page appears.
- 3 Click **Add** and complete the fields.

For example, in the standard 4-port configuration below:

In Bridge mode, **bvi0** is the logical interface corresponding to **lan0**.

These are the router (WAN-side) next-hops. If deployed in bonded 4-port bridge mode, then this is labeled, **bwan0 Next-hop IP**.

VLAN	Interface.VLAN	IP Netmask	wan0 Next-hop IP
<input type="checkbox"/>	bvi0.100	2.2.2.2/24	2.2.2.5
<input type="checkbox"/>	bvi0.200	50.8.80.200/24	50.8.80.35
<input type="checkbox"/>	bvi0.80	50.8.80.0/24	50.8.80.31
<input type="checkbox"/>	bvi1.180	50.9.180.0/24	50.9.180.51
<input type="checkbox"/>	bvi1.200	50.9.200.0/24	50.9.200.61

Remove Selected Add

Add VLAN

Interface.VLAN: (1..4095)

IP Netmask: /

wan0 Next-hop IP:

Apply Cancel

IP address of the Silver Peak appliance in this VLAN.

IP address of the WAN router in this VLAN.

- 4 Click **Apply**. Now all LAN-side packets destined for subnets in the new VLAN will be tagged.

Setting VLAN Tags in Outgoing WAN-side Packets

This is an optional task.

- ◆ **To set VLAN tags in outgoing WAN-side tunnel traffic**

- 1 Select **Configuration > Networking > IP Routes**.

- 2 Choose which **bvi** interface to set as the system next-hop. You can only choose one.

If you tag the “native” VLAN, add a static route on the WAN-side router directing all incoming traffic destined for the Silver Peak appliance IP to one of the VLAN IP interfaces.

The following pages display VLAN configuration in the context of the possible bridge deployments.

Shown for each bridge deployment are the **Configuration - System** page, the **Configuration - VLAN** page, and the **Configuration - IP Routes** page.

The deployments covered here are as follows:

- **2-Port Bridge** See page 239.
- **Standard 4-Port Bridge** See page 240.
- **Flat 4-Port Bridge** See page 241.
- **Bonded 4-Port Bridge** See page 242.

2-Port Bridge

This is the most common of the Bridge deployments.

A 2-port bridge deployment has:

- 1 WAN-side router next-hop
- 1 subnet

Configuration - System page

Deployment Mode

☐ Bonding

☐ Router

☒ Bridge

2-port

wan0

lan0

IP

Appliance IP / Netmask: 1.1.1.1 / 24

wan0 Next-hop IP: 1.1.1.2

lan0 Next-hop IP: 1.1.1.7 (optional)

☐ Propagate Link Down

Any VLAN deployed here has:

- a **bvi0** (logical bridge) interface
- 1 WAN-side router next-hop
- 1 subnet

Configuration - VLAN

VLAN

Interface.VLAN	IP Netmask	wan0 Next-hop IP
<input type="button" value="Remove Selected"/> <input type="button" value="Add"/>		

Add VLAN

Interface.VLAN: bvi0.200 (/1..4095)

IP Netmask: 10.10.10.0 / 24

wan0 Next-hop IP: 10.10.10.1

Configuration - VLAN

VLAN

Interface.VLAN	IP Netmask	wan0 Next-hop IP	wan1 Next-hop IP
<input type="checkbox"/> bvi0.100	10.10.10.0/24	10.10.10.1	10.10.10.2

Silver Peak

Appliance Time: 2009/08/22 15:46:37 US/Pacific Version: 3.1.0.0_27699

Logged in as: admin (logout)

Tallinn2 10.0.40.36

Configuration - IP Routes

Save Changes

Management

Destination IP Netmask	Next-hop IP	Interface
<input type="checkbox"/> default	10.0.40.33	mgmt0
<input type="checkbox"/> 10.0.40.32/27	0.0.0.0	mgmt0
<input type="checkbox"/> 169.254.0.0/16	0.0.0.0	mgmt1

WAN

Next-hop IP	Interface	Configured Role	Current Role
<input type="radio"/> 1.1.1.2	bvi0	active	down
<input checked="" type="radio"/> 10.10.10.1	bvi0.200	n/a	n/a

LAN

Destination IP Netmask	Next-hop IP	Metric
<input type="checkbox"/> default	1.1.1.7	10

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

Selecting this line ensures that packets destined for the WAN are tagged with VLAN 200.



Note If the “native” VLAN is tagged, add a static route on the WAN-side router directing all incoming traffic destined for the Silver Peak appliance IP to one of the VLAN IP interfaces.

Standard 4-Port Bridge

This is also known as *dual bridge mode*.

A standard 4-port bridge deployment has:

- 2 WAN-side router next-hops
- 2 subnets

Configuration - System page

Any VLAN deployed here has:

- **bvi0** and **bvi1** (logical bridge) interfaces
- 2 WAN-side router next-hops
- 2 subnets

Select the desired interface from the drop-down list.

VLAN	Interface.VLAN	IP Netmask	WAN0 Next-hop IP
<input type="checkbox"/>	bvi0.444	10.10.10.0/24	10.10.10.1

Destination IP Netmask	Next-hop IP	Interface
<input type="checkbox"/> default	10.0.40.33	mgmt0
<input type="checkbox"/> 10.0.40.32/27	0.0.0.0	mgmt0
<input type="checkbox"/> 169.254.0.0/16	0.0.0.0	mgmt1

Next-hop IP	Interface	Configured Role	Current Role
1.1.1.2	bvi0	active	down
10.10.10.1	bvi0.444	active	n/a

Next-hop IP	Interface	Configured Role	Current Role
1.1.3.2	bvi1	active	down

By default, the interfaces display as **active**.
You can, however, set one selected WAN interface to **backup**.



Note If the “native” VLAN is tagged, add a static route on the WAN-side router directing all incoming traffic destined for the Silver Peak appliance IP to one of the VLAN IP interfaces.

Flat 4-Port Bridge

A flat 4-port bridge deployment has:

- 2 WAN-side router next-hops
- 1 subnet

Configuration - System page

The 'Configuration - System' page shows the 'Flat 4-port' deployment mode. The 'Bridge' option is selected under 'Deployment Mode'. The 'Advanced' tab is active. The 'Appliance IP / Netmask' is set to 1.1.1.1 / 24. The 'wan0 Next-hop IP' is 1.1.1.2, 'wan1 Next-hop IP' is 1.1.1.4, 'lan0 Next-hop IP' is 1.1.1.7 (optional), and 'lan1 Next-hop IP' is empty (optional). The 'Propagate Link Down' checkbox is unchecked.

Any VLAN deployed here has:

- a **bvi0** (logical bridge) interface
- 2 WAN-side router next-hops
- 1 subnet

The 'Configuration - VLAN' page shows the 'Add VLAN' section. The 'Interface.VLAN' is set to bvi0.100 (1..4095). The 'IP Netmask' is 10.10.10.0 / 24. The 'wan0 Next-hop IP' is 10.10.10.1 and the 'wan1 Next-hop IP' is 10.10.10.2. The 'Apply' button is highlighted.

Below, the 'Configuration - VLAN' page shows the added VLAN in the table:

Interface.VLAN	IP Netmask	wan0 Next-hop IP	wan1 Next-hop IP
bvi0.100	10.10.10.0/24	10.10.10.1	10.10.10.2

The 'Configuration - IP Routes' page shows the 'WAN' and 'LAN' route tables. The 'WAN' table has the following entries:

Next-hop IP	Interface	Configured Role	Current Role
1.1.1.2	bvi0	active	down
1.1.1.4	bvi0	active	down
10.10.10.1	bvi0.100	active	n/a
10.10.10.2	bvi0.100	active	active

The 'LAN' table has the following entries:

Destination IP Netmask	Next-hop IP	Metric
default	1.1.1.7	10

By default, the interfaces display as **active / active**. You can, however, set either **bvi0** to **backup**.

Since this is the native VLAN, packets exiting to the WAN are not tagged.



Note If the “native” VLAN is tagged, add a static route on the WAN-side router directing all incoming traffic destined for the Silver Peak appliance IP to one of the VLAN IP interfaces.

Bonded 4-Port Bridge

Essentially, this is like the 2-port bridge, but with twice the bandwidth.

A bonded 4-port bridge deployment has:

- 1 WAN-side router next-hops
- 1 subnet

Configuration - System page

Any VLAN deployed here has:

- a **bvi0** (logical bridge) interface
- 1 WAN-side router next-hop
- 1 subnet

Interface.VLAN	IP Netmask	bwan0 Next-hop IP
bvi0.600	10.10.10.0/24	10.10.10.4

Destination IP Netmask	Next-hop IP	Interface	Configured Role	Current Role
default	10.0.40.33	mgmt0		
10.0.40.32/27	0.0.0.0	mgmt0		
169.254.0.0/16	0.0.0.0	mgmt1		

Destination IP Netmask	Next-hop IP	Metric
default	1.1.1.7	10

Here, the user has chosen to change the default selection of native VLAN to VLAN 600 instead. Packets leaving for the WAN will be tagged accordingly.



Note If the “native” VLAN is tagged, add a static route on the WAN-side router directing all incoming traffic destined for the Silver Peak appliance IP to one of the VLAN IP interfaces.



Reporting Historical Traffic

This chapter describes how to create reports and view statistics collected over a specific interval for applications, data reduction, bandwidth optimization, flow counts, latency, and packet loss.

In This Chapter

- **Overview** See page 244.
- **About Viewing Statistics** See page 244.
- **Viewing Application Historical Statistics** See page 248.
- **Viewing Reduction Statistics** See page 251.
- **Viewing Bandwidth Statistics** See page 253.
- **Viewing Flow Counts** See page 255.
- **Viewing Latency** See page 257.
- **Viewing Network Integrity** See page 258.
- **Viewing a Summary of All Historical Reports** See page 261.

Overview

Tunnel historical statistics provide information about cumulative and minimum/maximum/average values. These are available for preselected time intervals: **Current Day**, **Last 24 Hours**, [the previous 3 days, as in **Tuesday**, **Wednesday**, **Thursday**, **Last 7 Days**, and **Last 30 Days**. The Appliance Manager knows what day it is and adjusts the menu field accordingly.

All in all, historical statistics provide an optimal method for assessing the impact of the Silver Peak appliances.

The Appliance Manager provides statistics for tunnels, the physical network interfaces, applications, and connections, as follows:

- **Viewing Application Historical Statistics** See page 248.

The application historical statistics summarize the distribution of applications as a percentage of tunnel or pass-through traffic over a selected period of time. Within these statistics, you can determine the percentage of WAN-to-LAN traffic, the ratio of traffic at each side of the appliance, and calculate how much the traffic was optimized.

- **Viewing Reduction Statistics** See page 251.

Reduction statistics specify the number of bytes and/or packets received, processed, and transmitted by a tunnel in both the outbound (LAN-to-WAN) and inbound (WAN-to-LAN) directions.

- **Viewing Bandwidth Statistics** See page 253.

Bandwidth statistics summarize the overall inbound and outbound bandwidth improvements afforded by the Silver Peak appliance.

- **Viewing Flow Counts** See page 255.

Flow count statistics summarize the number of TCP flows versus non-TCP flows, in terms of minimum, peak, and average number for each time interval selected.

- **Viewing Latency** See page 257.

Latency statistics summarize the round-trip time of data in a Silver Peak tunnel.

- **Viewing Network Integrity** See page 258.

These statistics summarize the number of lost and/or out-of-order packets, before and after enabling Forward Error Correction (FEC) and Packet Order Correction (POC), respectively.

- **Viewing a Summary of All Historical Reports** See page 261.

Based on the filters selected, this page displays a set of all applicable historical reports.

Before discussing individual reports, the next section describes the basics of viewing reports.


About Viewing Statistics

This section discusses methods for viewing additional details about report charts and graphs. It includes:

- **Understanding Traffic Direction** See page 245.
- **Viewing Pie Charts** See page 246.
- **Selecting Time Periods** See page 247.
- **Exporting Table Data** See page 247.

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:



- **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"

Viewing Pie Charts

The Appliance Manager allows you to change the view of each pie chart independently.

♦ To access pie chart view options


- 1 Right-click in the quadrant of the pie chart. The contextual menu appears.
- 2 Select the option you want. You can choose options sequentially for a cumulative effect.




These two commands — **Enable Rotation** and **Enable Slicing Movement** — are mutually exclusive:

- By default, **Enable Slicing Movement** is active, and it's greyed out because you don't need to click on it.
- **Enable Rotation** is accessible because you can change to that state.



You can change back and forth between the two, building up a cumulative effect. If you want to return to the default state and view, just refresh the browser.



Enable Slicing Movement — to move a slice, click it or its label. To move it back, click again.



Enable Slicing Movement and **Enable Rotation** used sequentially.



Selecting Time Periods

Many reports let you filter data by choosing from predefined time periods. The period you choose defines the start time, duration, and type of interval.

Period	Description
Current Day	Hourly summaries beginning at midnight this morning
Last 24 Hours	Hourly summaries for the 24-hour period beginning 24 hours ago
[2 Days Ago]	Hourly summaries for the 24-hour period beginning on the named weekday 2 days ago
[3 Days Ago]	Hourly summaries for the 24-hour period beginning on the named weekday 3 days ago
[4 Days Ago]	Hourly summaries for the 24-hour period beginning on the named weekday 4 days ago
Last 7 Days	Daily summaries for the period beginning 7 full days ago
Last 30 Days	Daily summaries for the period beginning 30 full days ago

Exporting Table Data

Most reports provide access to the data with a **Table View** button. If they do, the **Export** button saves the statistics tables to a **.csv** file.

- Under the **Reporting** menu, you can export statistics for the following individual reports:
 - **Applications**
 - **Reduction**
 - **Bandwidth**
 - **Flow Counts**
 - **Latency**
 - **Network Integrity**
 - **Forward Error Correction**
 - **Packet Order Correction**
- Under the **Monitoring** menu, you can export statistics for the following reports. Only the **QoS** report doesn't have a table view.:
 - **Applications**
 - **Current Flows**
 - **Reduction**
 - **Bandwidth**
 - **Flow Counts**
 - **Latency**
 - **Network Integrity**
 - **Forward Error Correction**
 - **Packet Order Correction**

♦ To export a **.csv** file

- 1 In the report, click **Table View**. The table appears.
- 2 Click **Export** and either **Open** or **Save** the file, as needed.

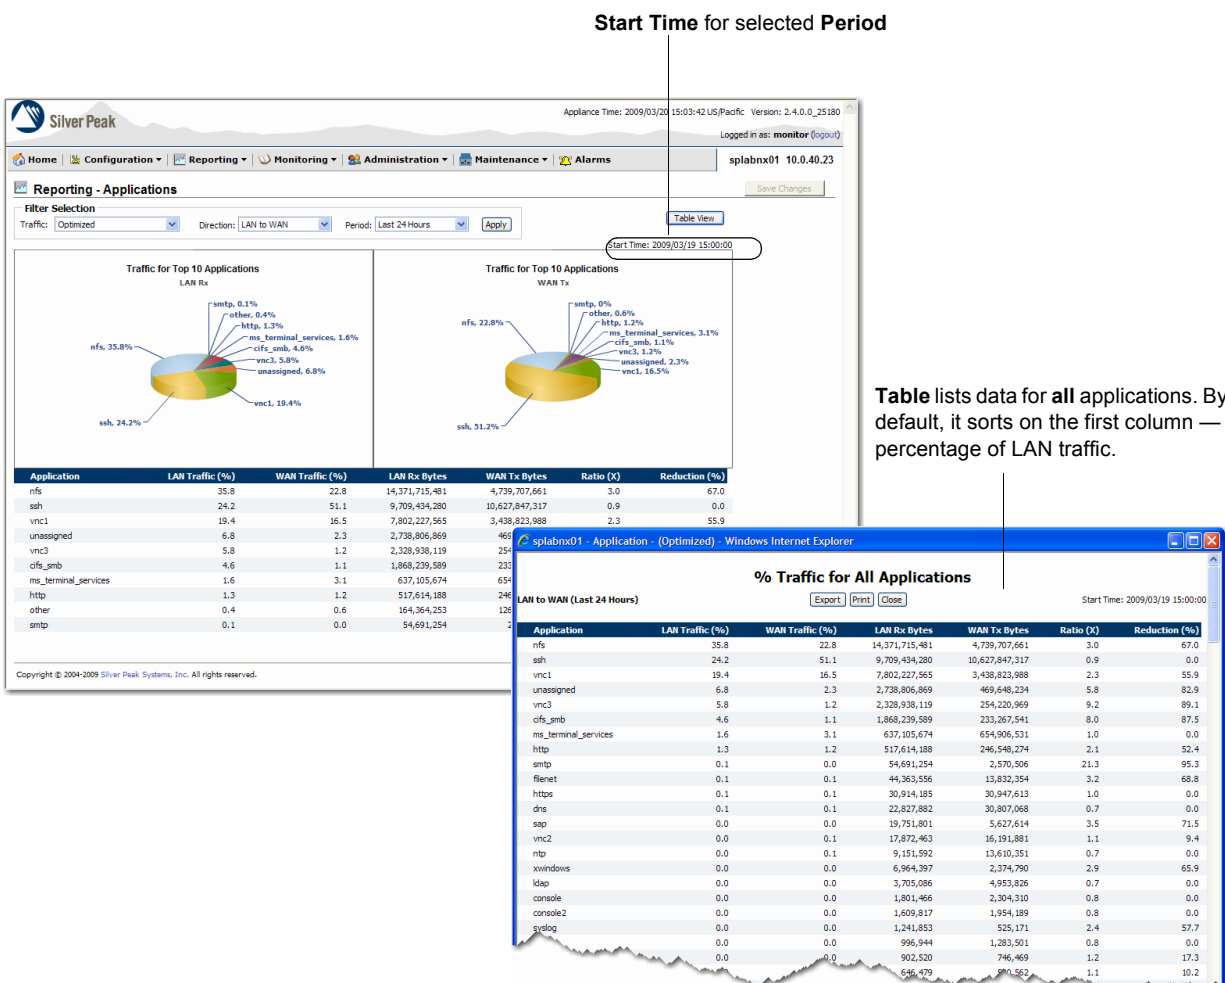
Viewing Application Historical Statistics

The **Reporting - Applications** page summarizes the distribution of applications as a percentage of traffic over a selected period of time. You can select for traffic that is optimized, tunnel-specific, or pass-through (shaped or unshaped).

It answers the following questions for a selectable time period:

- For all tunnels combined, what 10 applications account for the majority of traffic?
- For any given tunnel, what 10 applications account for the majority of traffic?
- What is the distribution of pass-through traffic (shaped, unshaped) across the top 10 applications?
- What application accounts for the largest percentage of my tunnel or pass-through traffic (shaped or unshaped)?
- Which applications account for more than $n\%$ of my bandwidth usage?
- When comparing the WAN- and LAN-side traffic, how are the application distributions different?
- Based on bytes, what is the ratio of LAN-to-WAN or WAN-to-LAN traffic for any given application?

A Sampling of Results



What Data Displays

The **Monitoring - Application Historical Stats** report displays the following statistics:

Table 12-1 Traffic Direction: LAN to WAN

Column	Definition
LAN Traffic (%)	Percentage of LAN traffic that this application comprises
WAN Traffic (%)	Percentage of WAN traffic that this application comprises
LAN Rx Bytes	Number of bytes received from the LAN side
WAN Tx Bytes	Number of bytes transmitted to the WAN side
Ratio (X)	$\frac{[\text{Bytes received from LAN}]}{[\text{Bytes transmitted to WAN}]}$
Reduction (%)	$\frac{[\text{Bytes received from LAN}] - [\text{Bytes transmitted to WAN}]}{\text{Bytes received from LAN}}$

Table 12-2 Traffic Direction: WAN to LAN

Column	Definition
LAN Traffic (%)	Percentage of LAN traffic that this application comprises
WAN Traffic (%)	Percentage of WAN traffic that this application comprises
LAN Tx Bytes	Number of bytes transmitted to the LAN side
WAN Rx Bytes	Number of bytes received from the WAN side
Ratio (X)	$\frac{[\text{Bytes transmitted to LAN}]}{[\text{Bytes received from WAN}]}$
Reduction (%)	$\frac{[\text{Bytes transmitted to LAN}] - [\text{Bytes received from WAN}]}{\text{Bytes transmitted to LAN}}$

♦ To view an Application (historical) report

- 1 From the **Reporting** menu, select **Applications**. Initially, the chart defaults to the last 24 hours' optimized LAN-to-WAN traffic.
- 2 In the **Filter Selection** section:
 - a Select the **Traffic** type. Options in the drop-down menu include:
 - **Optimized** – the sum of all optimized traffic. That is, all tunneled traffic.
 - individual tunnel name(s), listed alphabetically
 - **pass-through** for shaped, unoptimized traffic
 - **pass-through-unshaped** for unshaped, unoptimized traffic

- b** Select the **Direction** of the traffic. Options in the drop-down menu include:
 - **LAN to WAN** [default]
 - **WAN to LAN**
 - c** Select the time **Period**. Appliance Manager knows what day it is, so if today were Friday, then the choices would include **Current Day**, **Last 24 Hours**, **Thursday**, **Wednesday**, **Tuesday**, **Last 7 Days**, or **Last 30 Days**. The default is **Last 24 Hours**.
 - 3** Click **Apply** to update and display the chart.
 - 4** To view the table for **% Traffic for All Applications**, click **Table View**. This lists all applications, not just the top ten.

Viewing Reduction Statistics

The **Reporting - Reduction** page summarizes the reduction of (data) bytes afforded by the Silver Peak appliance.

The **Reduction** reports answer the following questions:

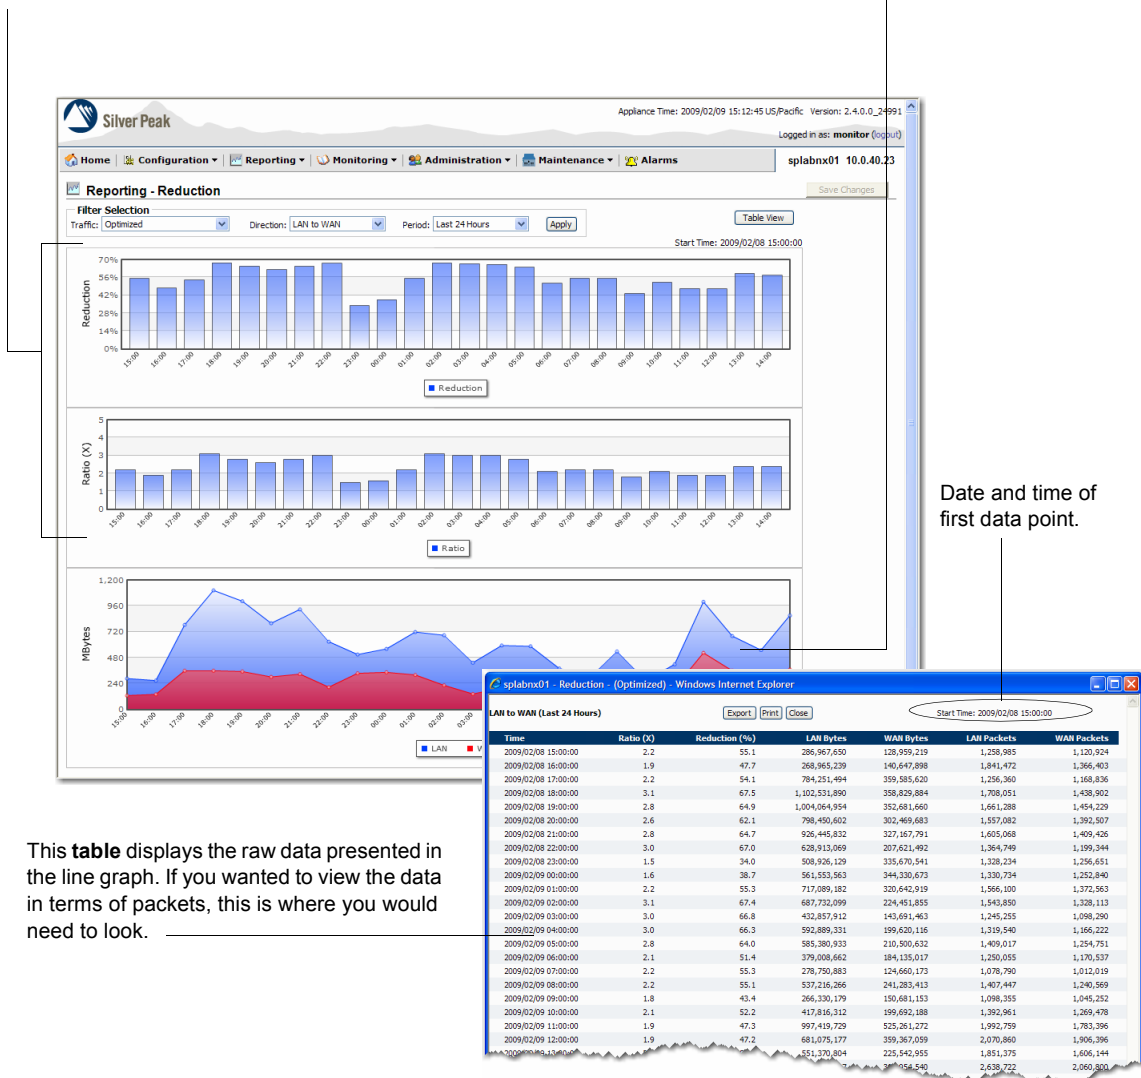
- How much data was sent and received in each time interval?
- What is the percent data reduction?
- What is the ratio of LAN to WAN (or WAN to LAN) traffic at any point in time?

A Sampling of Results

Both the **percent reduction** (shown in the top graph) and the **reduction ratio** (shown in the middle graph) are calculated from the megabytes data, shown in the bottom graph.

In the LAN to WAN direction, the ratio is calculated from: $[\text{LAN value}] / [\text{WAN value}]$.

This **line graph** shows the raw data, in megabytes, for all **optimized traffic**, for the last 24 hours. The direction of traffic is LAN to WAN.



This **table** displays the raw data presented in the line graph. If you wanted to view the data in terms of packets, this is where you would need to look.

♦ **To view a Reduction report**

- 1 From the **Reporting** menu, select **Reduction**. Initially, the chart defaults to the last 24 hours' optimized LAN-to-WAN traffic.
- 2 In the **Filter Selection** section:
 - a Select the **Traffic** type. Options in the drop-down menu include:
 - **Optimized** – the sum of all optimized traffic. That is, all tunneled traffic.
 - individual tunnel name(s), listed alphabetically
 - b Select the **Direction** of the traffic. Options in the drop-down menu include:
 - **LAN to WAN** [default]
 - **WAN to LAN**
 - **Bi-directional**
 - c Select the time **Period**. Appliance Manager knows what day it is, so if today were Friday, then the choices would include **Current Day**, **Last 24 Hours**, **Thursday**, **Wednesday**, **Tuesday**, **Last 7 Days**, or **Last 30 Days**. The default is **Last 24 Hours**.
- 3 Click **Apply** to update and display the chart.
- 4 To view the data displayed as a table, click **Table View**.

Viewing Bandwidth Statistics


The **Reporting - Bandwidth** page summarizes the overall inbound and outbound bandwidth improvements afforded by the Silver Peak appliance.

The **Bandwidth** report answers the following questions:


- How much has the bandwidth been optimized?
- At what rate was the data sent and received in each time interval?
- What is the ratio of LAN to WAN traffic at any point in time?
- What were the peak, average, and minimum data transfer rates?

This **LAN WAN bps** view shows the actual data rates — and the ratio of rates — for optimized traffic, for a 24-hour period beginning 4 days ago — on **Wednesday**. The direction of traffic is LAN to WAN.

By default, the charts display in kilobits or megabits per second. To view the data in terms of packets per second, view the table.



For the same time period and direction, this **LAN Min Max Avg** view shows the range of rates and their average.



By default, the charts display in megabits per second. To view the data in terms of packets per second, view the table.

Time	Ratio (X)	Reduction (%)	LAN bps	WAN bps	LAN pps	WAN pps
2009/08/12 00:00:00	1.0	1.0	1,003	992	2	2
2009/08/12 01:00:00	1.0	1.6	1,030	1,013	2	2
2009/08/12 02:00:00	1.0	4.1	1,054	1,011	2	2
2009/08/12 03:00:00	1.0	3.4	1,516	1,465	2	2
2009/08/12 04:00:00	1.1	5.3	1,048	992	2	2
2009/08/12 05:00:00	1.0	2.0	1,017	997	2	2
2009/08/12 06:00:00	1.0	1.0	994	984	2	2
2009/08/12 07:00:00	1.0	1.6	1,009	993	2	2
2009/08/12 08:00:00	1.0	3.8	1,106	1,064	2	2
2009/08/12 09:00:00	1.2	15.6	1,425	1,203	3	2
2009/08/12 10:00:00	1.1	9.9	1,599	1,441	3	2
2009/08/12 11:00:00	1.1	7.5	1,389	1,284	3	2
2009/08/12 12:00:00	1.1	6.4	1,170	1,095	3	2
2009/08/12 13:00:00	1.1	7.2	1,247	1,157	3	2
2009/08/12 14:00:00	1.1	6.5	1,802	1,685	3	2
2009/08/12 15:00:00	1.0	4.1	1,897	1,820	3	2
2009/08/12 16:00:00	1.2	14.3	2,356	2,020	3	2
2009/08/12 17:00:00	1.2	14.9	2,750	2,339	3	2
2009/08/12 18:00:00	0.9	0.0	1,180	1,259	3	2
2009/08/12 19:00:00	1.0	0.0	1,275	1,325	3	2
2009/08/12 20:00:00	1.0	0.0	1,275	1,299	3	2
2009/08/12 21:00:00	1.0	0.0	1,093	1,095	3	2
2009/08/12 22:00:00	1.0	2.4	1,044	1,020	2	2
2009/08/12 23:00:00	0.1	0.1	1,013	1,012	2	2

♦ **To view a Bandwidth optimization report**

- 1 From the **Reporting** menu, select **Bandwidth**. Initially, the chart defaults to the last 24 hours' optimized LAN-to-WAN traffic, expressed in megabits per second.

- 2 In the **Filter Selection** section:

- a Select the **Traffic** type. Options in the drop-down menu include:

- **Optimized** – the sum of all optimized traffic. That is, all tunneled traffic.
- individual tunnel name(s), listed alphabetically
- **pass-through** for shaped, unoptimized traffic
- **pass-through-unshaped** for unshaped, unoptimized traffic

- b In **Report Type**, select from the following:

LAN WAN bps	Transfer rates for LAN and WAN, in the selected direction. The chart option displays: <ul style="list-style-type: none"> • Ratio of LAN versus WAN transfer rates • LAN and WAN traffic rates in kilobits per second The charts adjust dynamically for bps, Kbps, Mbps, or Gbps.
LAN Min Max Avg	Minimum, average, and peak data rates for the LAN, charted in megabits per second (Mbps) and packets per second (pps)
WAN Min Max Avg	Minimum, average, and peak data rates for the WAN, charted separately in megabits per second (K bps) and packets per second (pps)

- c Select the **Direction** of the traffic. Options in the drop-down menu include:

- **LAN to WAN** [default]
- **WAN to LAN**
- **Bi-directional**

- d Select the time **Period**. Appliance Manager knows what day it is, so if today were Friday, then the choices would include **Current Day**, **Last 24 Hours**, **Thursday**, **Wednesday**, **Tuesday**, **Last 7 Days**, or **Last 30 Days**. The default is **Last 24 Hours**.

- 3 Click **Apply** to update and display the chart.

- 4 To view the data displayed as a table, click **Table View**.

Viewing Flow Counts


The **Reporting - Flow Counts** page summarizes the number of TCP flows versus non-TCP flows, in terms of minimum, peak, and average number for each time interval.

Within the TCP flows, it separates out how many flows were accelerated and how many were not. Since CIFS acceleration is a subset of TCP acceleration, that data is incorporated generically in the accelerated TCP flow data.

This report answers the following questions:

- How much of my traffic is TCP-based?
- How much of my TCP traffic is accelerated?

A Sampling of Results



This **line graph** displays the **Min/Max/Avg** values for **Optimized** traffic (that is, the cumulative data for all tunnels) over the last day.

Each type of **flow** displays as a separate chart. By definition, flows are considered bi-directional.

When you select **Table**, the Appliance Manager creates a separate data table for each flow type.

You can select the one you want from this line of links.

The **table** displays the raw data used to chart **Non TCP Min Max Avg**.

To export this table as a .csv (comma-separated values) file that you can open as an Excel spreadsheet, click **Export**. The **File Download** dialog box displays.

Time	Non TCP Min Max Avg	TCP Min Max Avg	TCP Accel Min Max Avg	TCP Non-Accel Min Max Avg
2009/03/19 17:00:00	852	961		1,099
2009/03/19 18:00:00	825	923		1,102
2009/03/19 19:00:00	722	864		1,064
2009/03/19 20:00:00	626	837		1,066
2009/03/19 21:00:00	735	907		1,101
2009/03/19 22:00:00	730	898		1,074
2009/03/20 00:00:00	651	824		1,044
2009/03/20 01:00:00	642	817		987
2009/03/20 02:00:00	754	890		990
2009/03/20 03:00:00	696	887		1,070
2009/03/20 04:00:00	656	864		1,060
2009/03/20 05:00:00	647	758		1,220
2009/03/20 06:00:00	587	746		1,024
2009/03/20 07:00:00	636	808		896
2009/03/20 08:00:00	662	807		1,021
2009/03/20 09:00:00	642	789		954
2009/03/20 10:00:00	569	762		965
2009/03/20 11:00:00	542	743		923
2009/03/20 12:00:00	603	829		900
2009/03/20 13:00:00	626	856		1,138
2009/03/20 14:00:00	636	943		1,162
2009/03/20 15:00:00	692			
2009/03/20 16:00:00				

♦ **To view a Flow Counts report**

- 1 From the **Reporting** menu, select **Flow Counts**. Initially, the chart defaults to the number of flows for optimized traffic over the last 24 hours.
- 2 In the **Filter Selection** section:
 - a Select the **Traffic** type. Options in the drop-down menu include:
 - **Optimized** – the sum of all optimized traffic. That is, all tunnelized traffic.
 - individual tunnel name(s), listed alphabetically
 - **pass-through** for shaped, unoptimized traffic
 - **pass-through-unshaped** for unshaped, unoptimized traffic
 - b From the **Period** menu, select the time period from which the report begins:

Appliance Manager knows what day it is, so if today were Friday, then the choices would include **Current Day**, **Last 24 Hours**, **Thursday**, **Wednesday**, **Tuesday**, **Last 7 Days**, or **Last 30 Days**. The default is **Last 24 Hours**.
- 3 Click **Apply** to update and display the chart.
- 4 To view the data displayed as a table, click **Table View**.

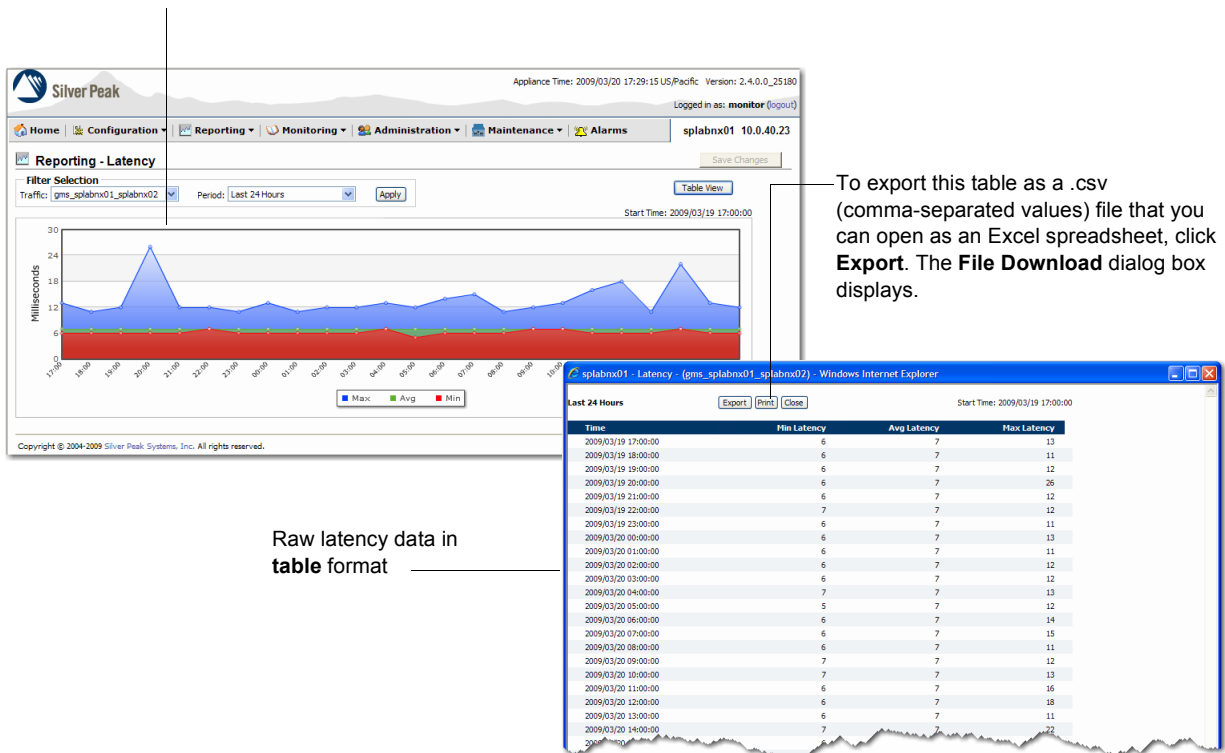
Viewing Latency

The **Reporting - Latency** page summarizes the round-trip time of data in a Silver Peak tunnel. It answers the following questions:

- How long does it take my data to get to the other end of the Silver Peak tunnel?
- What were the peak, average, and minimum time intervals?

A Sampling of Results

This **line graph** displays the minimum, average, and maximum round trip **latency** for tunnel **gms_splabnx01_splabnx02** for the last 24 hours.



♦ To view the Latency report

- 1 From the **Reporting** menu, select **Latency**. By default, the chart displays the last 24 hours' statistics for the first tunnel listed alphabetically.

By default, the **Report Type** is **Min Max Avg**. It displays the minimum, peak, and average latency for each time interval.

- 2 From the **Traffic** menu, select the name of the individual tunnel.
- 3 From the **Period** menu, select the time period from which the report begins:

Appliance Manager knows what day it is, so if today were Friday, then the choices would include **Current Day**, **Last 24 Hours**, **Thursday**, **Wednesday**, **Tuesday**, **Last 7 Days**, or **Last 30 Days**. The default is **Last 24 Hours**.

- 4 Click **Apply** to update and display the chart.
- 5 To view the data displayed as a table, click **Table View**.

Viewing Network Integrity

The **Reporting - Network Integrity** page summarizes the following on a per-tunnel basis:

- the number of packets lost before and after enabling Forward Error Correction (FEC)
- the number of out-of-order packets before and after enabling Packet Order Correction (POC)

It answers the following questions:

- How many errors were there before/after turning on Forward Error Correction?
- How many out-of-order packets were there before/after turning on Packet Order Correction?
- What were the peak, average, and minimum values?

A Sampling of Results

The top two **bar charts** display the average loss and peak loss, in percent, before and after enabling Forward Error Correction on tunnel **gms_splabnx01_splabnx02** during the last 7 days.

The lower two charts address average and peak out-of-order packets.



To export this table as a .csv (comma-separated values) file that you can open as an Excel spreadsheet, click **Export**. The **File Download** dialog box displays.

The raw FEC data in **table** form. There is a separate table for POC data.

splabnx01 - Reporting - Forward Error Correction (gms_splabnx01_splabnx02) - Windows Internet Explorer

Last 7 Days

[Pre Packet FEC Min Max Avg] [Post Packet FEC Min Max Avg]

Time	Avg Pre-FEC Loss (%)	Peak Pre-FEC Loss (%)	Min Pre-FEC Pkts Loss	Avg Pre-FEC Pkts Loss	Max Pre-FEC Pkts Loss	Avg Wan Rx Pkts
2009/03/13 00:00:00	0.000	0.000	0	0	0	34,677
2009/03/14 00:00:00	0.000	0.000	0	0	0	27,272
2009/03/15 00:00:00	0.000	0.000	0	0	0	28,063
2009/03/16 00:00:00	0.000	0.000	0	0	0	44,246
2009/03/17 00:00:00	0.000	0.000	0	0	0	60,921
2009/03/18 00:00:00	0.000	0.000	0	0	0	57,346
2009/03/19 00:00:00	0.003	0.000	0	0	737	53,833

♦ **To view the Network Integrity report**

- 1 From the **Reporting** menu, select **Network Integrity**. By default, the chart displays the last 24 hours' statistics for the first tunnel listed alphabetically.
- 2 From the **Traffic** menu, select the name of the individual tunnel.
- 3 From the **Period** menu, select the time period from which the report begins:
Appliance Manager knows what day it is, so if today were Friday, then the choices would include **Current Day**, **Last 24 Hours**, **Thursday**, **Wednesday**, **Tuesday**, **Last 7 Days**, or **Last 30 Days**. The default is **Last 24 Hours**.
- 4 Click **Apply** to update and display the chart.
- 5 To view the data displayed as a table, click **FEC Table View** or **POC Table View**, as desired.

In the **FEC Table View**, the Appliance Manager charts the following items:

Field	Definition
Avg Pre FEC (%)	Percentage of packets received from the WAN that had errors before FEC was applied
Min Pre Pkt FEC	Minimum number of packets received from the WAN that had errors before FEC was applied
Avg Pre Pkt FEC	Average number of packets received from the WAN that had errors before FEC was applied
Max Pre Pkt FEC	Maximum number of packets received from the WAN that had errors before FEC was applied
Avg Post FEC (%)	Percentage of packets received from the WAN having errors after FEC was applied
Min Post Pkt FEC	Minimum number of packets received from the WAN that had errors after FEC was applied
Avg Post Pkt FEC	Average number of packets received from the WAN that had errors after FEC was applied
Max Post Pkt FEC	Maximum number of packets received from the WAN that had errors after FEC was applied
Avg Wan Rx Pkts	Total number of WAN packets received



Tip If you enable **FEC** and see an increase in out-of-order packets (on the **Packet Order Correction** tab of the **Monitoring - Network Integrity** page), it indicates that you need to go back to the **Configuration - Tunnels** page and increase the **Reorder Wait** time.

In the **POC Table View**, Appliance Manager charts the following items:

Field	Definition
Avg Pre POC (%)	Percentage of packets received from the WAN that had errors before POC was applied
Min Pre Pkt POC	Minimum number of packets received from the WAN that had errors before POC was applied
Avg Pre Pkt POC	Average number of packets received from the WAN that had errors before POC was applied
Max Pre Pkt POC	Maximum number of packets received from the WAN that had errors before POC was applied

Field	Definition (Continued)
Avg Wan Rx Pkts	Total number of WAN packets received
Avg Post POC (%)	Percentage of packets received from the WAN having errors after POC was applied
Min Post Pkt POC	Minimum number of packets received from the WAN that had errors after POC was applied
Avg Post Pkt POC	Average number of packets received from the WAN that had errors after POC was applied
Max Post Pkt POC	Maximum number of packets received from the WAN that had errors after POC was applied
Avg Wan Rx Pkts	Average number of WAN packets received

For both the peak loss and the peak out-of-order packets, the Appliance Manager shows data at the 95th percentile. This serves to filter out non-representative spikes, for example, during reboots or network failures.

Viewing a Summary of All Historical Reports

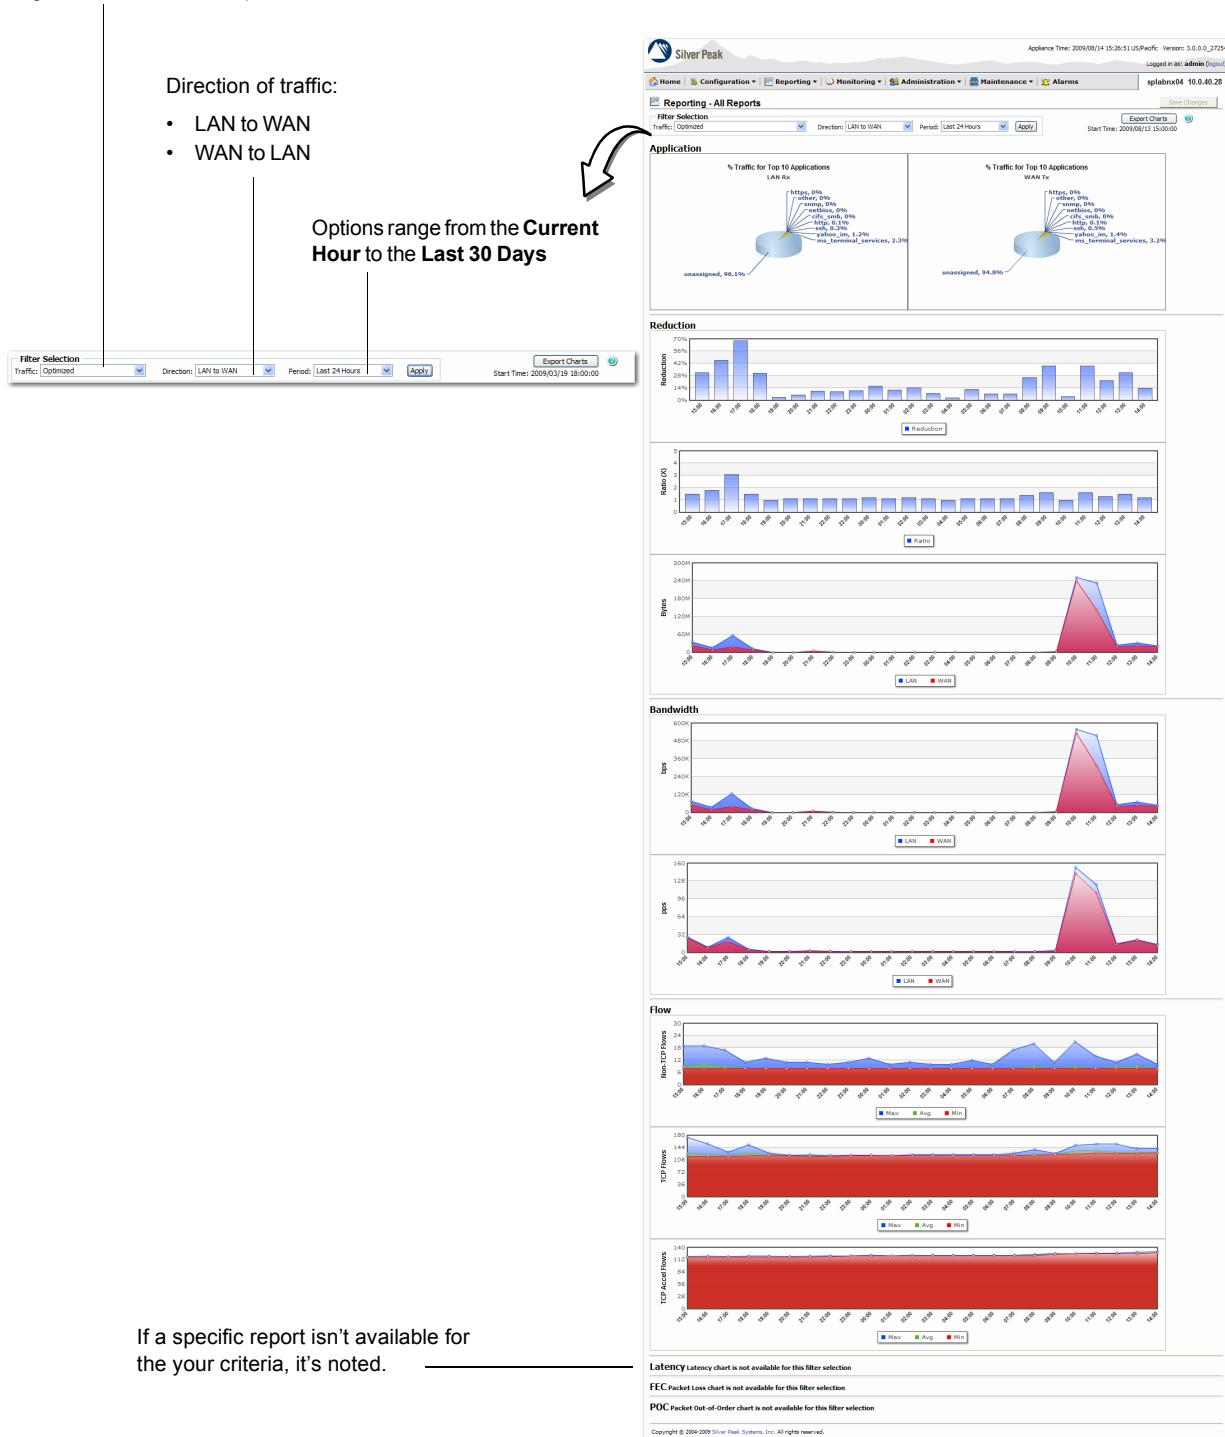
You can view a concatenated set of all the reports that meet your specified filter criteria.

Select all **Optimized** traffic, or specify a single tunnel from the drop-down list.

Direction of traffic:

- LAN to WAN
- WAN to LAN

Options range from the **Current Hour** to the **Last 30 Days**



If a specific report isn't available for your criteria, it's noted.

Exporting the report when using a 64-bit PC

When exporting this report, you need to choose the Microsoft Office Document Image Writer printer driver. However, if you're running Windows on a 64-bit machine, this driver isn't available.

If you have Acrobat Distiller installed, you can perform the following workaround, which results in a PDF file, instead of a .mdi file:

♦ To create a PDF file of your filtered results

- 1 From the browser's **File** menu, select **Print**. The **Print** dialog box appears.
- 2 Select **Print to File**. The **Print to File...** dialog box appears.
- 3 Name the file, adding the suffix, **.ps**. Then, from the **Save as type:** field, select **All Files (*.*)**.
- 4 Click **Save**.
- 5 Open Acrobat Distiller, navigate to the **.ps** file you created, and click **Open**. Distiller saves the file as a PDF file.



Monitoring Realtime Traffic

This chapter describes how to view realtime statistics. Generally, this includes the last hour's worth of collected data.

In This Chapter

- **Overview** See page 264.
- **About Viewing Statistics** See page 266.
- **Viewing Application Realtime Statistics** See page 270.
- **Viewing Current Flows** See page 273.
- **Viewing Tunnel QoS Statistics** See page 286.
- **Viewing Tunnel Realtime Statistics** See page 288.
- **Viewing Reduction Statistics** See page 293.
- **Viewing Bandwidth Statistics** See page 295.
- **Viewing Flow Counts** See page 297.
- **Viewing Latency Statistics** See page 299.
- **Viewing Network Integrity Statistics** See page 300.
- **Viewing Flow Redirection Statistics** See page 302.
- **Viewing NetFlow Statistics** See page 304.
- **Viewing Interface Statistics** See page 305.
- **Viewing Bridge Mode Statistics** See page 307.
- **Viewing IP Routes** See page 308.

Overview

With the **Monitoring** menus, the Appliance Manager provides the last hour's statistics for tunnels, the physical network interfaces, applications, and connections, as follows:

- **Viewing Application Realtime Statistics** See page 270.

Application realtime statistics summarize the percentage of traffic that can be accounted for, individually, by the "Top 10" applications.

Realtime statistics display the data accumulated since the last reboot. Additionally, you have the option of non-destructively clearing the counters to zero and viewing the delta values.

- **Viewing Current Flows** See page 273.

You can view a listing of 200 realtime connections, based on selectable filter criteria. These include end points, ports, traffic path, applications, protocol, Top Talkers, number of flows, and number of (mega)bytes. Additionally, you can customize which data columns display and view a flow's details.

- **Viewing Tunnel QoS Statistics** See page 286.

This is essentially a more detailed view of the tunnel realtime statistics (see next bullet). You can view the total number of bytes and packets transmitted and received, based on traffic class and/or WAN QoS [DSCP markings], for all tunnels or any individual tunnel. It also provides more granular information about types of packets, such as dropped, invalid, duplicate, etc.

Tunnel QoS statistics display the data accumulated since the last reboot. Additionally, you have the option of non-destructively clearing the counters to zero and viewing the delta values.

- **Viewing Tunnel Realtime Statistics** See page 288.

Tunnel realtime statistics specify the number of bytes and/or packets received, processed, and transmitted by a tunnel in both the outbound (LAN-to-WAN) and inbound (WAN-to-LAN) directions. They tally control packets, as well as accelerated versus non-accelerated traffic flow, round-trip latency, and packet loss before and after forward error correction. The Appliance Manager reports all these statistics as raw data.

Realtime statistics display the data accumulated since the last reboot. Additionally, you have the option of non-destructively clearing the counters to zero and viewing the delta values.

- **Viewing Reduction Statistics** See page 293.

Reduction statistics specify the number of bytes and/or packets received, processed, and transmitted by a tunnel in both the outbound (LAN-to-WAN) and inbound (WAN-to-LAN) directions.

- **Viewing Bandwidth Statistics** See page 295.

Bandwidth statistics summarize the overall inbound and outbound bandwidth improvements afforded by the Silver Peak appliance.

- **Viewing Flow Counts** See page 297.

This summarizes the number of TCP flows versus non-TCP flows. Within the TCP flows, it distinguishes between accelerated and non-accelerated flows.

- **Viewing Latency Statistics** See page 299.

Latency statistics summarize the round-trip time of data in a Silver Peak tunnel.

- **Viewing Network Integrity Statistics** See page 300.

These statistics summarize the number of lost and/or out-of-order packets, before and after enabling Forward Error Correction (FEC) and Packet Order Correction, respectively.

- **Viewing Flow Redirection Statistics** See page 302.

This shows the statistics collected, specific to the process when you allow two (or more) NX appliances to exchange flow ownership information and then redirect packets to the owner.

- **Viewing NetFlow Statistics** See page 304.

This displays how many NetFlow statistics the appliance exported to the collector(s). Stats are defined in terms of number of flows, and number of datagrams (packets) required to export those flows.

- **Viewing Interface Statistics** See page 305.

Interface statistics display generic performance data for the actual physical LAN, WAN, and management interfaces (primary and secondary).

Interface statistics display the data accumulated since the last reboot. Additionally, you have the option of non-destructively clearing the counters to zero and viewing the delta values.

- **Viewing Bridge Mode Statistics** See page 307.

This summarizes the data traffic traversing all the LAN and WAN interfaces, in a redundant bridge mode deployment.

- **Viewing IP Routes** See page 308.

This summarizes the next-hop reachability of the IP addresses, along with listing the source (WAN or LAN), state, uptime, WAN configured role, and WAN current role.

Before discussing individual reports, the next section describes the basics of viewing reports.


About Viewing Statistics

This section discusses methods for viewing additional details about report charts and graphs. It includes:

- **Understanding Traffic Direction** See page 266.
- **Viewing Counters Since Last Reboot** See page 266.
- **Clearing Counters Non-Destructively** See page 267.
- **Viewing Pie Charts** See page 268.
- **Exporting Table Data** See page 269.

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:



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

Viewing Counters Since Last Reboot

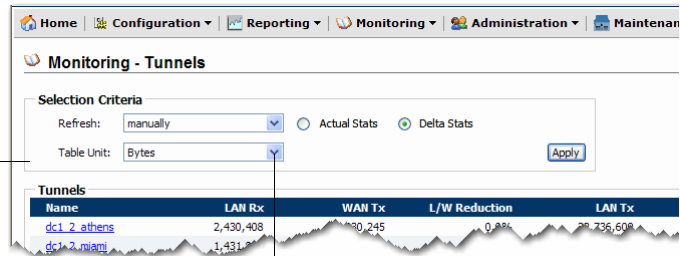
By default, the statistics that display in the following reports have accumulated since the last reboot: **Applications**, **Tunnel QoS**, **Tunnels**, **Flow Redirection**, **NetFlow**, and **Interfaces**.

To verify this, note that **Actual Stats** button is selected.

Clearing Counters Non-Destructively

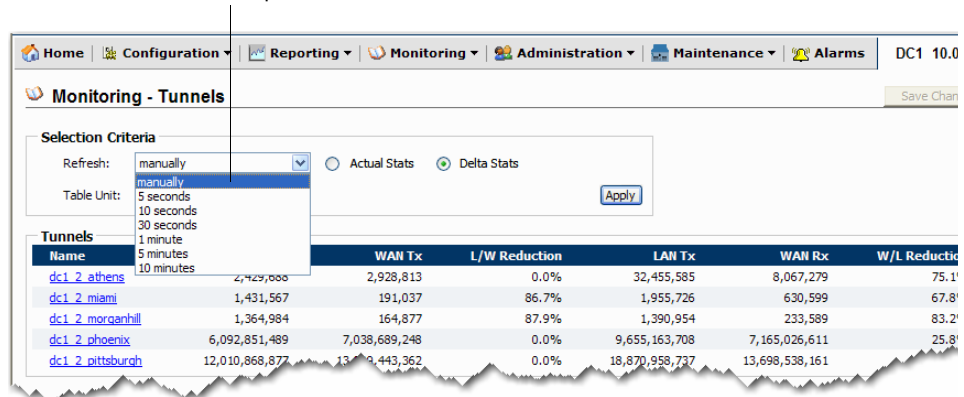
To non-destructively set the counters to zero, click **Delta Stats**. To update the statistics, you can manually refresh the page whenever you want. Or, you can select from one of the preconfigured auto-refresh intervals in the same menu.

To zero out counters non-destructively, select **Delta Stats**.



Select table's display units:
Bytes, MBytes, Pkts, or KPkts

If you set the **Refresh** menu to **manually**, click the browser's refresh icon as needed for a cumulative update.



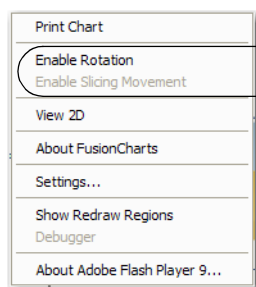
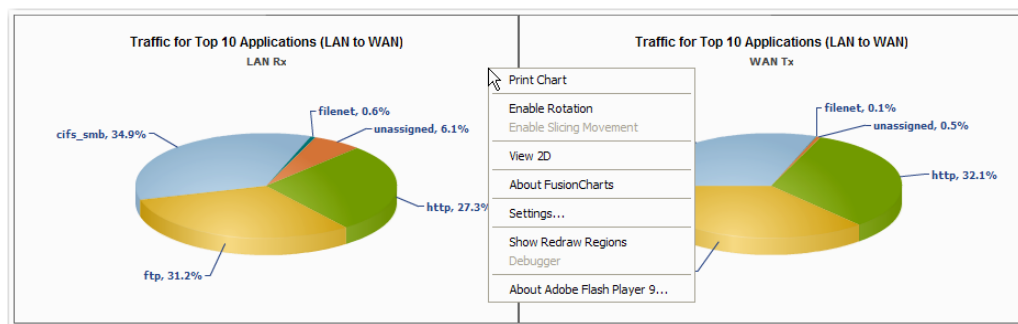
To restore the values since the last reboot, click **Actual Stats**.

Viewing Pie Charts

The Appliance Manager allows you to change the view of each pie chart independently.

♦ To access pie chart view options


- 1 Right-click in the quadrant of the pie chart. The contextual menu appears.
- 2 Select the option you want. You can choose options sequentially for a cumulative effect.




These two commands — **Enable Rotation** and **Enable Slicing Movement** — are mutually exclusive:

- By default, **Enable Slicing Movement** is active, and it's greyed out because you don't need to click on it.
- **Enable Rotation** is accessible because you can change to that state.


You can change back and forth between the two, building up a cumulative effect. If you want to return to the default state and view, just refresh the browser.




Enable Slicing Movement — to move a slice, click it. To move it back, click it again.



Enable Slicing Movement and **Enable Rotation** used sequentially.



Enable Rotation — to rotate the chart, click and drag to the left or right. Useful for repositioning when names are bunched or near the sides.



View 2D applied after slicing and rotating the original chart.

Exporting Table Data

Most reports provide access to the data with a **Table View** button. If they do, the **Export** button saves the statistics tables to a **.csv** file.

- Under the **Reporting** menu, you can export statistics for the following individual reports:
 - **Applications**
 - **Reduction**
 - **Bandwidth**
 - **Flow Counts**
 - **Latency**
 - **Network Integrity**
 - **Forward Error Correction**
 - **Packet Order Correction**
 - Under the **Monitoring** menu, you can export statistics for the following reports. Only the **QoS** report doesn't have a table view.:
 - **Applications**
 - **Current Flows**
 - **Reduction**
 - **Bandwidth**
 - **Flow Counts**
 - **Latency**
 - **Network Integrity**
 - **Forward Error Correction**
 - **Packet Order Correction**
- ♦ **To export a .csv file**
- 1 In the report, click **Table View**. The table appears.
 - 2 Click **Export** and either **Open** or **Save** the file, as needed.

Viewing Application Realtime Statistics

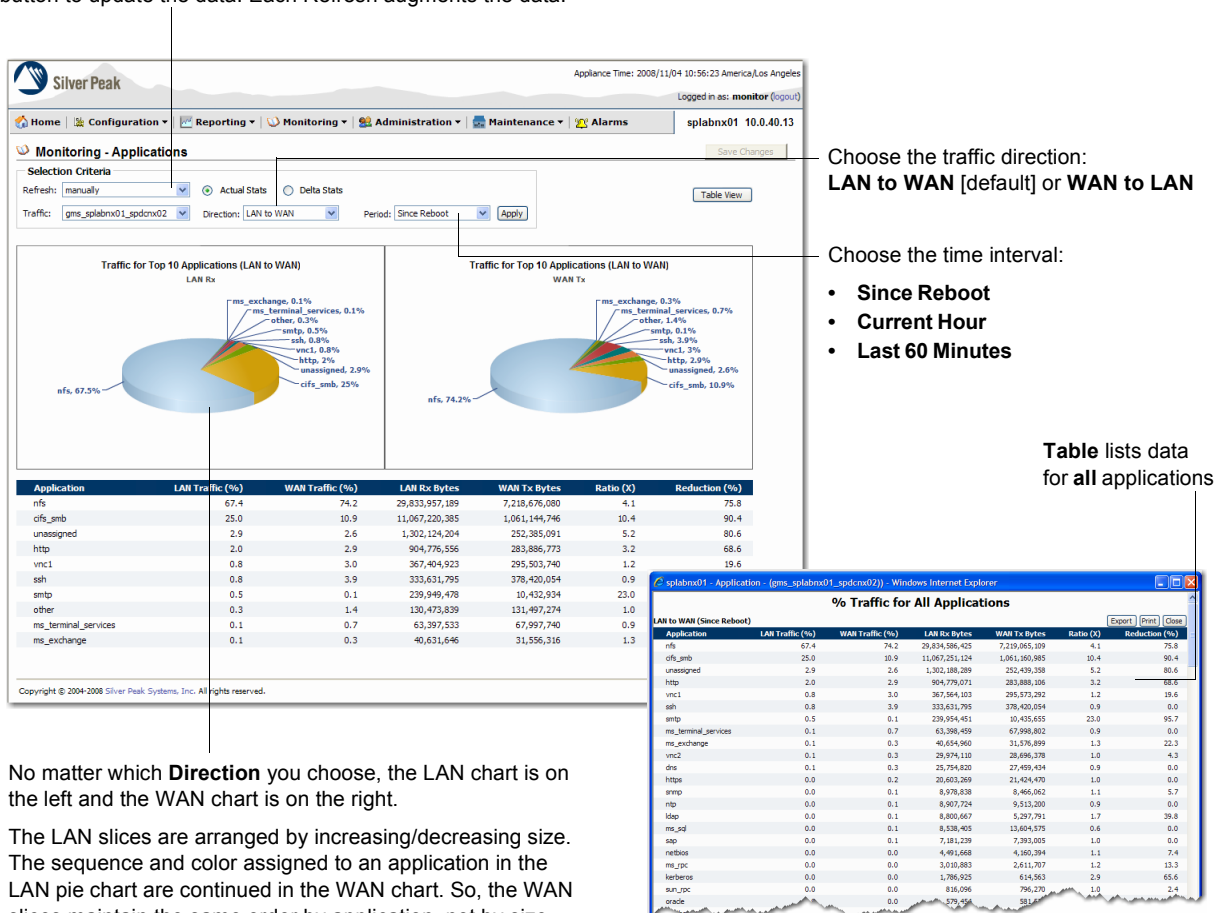
The **Monitoring - Application Stats** page summarizes the percentage of traffic that can be accounted for, individually, by the “Top 10” applications.

The reports answer the following questions:

- For any given tunnel, what 10 applications account for the majority of traffic since the last reboot?
- For any given tunnel, what 10 applications account for the majority of traffic since I non-destructively cleared the counters and refreshed the screen?
- What is the distribution of pass-through traffic (shaped, unshaped) across the top 10 applications since the last reboot?
- What is the distribution of pass-through traffic (shaped, unshaped) across the top 10 applications since I non-destructively cleared the counters and refreshed the screen?
- What application accounts for the largest percentage of my tunnel or pass-through traffic?
- Which applications account for more than $n\%$ of my bandwidth?
- When comparing the WAN- and LAN-side traffic, how are the application distributions different?
- Based on bytes, what is the ratio of LAN-to-WAN or WAN-to-LAN traffic for any given application?

A Sampling of Results

When **Refresh** is set to **manually**, use the browser's **Refresh** button to update the data. Each Refresh augments the data.



No matter which **Direction** you choose, the LAN chart is on the left and the WAN chart is on the right.

The LAN slices are arranged by increasing/decreasing size. The sequence and color assigned to an application in the LAN pie chart are continued in the WAN chart. So, the WAN slices maintain the same order by application, not by size. This allows for direct visual comparison.

What Data Displays

The **Monitoring - Applications** report displays the following statistics for each listed application.

Table 13-1 Traffic Direction: LAN to WAN

Column	Definition
LAN Traffic (%)	Percentage of LAN traffic that this application comprises
WAN Traffic (%)	Percentage of WAN traffic that this application comprises
LAN Rx Bytes	Number of bytes received from the LAN side
WAN Tx Bytes	Number of bytes transmitted to the WAN side
Ratio (X)	$\frac{[\text{Bytes received from LAN}]}{[\text{Bytes transmitted to WAN}]}$
Reduction (%)	$\frac{[\text{Bytes received from LAN}] - [\text{Bytes transmitted to WAN}]}{\text{Bytes received from LAN}}$

Table 13-2 Traffic Direction: WAN to LAN

Column	Definition
LAN Traffic (%)	Percentage of LAN traffic that this application comprises
WAN Traffic (%)	Percentage of WAN traffic that this application comprises
LAN Tx Bytes	Number of bytes transmitted to the LAN side
WAN Rx Bytes	Number of bytes received from the WAN side
Ratio (X)	$\frac{[\text{Bytes transmitted to LAN}]}{[\text{Bytes received from WAN}]}$
Reduction (%)	$\frac{[\text{Bytes transmitted to LAN}] - [\text{Bytes received from WAN}]}{\text{Bytes transmitted to LAN}}$

♦ To view the Application Statistics report

- 1 On the **Monitoring** menu, select **Applications**. Initially, the chart defaults to LAN-to-WAN traffic for the first alphabetically listed tunnel.
- 2 In the **Selection Criteria** section:
 - a Choose the page **Refresh** interval:
 - **If you want to refresh the screen manually**, select **manually** from the **Refresh** drop-down menu, and click **Apply** as needed to update the data.
 - **If you want the screen to refresh at a preselected time interval**, select one of the following from the **Refresh** drop-down menu: **30 seconds**, **1 minute**, **5 minutes**, or **10 minutes**.

- b** Select either the **Actual Stats** or the **Delta Stats** option.
 - **Actual Stats** displays the statistics that have accumulated since the last reboot.
 - **Delta Stats** non-destructively resets the counters to zero. Each refresh cumulatively increments the statistics since the counter reset.
 - c** Select the **Traffic** type. Options in the drop-down menu include:
 - individual tunnel name(s), listed alphabetically
 - **pass-through** for shaped, unoptimized traffic
 - **pass-through-unshaped** for unshaped, unoptimized traffic
 - d** Select the **Direction** of the traffic. Options in the drop-down menu include:
 - **LAN to WAN** [default]
 - **WAN to LAN**
 - e** Select the time **Period**. Options in the drop-down menu include:
 - **Since Reboot** [default]
 - **Current Hour** [begins at the top of the hour, as in 11:00]
 - **Last 60 Minutes**
 - 3** Click **Apply** to display the traffic for the **Top 10 Applications**.
 - 4** To display a table listing the statistics for all applications, click **Table View**.

Viewing Current Flows

The **Monitoring - Current Flows** page retrieves a listing of up to 200 realtime connections based on selectable filter criteria.

- These include end points, ports, traffic path, applications, protocol, Top Talkers, number of flows, and number of (mega)bytes.
- This page provides the ability to filter before waiting on flows to populate. No flows populate until you click **Apply**.
- When you first filter and click **Apply** to generate a list of flows, the page displays a default set of parameter columns, along with individual links to flow details.
- You can customize the Current Flows columns so that you can choose which of the many columns you want to view onscreen.

It answers questions like:

- What flows is the Silver Peak appliance optimizing?
- Which flows are the Top Talkers, that is, which flows are generating the most traffic?
- Which CIFS flows are being accelerated? Or could be?
- Which CIFS flows have SMB signing enabled on the server?
- Which TCP flows are being accelerated? Or could be?
- Is any TCP flow asymmetric?
- What's the performance improvement on a "per flow" basis?
- What SET actions are the Route, QoS, and Optimization policies applying to the flow?
- Is any tunnelized traffic experiencing reduced functionality as a result of two nodes having mismatched software versions?

Selecting Filters

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

Select from the following **Traffic** types

Appliance Time: 2009/03/23 17:13:56 US/Pacific Version: 2.4.1.0_25504
Logged in as: **monitor** (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms | splabnx01 10.0.40.23

Monitoring - Current Flows

Filter Criteria

Two End Points

IP Address1: 0.0.0.0 Port1: 0

IP Address2: 0.0.0.0 Port2: 0

Traffic: All Traffic

Application: All Applications

Filter: None

Max Matches: 100 (1..200)

Apply Clear Filter

Click Apply button to filter current flows.

Two additional **Filter** fields appear if you choose an option other than **None**.

This field's options include built-in **applications**, custom applications, and user-created application groups.

Filter: None

Top Talkers: 10

Filter: None

Greater than: 10

Filter: Protocol

Equal

ip

ip

tcp

udp

ah

egp

eigrp

encap

esp

etherip

fc

gre

icmp

idpr

idpr-cmtp

idpr

Most Bandwidth Consumed (Top Talkers)

When you select a LAN or WAN option, the filter **defaults** to the ten **Top Talker** flows. These are the flows consuming the most bandwidth, in bytes. You can filter for a maximum of 200 flows.

Relative Bandwidth

You can also filter for flows relative to a specific number of bytes (default), with **Equal**, **Greater Than**, or **Less Than**. For example, to filter on:

- 1000 bytes, enter **1000**
- 200 megabytes, enter **200m** or **200M** (no space)
- 45.3 gigabytes, enter **45.3g** or **45.3G** (no space)
- 7 terabytes, enter **7t** or **7T** (no space)

Protocol-Specific

Filter on a protocol to either isolate it (**Equal**) or exclude it (**Not Equal**).



Tip Customize the results columns to show the data that you're filtering. For more detail, see *"To customize the screen display"* on page 276.

Customizing Which Columns Display

Following are some customization guidelines:

- The first time you click **Apply**, a default set of columns displays results. This includes:

IP1	Protocol
Port 1	L/W Ratio (X)
IP2	W/L Ratio (X)
Port 2	SMB Signed
Tx Action	Up Time
Application	Details

The items in **blue** always display, whether or not you customize your screen display.

- Customizations persist for the specific user during a given browser session.** If you log out and in during the same browser session, your customizations persevere. If you quit the browser and subsequently restart it, the screen display automatically resets to the default.
- You can view as many of the possible columns as you want. You may have to scroll, though.
- When you **Export** the data from a given screen, **only the data belonging to the displayed columns is included** in the .csv file.
- Customize** and **Export** functions are accessible to all users.

This screen shows the default set of columns.

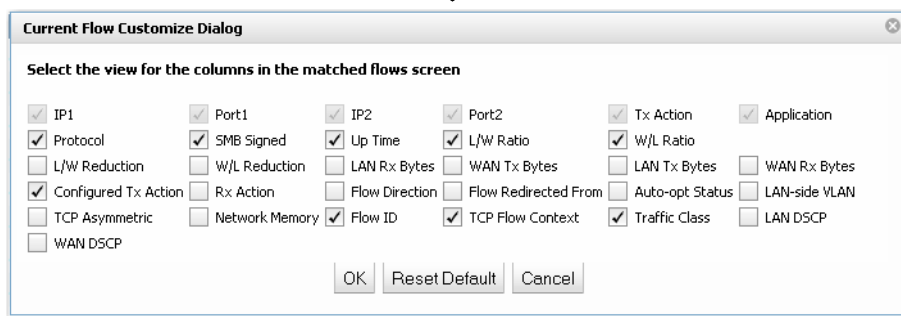
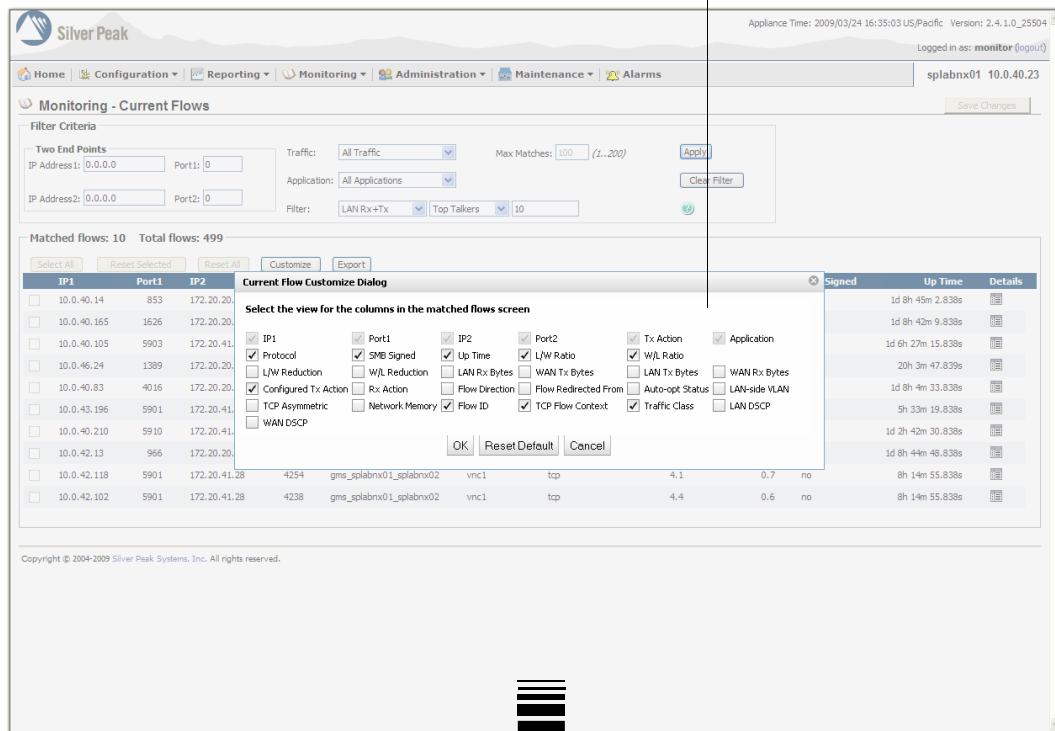
The screenshot shows the Silver Peak Monitoring - Current Flows interface. At the top, there's a navigation bar with tabs: Home, Configuration, Reporting, Monitoring (selected), Administration, Maintenance, and Alarms. Below the navigation bar, there's a filter criteria section with fields for IP Address1, Port1, IP Address2, Port2, Traffic (All Traffic), Application (All Applications), and Filter (None). There are Apply and Reset buttons. Below the filter criteria, it says "Matched flows: 100 Total flows: 498". There are buttons for Select All, Reset Selected, Reset All, Customize, and Export. The main table displays the following columns: IP1, Port1, IP2, Port2, Tx Action, Configured Tx Action, Application, Protocol, L/W Ratio (X), W/L Ratio (X), SMB Signed, Flow ID, TCP Flow Context, Traffic Class, Up Time, and Details. The table contains several rows of data, including flows for vnc1, ssh, cifs_smb, cifs_server, unassigned, dns, and http.

To specify which columns display or not, click **Customize**.

◆ To customize the screen display

- 1 To access the **Current Flow Customize Dialog** box, click **Customize**.

This is the default set. Greyed-out boxes cannot be unchecked.



- 2 To make your selections, check/uncheck boxes as needed.
- 3 Click **OK**. The screen refreshes with the new columns.

◆ To reset to the default display

- 1 To access the **Current Flow Customize Dialog** box, click **Customize**.
- 2 Click **Reset Default**. The screen display resets.

♦ **To customize and review a Current Flows report**

- 1 From the **Monitoring** menu, select **Current Flows**. Initially, the page defaults to all endpoints with **All Traffic** and **All Applications** and no additional filtering.

- When you enter a specific endpoint, the page returns connections 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.

Select the kind of traffic you want.

Choose from all built-in and user-defined applications or application groups, or choose **All Applications**.

The filter defaults to **100** matches. You can specify up to 200.

Click Apply button to filter current flows.

Field	Description
IP Address1 (2) / Port1 (2)	<p>An endpoint(s) on which you want to filter:</p> <ul style="list-style-type: none"> • When you enter a specific endpoint, the page returns connections 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.
Traffic	<p>Select the type of traffic connections you want to retrieve:</p> <ul style="list-style-type: none"> • All Traffic – all optimized and pass-through traffic. • Optimized – the sum of all optimized traffic. That is, all tunnelized traffic. • pass-through – all unoptimized, shaped traffic. • pass-through-unshaped – all unoptimized, unshaped traffic. • [a named Tunnel] – that specific tunnel's optimized traffic.
Application	<p>Select which standard or user-defined application (or application group) to use as a filter criteria. The default value is All Applications.</p>

- a If you keep the default values and click **Apply**, the page retrieves a default of 100 of the total current flows.

Here, the Appliance Manager retrieved 100 out of 498 flows.

Appliance Time: 2009/03/24 16:35:03 US/Pacific Version: 2.4.1.0_25504
Logged in as: admin (logout)

Home Configuration Reporting Monitoring Administration Maintenance Alarms splabnx01 10.0.40.13 Save Changes

Monitoring - Current Flows

Filter Criteria

Two End Points
IP Address1: 0.0.0.0 Port1: 0 Traffic: All Traffic Max Matches: 100 (1..200) Apply
IP Address2: 0.0.0.0 Port2: 0 Application: All Applications Clear Filter
Filter: None

Matched flows: 100 Total flows: 498

IP1	Port1	IP2	Port2	Tx Action	Configured Tx Action	Application	Protocol	L/W Ratio (X)	W/L Ratio (X)	SHB Signed	Flow ID	TCP Flow Context	Traffic Class	Up Time	Details
10.0.40.236	5901	172.20.28.59	2192	gms_splabnx01_spdcmv02	gms_splabnx01_spdcmv02	vnc1	tcp [acce]	2.8	0.9	no	8040	0x2b703da7a2	3	6h 5m 5s	
10.0.40.176	22	172.20.41.128	1456	gms_splabnx01_sphqrv02	gms_splabnx01_sphqrv02	ssh	tcp [acce]	0.9	0.5	no	46581	0x2b70675252	4	2d 6h 4m 44s	
10.0.42.185	4763	172.20.20.64	445	gms_splabnx01_spdcmv02	gms_splabnx01_spdcmv02	cifs_smb	cifs [acce]	1.8	3.1	no	55260	0x2b7052178a	5	15d 2h 53m 0s	
10.0.40.18	139	172.20.20.46	1971	gms_splabnx01_spdcmv02	gms_splabnx01_spdcmv02	cifs_smb	cifs [server]	1.0	0.9	no	51391	0x2b70535422	5	2m 3s	
10.0.40.17	51230	172.20.41.128	4130	gms_splabnx01_sphqrv02	gms_splabnx01_sphqrv02	unassigned	tcp [acce]	0.6	0.7	no	34945	0x2b702fda02	1	1d 1h 4m 24s	
10.0.42.53	47865	172.20.20.57	53	gms_splabnx01_spdcmv02	gms_splabnx01_spdcmv02	dns	udp	0.8	0.8	no	51293	0x0	4	2m 22s	
10.0.41.134	22	172.20.41.38	35189	gms_splabnx01_sphqrv02	gms_splabnx01_sphqrv02	ssh	tcp [acce]	0.6	0.4	no	12429	0x2b704e2706	4	1d 9h 53m 26s	
10.0.41.117	80	172.20.41.28	4368	gms_splabnx01_sphqrv02	gms_splabnx01_sphqrv02	http	tcp [acce]	16.9	1.5	no	51814	0x2b7052f67a	4	25s	

A pink background indicates that the SET action(s) for this flow were changed after the flow began. Therefore, this is now a stale connection.

Appliance Time: 2010/05/07 18:13:14 US/Pacific Version: 3.2.0.0_31090
Logged in as: admin (logout)

Home Configuration Reporting Monitoring Administration Maintenance Alarms haifa 10.0.40.106 Save Changes

Monitoring - Current Flows

Filter Criteria

Two End Points
IP Address1: 0.0.0.0 Port1: 0 Traffic: All Traffic Max Flows: 100 (1..200) Apply
IP Address2: 0.0.0.0 Port2: 0 Application: All Applications Clear Filter
Filter: None

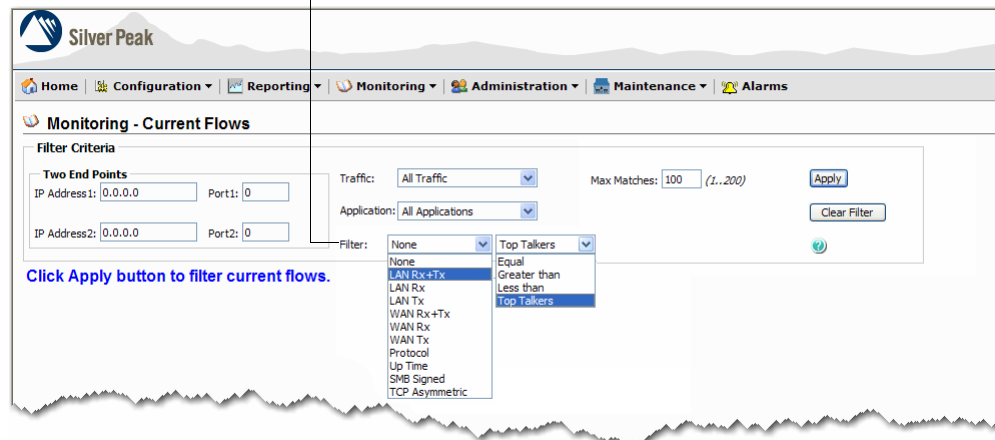
Returned flows: 9 Matched flows: 9 Total flows: 9

IP1	Port1	IP2	Port2	Tx Action	Application	Protocol	L/W Ratio (X)	W/L Ratio (X)	SHB Signed	Up Time	Details
10.1.205.50	80	10.1.206.50	53567	toCairo	abc	tcp	1.0	0.7	no	14h 1m 3.221s	
10.1.205.50	109	10.1.206.50	53575	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.222s	
10.1.205.50	389	10.1.206.50	53572	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.222s	
10.1.205.50	26000	10.1.206.50	53568	toCairo	abc	tcp	1.0	0.7	no	14h 1m 3.221s	
10.1.205.50	666	10.1.206.50	53569	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.221s	
10.1.205.50	443	10.1.206.50	53570	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.221s	
10.1.205.50	1521	10.1.206.50	53574	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.221s	
10.1.205.50	5055	10.1.206.50	53573	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.221s	
10.1.205.50	2049	10.1.206.50	53571	toCairo	abc	tcp	1.0	0.8	no	14h 1m 3.221s	

An orange background indicates that the tunnelized flow is experiencing reduced functionality as a result of two nodes having mismatched software versions. This tunnel compatibility mode offers some basic benefits, but not the full set of optimizations. For the list, see ["Tunnel Compatibility Mode" on page 126](#).

- b** To further refine the results, before or after clicking **Apply**, select options from the **Filter** drop-down list and fields.

As soon as you select an option besides **None** in the **Filter** box, two other option boxes appear to the right.



2 Click **Apply**. The results display.

Flows that use **TCP protocol** have selectable (dark edged) check boxes and can be reset, if you have **admin** privileges. You can select them individually, or click **Select All**.

Resets all flows you've selected in the result display.

Note: This button is greyed out because the user is logged on with the more limited **monitor** privileges or because no flows are selected.

Resets **all TCP flows in the appliance**, whether or not they're displayed in the list of matches.

This is a service-affecting step.

Defaults to 100 matched flows. The maximum is 200.

- To refresh the results list, click **Apply**.
- To revert to default filter values, click **Clear Filter**.

Filter Criteria

Two End Points
 IP Address1: 0.0.0.0 Port1: 0
 IP Address2: 0.0.0.0 Port2: 0

Traffic: All Traffic Max Matches: 100 (1..200)
 Application: All Applications
 Filter: LAN Rx+Tx Top Talkers 10

Matched flows: 10 Total flows: 654

Select All Reset Selected Reset All Customize Export

	IP1	Port1	IP2	Port2	Tx Action	Configured Tx Action	Application	Protocol	L/W Ratio (X)	W/L Ratio (X)	SMB Signed	Flow ID	TCP Flow Context	Traffic Class	Up Time	Details
<input type="checkbox"/>	10.0.42.9	800	172.20.20.54	2049	gms_splabnx01_spdcnx02	gms_splabnx01_spdcnx02	nfs	tcp [accel]	1.2	3.8	no	33575	0x2B700AE1AA	5	17d 0h 34m 56s	
<input type="checkbox"/>	10.0.40.105	2055	10.0.100.115	32768	none	none	filenet	udp	N/A	1.0	no	105	0x0	None	28d 14h 7m 33s	
<input type="checkbox"/>	10.0.40.105	5901	172.20.41.99	1076	gms_splabnx01_sphqnx02	gms_splabnx01_sphqnx02	vnc1	tcp [accel]	5.6	0.5	no	13522	0x2B7020F736	3	23d 2h 36m 29s	
<input type="checkbox"/>	10.0.46.52	3605	172.20.20.54	445	gms_splabnx01_spdcnx02	gms_splabnx01_spdcnx02	cifs_smb	cifs [accel]	1.0	4.4	no	7560	0x2B703CDFAA	5	11d 15h 11m 58s	
<input type="checkbox"/>	10.0.41.198	9996	10.0.100.123	32768	none	none	unassigned	udp	N/A	1.0	no	8646	0x0	None	19d 0h 40m 54s	
<input type="checkbox"/>	10.0.41.183	799	172.20.20.54	2049	gms_splabnx01_spdcnx02	gms_splabnx01_spdcnx02	nfs	tcp [accel]	1.3	3.7	no	830	0x2B706CB352	5	28m 2s	
<input type="checkbox"/>	10.0.40.150	906	172.20.20.64	2049	gms_splabnx01_spdcnx02	gms_splabnx01_spdcnx02	nfs	tcp [accel]	3.0	2.4	no	65061	0x2B7069412E	5	33m 0s	
<input type="checkbox"/>	10.0.40.83	1272	172.20.20.54	445	gms_splabnx01_spdcnx02	gms_splabnx01_spdcnx02	cifs_smb	cifs [accel]	2.5	27.0	no	44142	0x2B701FB336	5	14d 3h 49m 10s	
<input type="checkbox"/>	10.0.41.198	1066	172.20.20.54	445	gms_splabnx01_spdcnx02	gms_splabnx01_spdcnx02	cifs_smb	cifs [accel]	1.1	2.6	no	24081	0x2B70117B46	5	8d 17h 34m 53s	
<input type="checkbox"/>	10.0.41.198	9996	10.0.100.115	32769	none	none	unassigned	udp	N/A	1.0	no	106	0x0	None	28d 14h 7m 33s	

Copyright © 2004-2008 Silver Peak Systems, Inc. All rights reserved.

Clicking the icon in the **Details** column displays a detailed flow report. However, if the flow has already terminated or timed out, no report displays.

Current Flow Details

Following is an example of a **Monitoring - Current Flows** page, with customized columns, followed by the results of clicking **Details** for the first flow listed.

Monitoring - Current Flows

Filter Criteria

Two End Points
 IP Address1: 0.0.0.0 Port1: 0
 IP Address2: 0.0.0.0 Port2: 0

Traffic: All Traffic Max Matches: 100 (1..200)
 Application: All Applications
 Filter: LAN Rx+Tx Top Talkers 10

Matched flows: 10 Total flows: 504

IP1	Port1	IP2	Port2	Tx Action	Configured Tx Action	Application	Protocol	L/W Ratio (X)	L/W Reduction (%)	W/L Ratio (X)	LAN Rx Bytes	LAN Tx Bytes	SMB Signed	Up Time	Details
10.0.41.198	9996	10.0.100.115	32769	none	none	unassigned	udp	N/A	0	1.0	0	3,003,596	no	34d 0h 3m 20s	
10.0.41.133	5901	172.20.41.66	4103	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	vnc1	tcp [accel]	6.0	83.3	0.5	216,379,014	11,307,300	no	7d 4h 47m 21s	
10.0.41.183	799	172.20.20.54	2049	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	nfs	tcp [accel]	1.3	21.4	3.8	76,559,244	179,381,240	no	21m 50s	
10.0.42.249	5901	172.20.41.57	32955	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	vnc1	tcp [accel]	1.2	19.4	0.6	46,612,325	11,249,510	no	12h 16m 27s	
10.0.40.105	5901	172.20.41.99	1097	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	vnc1	tcp [accel]	5.9	83.1	0.5	215,289,239	11,574,800	no	12h 26m 11s	
10.0.40.17	49931	172.20.41.0	3067	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	unassigned	tcp [accel]	9.2	89.1	0.6	69,879,903	2,359,558	no	13h 2m 0s	
10.0.40.150	906	172.20.20.64	2049	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	nfs	tcp [accel]	1.6	35.9	2.5	35,723,200	55,139,564	no	9m 41s	
10.0.46.11	5901	172.20.28.36	2046	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	vnc1	tcp [accel]	1.5	32.8	1.0	7,449,625	3,930,528	no	35m 57s	
10.0.40.17	49927	172.20.41.0	3068	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	unassigned	tcp [accel]	6.0	83.3	0.6	28,466,309	1,811,384	no	13h 2m 0s	
10.0.41.3	2001	172.20.41.71	4214	gms_splabnx01_sphqrx02	gms_splabnx01_sphqrx02	unassigned	tcp [accel]	0.4	0	0.4	5,374,026	2,694,320	no	1d 7h 48m 38s	

Copyright © 2004-2008 Silver Peak Systems, Inc. All rights reserved.

Clicking the icon in the **Details** column displays a detailed flow report. However, if the flow has already terminated or timed out, no report displays, and you get an error message..

splabnx01 - Monitoring - Current Flow Detail - Windows Internet Explorer

Route Information

Map Name: gms_RouteMap
 Priority in Map: 1040
 Endpoint IP1: 10.0.41.133
 Endpoint IP2: 172.20.41.66
 Port1: 5901
 Port2: 4103
 Configured Tx Action: gms_splabnx01_sphqrx02
 Tx Action: gms_splabnx01_sphqrx02
 Rx Action: gms_splabnx01_sphqrx02
 Application: vnc1
 Protocol: tcp [accel]
 Using Stale Map Entry: no
 Flow Direction: WAN to LAN
 Flow Redirected From: None
 Auto-opt Status: Policy Routed
 Auto-opt Transit Node: 10.0.100.110, 10.0.100.72
 LAN-side VLAN: None

Stats Information

LAN to WAN Ratio (X): 6.0
 WAN to LAN Ratio (X): 0.5
 LAN to WAN Reduction: 83.3%
 WAN to LAN Reduction: 0%
 LAN Rx MBytes: 216
 WAN Tx MBytes: 36
 LAN Tx MBytes: 11
 WAN Rx MBytes: 23
 Flow Up Time: 7d 4h 47m 21s
 Flow ID: 3216
 TCP Flow Context: 0x2B70243066

Optimization Information

Map Name: map1
 Priority in Map: default
 TCP Accelerated Configured: yes
 TCP Accelerated Status: yes
 TCP Asymmetric: no
 CIFS Accelerated Configured: yes
 CIFS Accelerated Status: no
 SMB signed: no
 Network Memory: yes
 Payload: yes
 Using Stale Map Entry: no

QoS Information

Map Name: gms_QosMap
 Priority in Map: 40
 Traffic Class: 3
 LAN DSCP: trust-lan
 WAN DSCP: trust-lan
 Using Stale Map Entry: no

Refresh Close

- The (added) green highlighting shows which fields **always** display as columns.
- The (added) yellow highlighting shows which fields you can include in the screen display, by using the **Customize** button.

If a flow is not fully optimized because of a software version mismatch between nodes, then the **Details** page annotates this *reduced functionality*.

haifa - Monitoring - Current Flow Detail - Windows Internet Explorer

Route Information	
Map Name	map1
Priority in Map	default
Endpoint IP1	10.1.205.50
Endpoint IP2	10.1.206.50
Port1	389
Port2	53572
Configured Tx Action	toCairo
Tx Action	toCairo
Rx Action	toCairo
Application/Priority	abc/100
Protocol	tcp
Using Stale Map Entry	no
Flow Direction	LAN to WAN
Flow Redirected From	None
Auto-opt Status	Auto Routed
Auto-opt Transit Node	10.1.206.20
LAN-side VLAN	None

Stats Information	
LAN to WAN Ratio (X)	1.0
WAN to LAN Ratio (X)	0.8
LAN to WAN Reduction	0%
WAN to LAN Reduction	0%
LAN Rx Bytes	6,725,911,620
WAN Tx Bytes	6,913,637,613
LAN Tx Bytes	154,494,428
WAN Rx Bytes	194,931,623
Flow Up Time	14h 1m 3.222s
Flow ID	1093
TCP Flow Context	0x2AAB38CECF36A

Optimization Information	
Map Name	map1
Priority in Map	default
TCP Accelerated Configured	no
TCP Accelerated Status	no [Reduced Functionality]
TCP Asymmetric	no
CIFS Accelerated Configured	no
CIFS Accelerated Status	no
SMB Signed	no
Network Memory	no [Reduced Functionality]
Payload	no
Using Stale Map Entry	no

QoS Information	
Map Name	map1
Priority in Map	default
Traffic Class	1
LAN DSCP	trust-lan
WAN DSCP	trust-lan
Using Stale Map Entry	no

Refresh Close

In the following definitions table, the red rows correspond to information that is **only** available in the **Current Flows Details** page. The rest are available by default.

Field	Definition
Route Information	
Map Name	The name of the Route Policy.
Priority in Map	The number of the entry in the Route Policy that the flow matches.
Endpoint IP1 (or IP2)	IP Address of one end of the flow. This will either be the specific address you filtered for or the other end of a connection to it.
Port1 (or Port2)	Number of the port the application is using. A value of 0 means "any port".
Configured Tx Action	The SET action configured in the Route Policy's Tunnel field.
Tx Action	How the traffic is actually being transmitted.
Rx Action	By what path or method the appliance is receiving this flow's traffic.
Application	Name of the application to which that flow's traffic belongs.
Protocol	<p>The Protocol options include the following:</p> <ul style="list-style-type: none"> • tcp — unaccelerated TCP • udp — UDP • tcp [accel] — accelerated TCP • cifs [accel] — this CIFS flow is on the client side • cifs [server] — this CIFS flow is on the server side, so disregard. CIFS only displays as "accelerated" on the client-side appliance, not on the server side. • cifs — when we don't know whether a CIFS connection is client-side or server-side (for example, in a hairpinned connection). • other ip protocols — like icmp, etc.
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 LAN to WAN , or WAN to LAN .
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 Optimized.
Auto-opt Transit Mode	The IP addresses of the hops between this appliance and the other end of the connection.
LAN-side VLAN	Whether or not the flow is part of a VLAN configured in Appliance Manager..
Optimization Information	
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 Accelerated Configured	Whether or not TCP acceleration is configured in the Optimization Policy.
TCP Accelerated Status	Whether the flow is TCP accelerated [Yes] or not [No].
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.
CIFS Accelerated Configured	Whether or not CIFS acceleration is configured in the Optimization Policy. [Yes/No]
CIFS Accelerated Status	Whether or not CIFS is being accelerated or not. [Yes/No]

Field	Definition (Continued)
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.
Network Memory	Whether or not Network Memory is turned on.
Payload	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.
Stats Information	
LAN to WAN Ratio (X)	For the outbound traffic, a ratio of the LAN Rx Bytes divided by the WAN Tx 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.
WAN to LAN Ratio (X)	For the inbound traffic, a ratio of the WAN Rx Bytes divided by the LAN Tx Bytes .
LAN to WAN Reduction	The percentage by which outbound traffic is reduced in size.
WAN to LAN Reduction	The percentage of optimization in the WAN-to-LAN direction.
LAN Rx Bytes	Total number of bytes received from the LAN [outbound traffic]
WAN Tx Bytes	Total number of bytes sent to the WAN [outbound traffic]
LAN Tx Bytes	Total number of bytes sent to the LAN [inbound traffic]
WAN Rx Bytes	Total number of bytes 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.
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: <ul style="list-style-type: none"> • If the traffic is tunnelized, the traffic class is tunnel-specific. • If the traffic is pass-through (shaped), the traffic class is based on the pass-through (shaped) configuration.
LAN DSCP	The LAN DSCP marking that the QoS policy assigned to the flow, based on the MATCH conditions satisfied.
WAN DSCP	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.

Resetting Flows to Improve Performance

In the list of (up to 200) results, you can look for the flows that aren't being accelerated, but *could* be. Generally, this means flows that use TCP protocol and are not TCP-accelerated:

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

The arrows show traffic that **could be** TCP-accelerated.

Accessible (not dimmed) check boxes indicate that resetting the flows would terminate the flows and then reestablish them **with** TCP acceleration, *if* the flow would now fit the Route Policy's MATCH criteria for a tunnel.

Appliance Name: 2009/03/24 16:35:03 US/Pacific

Version: 2.4.1.0_25504

Home

Configuration

Reporting

Monitoring

Administration

Maintenance

Alarms

Monitoring - Current Flows

Filter Criteria

Two End Points

IP Address1: 0.0.0.0

Port1: 0

IP Address2: 0.0.0.0

Port2: 0

Traffic: All Traffic

Application: All Applications

Filter: None

Max Matches: 100 (1..200)

Apply

Clear Filter

Matched flows: 100

Total flows: 430

Select All

Reset Selected

Reset All

Customize

Export

IP1	Port1	IP2	Port2	Tx Action	Application	Protocol	L/W Ratio (X)	W/L Ratio (X)	SHB Signed	Up Time	Details	
<input type="checkbox"/>	10.0.40.208	22	10.0.60.68	33977	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.5	0.4	no	4d 7h 23m 18s	
<input type="checkbox"/>	10.0.41.7	32770	172.20.20.10	53	gms_splabmx01_gphqnx02	dns	udp	0.8	1.0	no	7s	
<input type="checkbox"/>	10.0.42.102	123	172.20.20.100	123	gms_splabmx01_gphqnx02	ntp	udp	0.7	N/A	no	1m 3s	
<input type="checkbox"/>	10.0.43.74	32773	172.20.20.57	53	gms_splabmx01_gphqnx02	dns	udp	0.7	0.9	no	2m 15s	
<input type="checkbox"/>	10.0.60.64	47806	172.19.10.218	443	gms_splabmx01_gphqnx02	https	tcp	0.8	1.0	no	2m 2s	
<input type="checkbox"/>	172.19.10.223	3012	172.20.41.68	1494	gms_splabmx01_gphqnx02	citrix	tcp	17.8	1.0	no	2m 4s	
<input type="checkbox"/>	10.0.42.46	123	172.20.20.10	123	gms_splabmx01_gphqnx02	ntp	udp	0.7	0.7	no	1m 52s	
<input type="checkbox"/>	10.0.42.249	22	172.20.41.57	32902	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.3	0.3	no	4d 8h 57m 15s	
<input type="checkbox"/>	10.0.42.13	22	172.20.20.64	57842	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.5	0.5	no	6h 9m 43s	
<input type="checkbox"/>	10.0.40.10	443	172.20.28.41	2384	gms_splabmx01_gphqnx02	https	tcp [accel]	1.0	0.8	no	3m 28s	
<input type="checkbox"/>	10.0.40.102	22	10.0.61.12	47662	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.3	0.3	no	21d 8h 10m 54s	
<input type="checkbox"/>	10.0.40.17	51230	172.20.41.0	2947	gms_splabmx01_gphqnx02	unassigned	tcp [accel]	0.6	0.7	no	4d 9h 13m 27s	
<input type="checkbox"/>	10.0.40.21	22	172.20.41.57	32861	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.3	0.3	no	5d 11h 8m 43s	
<input type="checkbox"/>	10.0.41.134	22	172.20.41.38	35189	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.3	0.5	no	7d 11h 51m 10s	
<input type="checkbox"/>	10.0.40.21	22	172.20.41.57	32862	gms_splabmx01_gphqnx02	ssh	tcp [accel]	0.3	0.3	no	5d 10h 51m 35s	
<input type="checkbox"/>	10.0.40.10	443	172.20.28.41	2384	gms_splabmx01_gphqnx02	https	tcp	0.5	1.0	no	3m 28s	

Viewing Tunnel QoS Statistics

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

It answers the questions:

- How many dropped and/or out-of-order packets are seen on the appliance on a Traffic Class or DSCP basis?
- Is QoS working correctly?
- Are there problems in the network related to insufficient re-order wait times?

◆ To view a Tunnel QoS Statistics report

- 1 From the **Monitoring** menu, select **QoS**. Initially, the chart defaults to all **Traffic Classes** and all **DSCP** markings for the first tunnel listed alphabetically. In other words, the results are not filtered.

Choose an individual tunnel.

You can choose **all** traffic classes, or any one from **1** through **10**. The default is **all**.

You can choose **all** DSCP markings, or any one individually. The default is **all**.

Appliance Time: 2009/08/14 16:10:00 US/Pacific Version: 3.0.0.0_27254
Logged in as: admin (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms | splabnx04 10.0.40.28

Monitoring - Tunnel QoS

Save Changes

Selection Criteria

Refresh: manually ☒ Actual Stats ☐ Delta Stats

Tunnel: labnx03 Traffic Class: All DSCP: All Apply

LAN Rx Bytes:	44,870,344,428
LAN Rx Pkts:	52,705,882
LAN Tx Bytes:	19,323,439,772
LAN Tx Pkts:	45,873,463
LAN Rx Dropped Pkts:	15,588
QoS Class Total Bytes Exceeded:	0
QoS Class Total Pkts Exceeded:	0
QoS Class Per-Flow Bytes Exceeded:	0
QoS Class Per-Flow Pkts Exceeded:	0
QoS Class Queue Time Exceeded:	15,588

WAN Tx Bytes:	40,187,666,217
WAN Tx Pkts:	46,298,108
WAN Rx Bytes:	8,769,203,663
WAN Rx Pkts:	41,707,422

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

- 2 Select the **Refresh** interval.
- 3 Select **Actual Stats** or **Delta Stats**.
- 4 Select options from the **Tunnel**, **Traffic Class**, and **DSCP** drop-down menus.
- 5 Click **Apply**. The tunnel QoS statistics update.

When you select a tunnel from the table, the Appliance Manager provides two detailed reports:

- **QoS Statistics on Traffic Class**
- **QoS Statistics on DSCP**

The **QoS Statistics on Traffic Class** area displays the following information:

Field	Definition
LAN Rx Bytes	Number of bytes received from the LAN
LAN Rx Pkts	Number of packets received from the LAN
LAN Tx Bytes	Number of bytes sent to the LAN
LAN Tx Pkts	Number of packets sent to the LAN
Lan Rx Dropped Pkts	Number of packets received from the LAN that were dropped
QoS Class Total Bytes Exceeded	Number of packets that exceeded the byte queue limit for a traffic class
QoS Class Total Pkts Exceeded	Number of packets that exceeded the packet queue limit for a traffic class
QoS Class Per-Flow Bytes Exceeded	Number of packets that exceeded the bytes-per-flow queue limit for a traffic class
QoS Class Per-Flow Pkts Exceeded	Number of packets that exceeded the packets-per-flow queue limit for a traffic class
QoS Class Queue Time Exceeded	Number of packets that exceeded the maximum queue wait time for a traffic class

The **QoS Statistics on DSCP** area displays the following information:

Field	Definition
WAN Tx Bytes	Number of bytes sent to the WAN
WAN Tx Pkts	Number of packets sent to the WAN
WAN Rx Bytes	Number of bytes received from the WAN
WAN Rx Pkts	Number of packets received from the WAN

Viewing Tunnel Realtime Statistics

The **Monitoring - Tunnels** page summarizes the overall inbound and outbound traffic statistics for the tunnels since the last reboot. The **Tunnels** table condenses the total LAN and WAN counters.

When you select a tunnel from the table, the Appliance Manager provides the following detailed report:

Name	Description
LAN/WAN Statistics	<ul style="list-style-type: none"> Specifies the number of bytes and packets received, processed, and transmitted by a Silver Peak tunnel in both the outbound (LAN-to-WAN) and inbound (WAN-to-LAN) directions. Statistics are separated for inbound and outbound traffic.
Flows/Latency/ Packet Correction Statistics	<ul style="list-style-type: none"> Specifies packets by TCP flow versus non-TCP flow. Packets in a TCP flow are further sorted by whether or not they're accelerated. Displays round trip latency time in milliseconds (minimum, maximum, and average). Displays how many received packets were lost before and after Forward Error Correction (FEC). Displays how many received packets were out-of-order before and after Packet Order Correction (POC). Statistics represent combined, bi-directional data.

The default display unit is **Bytes**. If you want, you can choose **MBytes**, **Pkts** [packets], or **KPkts** instead. If you select a specific tunnel and then change the units, the page refreshes to display the table only.

When you choose **manually**, click the browser's **Refresh** icon to view up-to-the-minute data.

How long the tunnel has been up

For detail, click a tunnel name.

OUTBOUND traffic
(Transmit LAN to WAN)

INBOUND traffic
(Receive WAN to LAN)

Appliance Time: 2008/05/14 15:50:59 US/Pacific
Logged in as: admin (logout)
Santiago 10.0.40.111

Monitoring - Tunnels

Selection Criteria
Refresh: manually Actual Stats Delta Stats
Table Unit: Bytes Apply

Tunnels	LAN Rx	WAN Tx	L/W Reduction	LAN Tx	WAN Rx	W/L Reduction	Up Time
11	248,375,650,204	73,142,321,509	70.6%	4,695,466,910	3,960,892,893	15.6%	2d 6h 26m 13s

Copyright © 2004-2008 Silver Peak Systems, Inc. All rights reserved.

LAN Rx - WAN Tx
LAN Rx

WAN Rx - LAN Tx
WAN Rx

The Appliance Manager reports all realtime tunnel statistics as raw data.

◆ **To view a specific tunnel's detailed statistics**

Click the tunnel's **Name**. Typically, the following displays:

The screenshot shows the Silver Peak Appliance Manager interface. The top navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as admin. The main section is titled "Monitoring - Tunnels" and includes a "Selection Criteria" panel with a "Refresh" dropdown set to "manually" and "Table Unit" set to "Bytes". Below this is a table of tunnels with columns: Name, LAN Rx, WAN Tx, L/W Reduction, LAN Tx, WAN Rx, W/L Reduction, and Up Time. The first tunnel, "t1", is highlighted. To the left of the tunnel details, there are two blue arrows pointing to the "LAN/WAN Statistics" and "Flows/Latency/Packet Correction Statistics" sections.

Tunnel Statistics : t1 [Close]

LAN/WAN Statistics

LAN Rx Bytes:	248,375,734,720	WAN Tx Bytes:	73,142,410,729
LAN Rx Pkts:	209,667,517	WAN Tx Pkts:	69,911,328
LAN Tx Bytes:	4,695,473,862	WAN Rx Bytes:	3,960,903,371
LAN Tx Pkts:	116,923,117	WAN Rx Pkts:	35,383,653

Flows/Latency/Packet Correction Statistics

Traffic Flows

Non-TCP Flows:	1	TCP Flows:	1
TCP Accel Flows:	0	TCP Host Accel Flows:	1

Round Trip Latency

Average:	15	Minimum:	10
Maximum:	20		

Rx Packet Correction

Pre FEC Loss:	176,856	Post FEC Loss:	0
Pre POC Out-of-Order:	10,759,709	Post POC Out-of-Order:	1,231

Copyright © 2004-2008 Silver Peak Systems, Inc. All rights reserved.

In **tunnel compatibility mode**, a tunnel performs a reduced set of optimizations. The **Monitoring - Tunnels** page notes that in the **Status** column and highlights the tunnel with yellow.

The screenshot shows the Silver Peak Appliance Manager interface with the "Monitoring - Tunnels" page. The "Selection Criteria" panel is the same as in the previous screenshot. The table of tunnels includes a new "Status" column. The tunnel "toCaro" is highlighted in yellow, indicating it is in compatibility mode. Its status is "up - active - reduced functionality".

Name	Status	LAN Rx	WAN Tx	L/W Reduction	LAN Tx	WAN Rx	W/L Reduction	Up Time
toAnakin	down	0	0	0.0%	0	0	0.0%	0s
toArwen	down	0	0	0.0%	0	0	0.0%	0s
toCaro	up - active - reduced functionality	1,669,134,586,904	1,715,698,808,504	0.0%	35,332,154,460	45,192,137,992	0.0%	25d 6h 44m 35s
toNewYork	down	0	0	0.0%	0	0	0.0%	0s
toTaipei	down	0	0	0.0%	0	0	0.0%	0s

For the list of preserved functionalities and disabled optimizations, see ["Tunnel Compatibility Mode" on page 126](#).

The Appliance Manager organizes an individual tunnel's statistics into two parts:

- **LAN/WAN Statistics** See page 290.
- **Flows / Latency / Packet Correction Statistics** See page 291.

LAN/WAN Statistics

The **LAN/WAN Statistics** summarize realtime data directly related to the Silver Peak tunnel's processing. These statistics answer the following questions:

- For any given tunnel, how many bytes (or packets) did the tunnel receive and subsequently transmit?
- Which tunnels have processed the most traffic? The least traffic?
- What error types and quantities were encountered for traffic inbound from the WAN?
- What error types and quantities were encountered for traffic inbound from the LAN?

Tunnel Statistics : t1 Close			
LAN/WAN Statistics			
LAN Rx Bytes:	248,375,734,720	WAN Tx Bytes:	73,142,410,729
LAN Rx Pkts:	209,667,517	WAN Tx Pkts:	69,911,328
LAN Tx Bytes:	4,695,473,862	WAN Rx Bytes:	3,960,903,371
LAN Tx Pkts:	116,923,117	WAN Rx Pkts:	35,383,653

The **LAN/WAN Statistics** area displays the following information:

Section	Field	Definition
LAN Rx (outbound traffic)	Rx Bytes	Number of bytes received from the LAN.
	Rx Pkts	Number of packets received from the LAN.
WAN Tx (outbound traffic)	Tx Bytes	Number of bytes sent to the WAN.
	Tx Pkts	Number of packets sent to the WAN.
LAN Tx (inbound traffic)	Tx Bytes	Number of bytes sent to the LAN.
	Tx Pkts	Number of packets sent to the LAN.
WAN Rx (inbound traffic)	Rx Bytes	Number of bytes received from the WAN.
	Rx Pkts	Number of packets received from the WAN.

Flows / Latency / Packet Correction Statistics

The **Flows / Latency / Packet Correction Statistics** area includes other general statistics. It answers the questions:

- How many of the traffic flows are based on TCP and how many are not?
- How much of the TCP flow was accelerated?
- What is the minimum, average, and peak latency in milliseconds?
- How many packets were lost before Forward Error Correction (FEC), and how many were lost after?
- How many out-of-order packets were there before and after Packet Order Correction?

Flows/Latency/Packet Correction Statistics			
Traffic Flows			
Non-TCP Flows:	1	TCP Flows:	1
TCP Accel Flows:	0	TCP Non-Accel Flows:	1
Round Trip Latency			
Average:	15		
Maximum:	20	Minimum:	10
Rx Packet Correction			
Pre FEC Loss:	176,856	Post FEC Loss:	0
Pre POC Out-of-Order:	10,759,709	Post POC Out-of-Order:	1,231

TCP Flow / Latency / Packet Loss Statistics area displays the following statistics:

Section	Field	Definition
Traffic Flows	Non-TCP Flows	Number of flows that are not TCP-based.
	TCP Flows	Number of flows that are TCP-based.
	TCP Accel Flows	Number of TCP flows that are accelerated. Since CIFS acceleration is a subset of TCP acceleration, they are included herein.
	TCP Non-Accel Flow	Number of TCP flows that are not accelerated.
Round Trip Latency	Average	Length of the average round trip latency, in milliseconds.
	Maximum	Length of the peak round trip latency, in milliseconds.
	Minimum	Length of the shortest round trip latency, in milliseconds.
Rx Packet Correction	Pre FEC Loss	Number of packets lost before Forward Error Correction (FEC).
	Post FEC Loss	Number of packets lost after Forward Error Correction (FEC).
	Pre POC Out-of-Order	Number of out-of-order packets before Packet Order Correction (POC).
	Post POC Out-of-Order	Number of out-of-order packets after Packet Order Correction (POC).

Enabling **Forward Error Correction (FEC)** can sometimes result in the creation of additional **Out-of-Order Packets (OOP)**. After enabling **FEC**, look for **OOP** and adjust if necessary, as follows:

- 1 Go to **Monitoring > Network Integrity > Packet Order Correction**.
- 2 If out-of-order packets exist, then you'll need to try another **Reorder Wait** time for the tunnel in question.
- 3 Go to **Configuration > Tunnels**, and select the tunnel in question.
- 4 At first, set the **Reorder Wait** time to **10ms**, and save the configuration.
- 5 Return to **Monitoring > Network Integrity > Packet Order Correction** to see if the out-of-order packets have been eliminated.
- 6 If there are still out-of-order packets, then go back to the tunnel configuration and increase the **Reorder Wait** time.
- 7 Repeat Steps 5 and 6 until there are no more out-of-order packets.



Tip To view these same tunnel statistics for preselected time intervals, go to the **Reporting** menu and select from the following: **Flow Counts**, **Latency**, or **Network Integrity**.

For more information, see “[Viewing Flow Counts](#)” on page 255, “[Viewing Latency](#)” on page 257, or “[Viewing Network Integrity](#)” on page 258.

Viewing Reduction Statistics

The **Monitoring - Reduction** page summarizes the reduction of (data) bytes afforded by the Silver Peak appliance over three selectable time periods: realtime (every 6 seconds), and minute by minute over the current hour or the last 60 minutes.

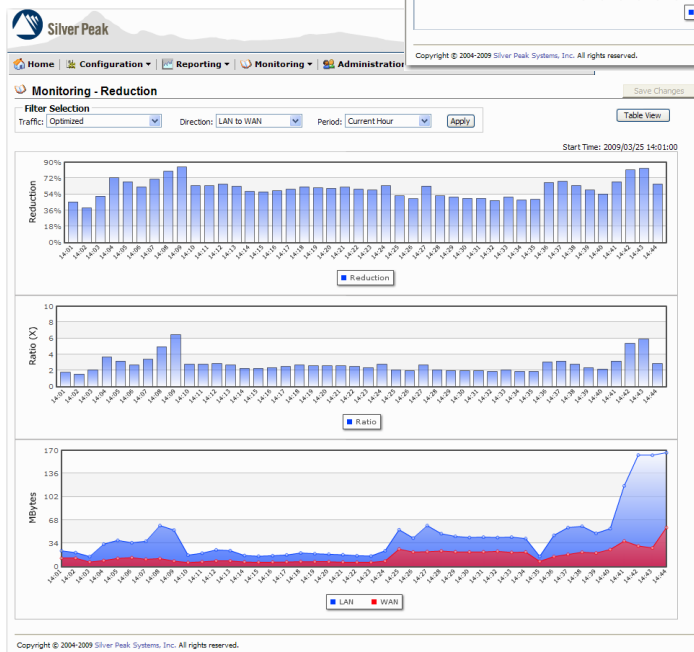

The **Reduction** report answers the following questions:

- How much data is traveling in real-time?
- How much data was sent and received for each minute in the current hour or last 60 minutes?
- What is the ratio of LAN to WAN (or WAN to LAN) traffic at any point in time?

A Sampling of Results

This **Realtime line chart** shows the raw LAN and WAN data, in bytes, for tunnel **gms_splabnx01_splabnx02**, every ~6 seconds. Data accrues at the right side of the chart.

The green line shows the ratio of the two. In the **LAN to WAN** direction, the ratio is calculated from: $\frac{[\text{LAN value}]}{[\text{WAN value}]}$.



Current Hour and Last 60 Minutes charts display percent reduction, ratio (x), and megabytes (LAN and WAN). Both charts provide data table access.

Date and time of first data point.

Start Time: 2009/03/25 14:01:00

n (%)	LAN Bytes	WAN Bytes	LAN Packets	WAN Packets
45.8	23,298,025	12,630,836	76,650	65,673
39.0	20,574,282	12,540,913	73,377	64,856
52.1	14,911,326	7,180,354	67,875	60,158
72.6	33,023,193	9,043,105	77,941	66,088
68.4	38,576,698	12,199,144	87,749	74,096
62.4	35,180,254	13,223,531	80,511	71,041
70.9	37,327,493	10,873,233	80,909	64,523
80.0	60,101,861	12,048,318	97,281	70,069
84.6	53,431,580	8,239,269	91,029	60,309
63.9	16,920,408	6,180,971	61,923	54,165
63.7	19,832,414	7,203,211	69,981	61,103
2009/03/25 14:12:00	2.9	65.5	24,662,996	8,519,310
2009/03/25 14:13:00	2.7	63.1	23,515,571	8,672,820
2009/03/25 14:14:00	2.3	57.3	16,510,017	7,052,709
2009/03/25 14:15:00	2.3	57.1	15,351,849	6,590,557
2009/03/25 14:16:00	2.4	58.0	16,130,538	6,769,008
2009/03/25 14:17:00	2.5	59.8	17,252,444	6,937,388
2009/03/25 14:18:00	2.7	62.7	20,241,644	7,542,132
2009/03/25 14:19:00	2.6	61.4	19,180,565	7,405,174
2009/03/25 14:20:00	2.6	60.9	18,360,130	7,177,044
2009/03/25 14:21:00	2.6	62.0	17,662,237	6,710,976
2009/03/25 14:22:00	2.5	59.8	16,164,641	6,481,135
2009/03/25 14:23:00	2.5	59.8	15,164,641	6,481,135
2009/03/25 14:24:00	2.5	59.8	14,164,641	6,481,135
2009/03/25 14:25:00	2.5	59.8	13,164,641	6,481,135

This **table** displays the raw data presented in the chart. If you wanted to view the data in terms of packets, this is where you would need to look.

♦ **To view a Reduction report**

- 1 From the **Monitoring** menu, select **Reduction**. Initially, the chart defaults to real-time LAN-to-WAN traffic for the first tunnel, as listed alphabetically.
- 2 In the **Filter Selection** section:
 - a Select the **Traffic** type. Options in the drop-down menu include:
 - **Optimized** – the sum of all optimized traffic. That is, all tunneled traffic.
 - individual tunnel name(s), listed alphabetically
 - b Select the **Direction** of the traffic. Options in the drop-down menu include:
 - **LAN to WAN** [default]
 - **WAN to LAN**
 - **Bi-directional**
 - c Select the time **Period**. Options in the drop-down menu include:
 - **Realtime** [default]
 - **Current Hour**
 - **Last 60 Minutes**
- 3 Click **Apply** to update and display the chart.
- 4 To view minute-by-minute data displayed as a table, click **Table View**.

Viewing Bandwidth Statistics

The **Monitoring - Bandwidth** page summarizes the inbound and outbound bandwidth improvements afforded by the Silver Peak appliance over three selectable time periods: realtime (every 6 seconds), and minute by minute over the current hour or the last 60 minutes.

The **Bandwidth** report answers the following questions:


- How much has the bandwidth been optimized?
- At what rate was the data sent and received in each time interval?
- What is the ratio of LAN to WAN traffic at any point in time?

A Sampling of Results

This **Realtime** view shows the actual data rates — and the ratio of rates — for a single tunnel over the last hour. The direction of traffic is LAN to WAN. The chart adds a data point roughly every six seconds.

You can also view a chart for the **Current Hour**.


By default, the charts display in megabits per second. To view the data in terms of packets per second, view the table.



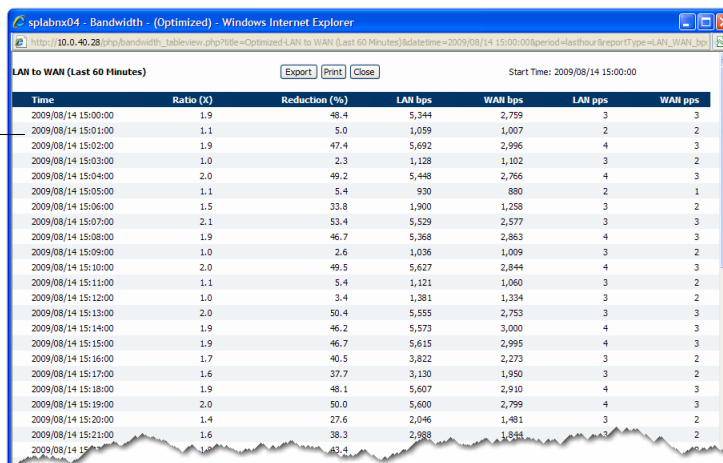
This **Last 60 Minutes** view shows the actual data rates — and the ratio of rates — for all **optimized** traffic over the last hour. The direction of traffic is LAN to WAN.

You can also view a chart for the **Current Hour**.

By default, the charts display in megabits per second. To view the data in terms of packets per second, view the table.



By default, the charts display in megabits per second. To view the data in terms of packets per second, view the table.



♦ To view a Bandwidth optimization report

- 1 From the **Monitoring** menu, select **Bandwidth**. Initially, the chart defaults to the last hour's optimized LAN-to-WAN traffic, expressed in megabits per second.
- 2 In the **Filter Selection** section:
 - a Select the **Traffic** type. Options in the drop-down menu include:
 - **Optimized** – the sum of all optimized traffic. That is, all tunnelized traffic.
 - individual tunnel name(s), listed alphabetically
 - **pass-through** for shaped, unoptimized traffic
 - **pass-through-unshaped** for unshaped, unoptimized traffic
 - b Select the **Direction** of the traffic. Options in the drop-down menu include:
 - **LAN to WAN** [default]
 - **WAN to LAN**
 - **Bi-directional**
 - c Select the time **Period**. Options in the drop-down menu include:
 - **Realtime** [default]
 - **Current Hour**
 - **Last 60 Minutes**
- 3 Click **Apply** to update and display the chart.
- 4 To view minute-by-minute data displayed as a table, click **Table View**.

Viewing Flow Counts

The **Monitoring - Flow Counts** page summarizes the number of TCP flows versus non-TCP flows. Within the TCP flows, its bar charts separate out how many flows were accelerated and how many were not. Since CIFS acceleration is a subset of TCP acceleration, that data is incorporated generically in the accelerated TCP flow data.

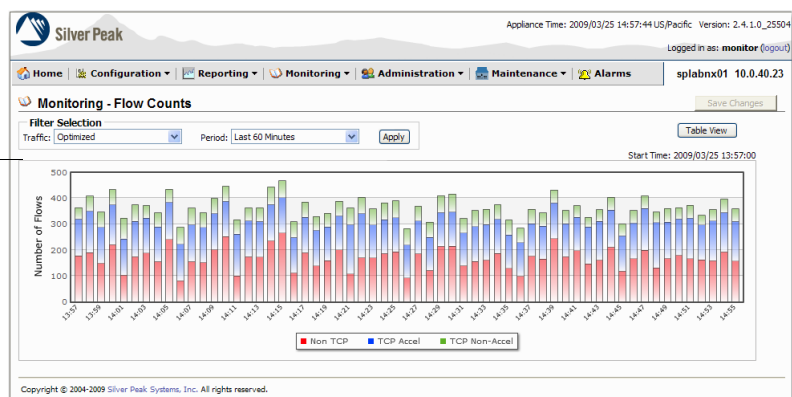
This report answers the following questions:

- How much of my traffic is TCP-based?
- How much of my TCP traffic is accelerated?

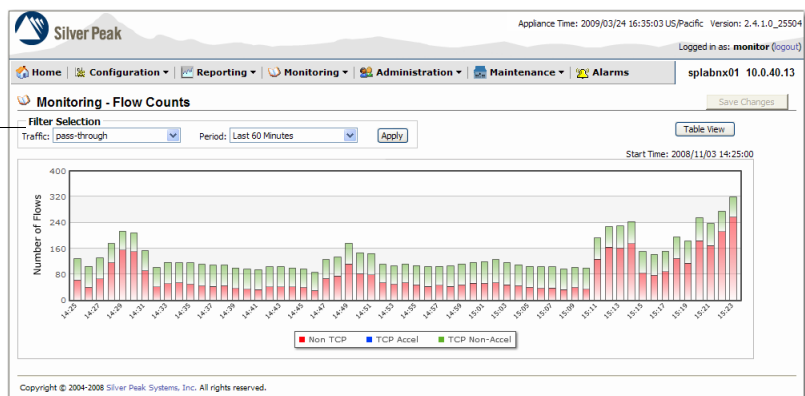
A Sampling of Results

This **bar chart** displays the values for optimized traffic, minute by minute, over the last 60 minutes. You can also choose to view data for a specific tunnel.

By definition, flows are considered bi-directional.



This **bar chart** displays the values for shaped passthrough traffic, minute by minute, over the last 60 minutes.



To export this table as a .csv (comma-separated values) file that you can open as an Excel spreadsheet, click **Export**. The **File Download** dialog box displays.

When the statistics were sampled

Here, **Table View** displays all 4 traffic flows.

splabnx01 - Monitoring - Flows (Optimized) - Windows Internet Explorer

Last 60 Minutes

Export Print Close

Start Time: 2009/03/25 13:57:00

Time	Non-TCP Flows	TCP Flows	TCP Accel Flows	TCP Non-Accel Flow
2009/03/25 13:57:00	179	185	141	44
2009/03/25 13:58:00	191	217	159	58
2009/03/25 13:59:00	149	198	139	59
2009/03/25 14:00:00	221	212	154	58
2009/03/25 14:01:00	104	220	140	80
2009/03/25 14:02:00	175	200	137	63
2009/03/25 14:03:00	189	183	135	48
2009/03/25 14:04:00	156	188	134	54
2009/03/25 14:05:00	241	193	143	50
2009/03/25 14:06:00	82	206	141	65
2009/03/25 14:07:00	156	206	142	64
2009/03/25 14:08:00	152	192	135	57
				58

♦ **To view a Flow report**

- 1 From the **Monitoring** menu, select **Flow Counts**. The chart defaults to the actual data for **optimized** traffic over the last hour.
- 2 In the **Filter Selection** section:
 - a Select the **Traffic** type. Options in the drop-down menu include:
 - **Optimized** – the sum of all optimized traffic. That is, all tunnelized traffic.
 - individual tunnel name(s), listed alphabetically
 - **pass-through** for shaped, unoptimized traffic
 - **pass-through-unshaped** for unshaped, unoptimized traffic
 - b Select the time **Period**. Options in the drop-down menu include:
 - **Realtime** [default]
 - **Current Hour**
 - **Last 60 Minutes**
- 3 Click **Apply** to update and display the chart.
- 4 To view the table, click **Table View**. The table charts four items:


Field	Definition
Non TCP Flows	Number of flows that are not TCP-based.
TCP Flows	Number of flows that are TCP-based.
TCP Accel Flows	Number of TCP flows that are accelerated. Since CIFS acceleration is a subset of TCP acceleration, they are included.
TCP Non-Accel Flows	Number of flows that are not TCP-accelerated.

Viewing Latency Statistics

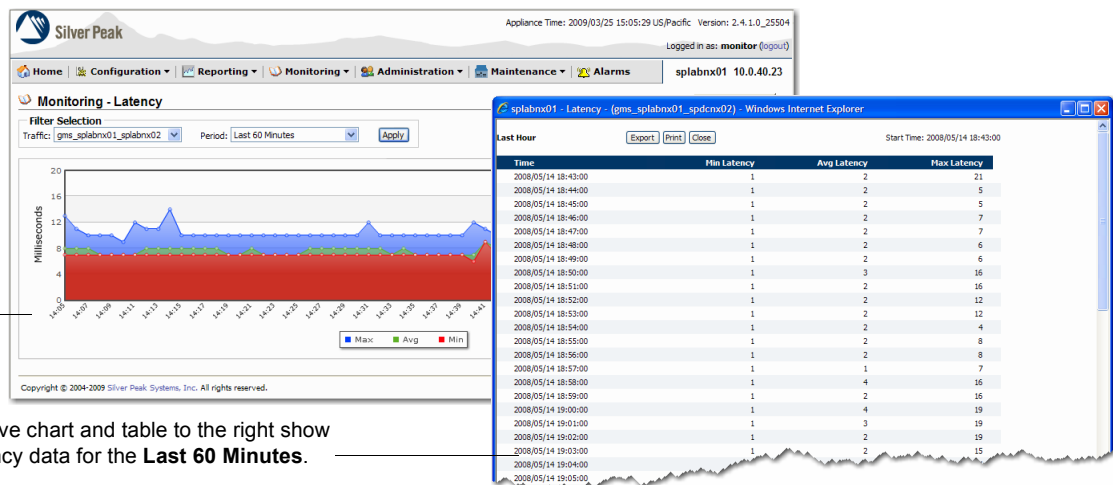
The **Monitoring - Latency** page summarizes the round-trip time of data in a Silver Peak tunnel during three selectable time periods: realtime (every 6 seconds), and minute by minute over the current hour or the last 60 minutes. It answers the following questions:

- How long does it take my data to get to the other end of the Silver Peak tunnel?
- What were the peak, average, and minimum time intervals?

A Sampling of Results



Roughly every 6 seconds, this **line chart** adds data points for the minimum, average, and maximum round trip **latency** for tunnel **gms_splabnx01_splabnx02**.



The above chart and table to the right show the latency data for the **Last 60 Minutes**.

◆ To view the Latency report

- 1 From the **Monitoring** menu, select **Latency**. By default, the chart displays the last hour's statistics for the first tunnel listed alphabetically.

By default, the **Report Type** is **Min Max Avg**. It displays the minimum, peak, and average latency for each time interval.

- 2 From the **Traffic** menu, select the name of the individual tunnel.
- 3 From the **Period** menu, select the time interval.
- 4 Click **Apply** to update and display the chart.

To view the data displayed as a table, click **Table View**.

Viewing Network Integrity Statistics

The **Monitoring - Network Integrity** page summarizes the following statistics by tunnel:

- the number of packets lost before and after enabling **Forward Error Correction (FEC)**
- the number of out-of-order packets before and after enabling **Packet Order Correction (POC)**.

The bar charts show data for three selectable time periods: realtime (every 6 seconds), and minute by minute over the current hour or the last 60 minutes.

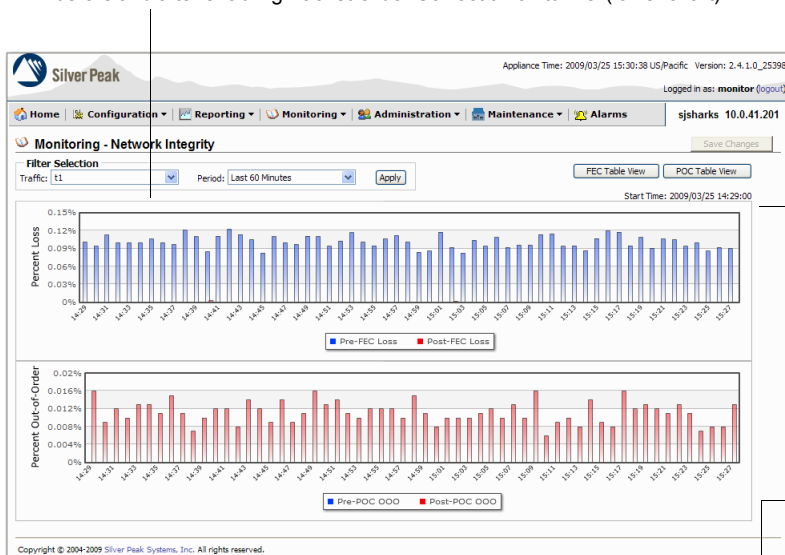
It answers the following questions:

- How many errors were there before and after turning on Forward Error Correction?
- How many out-of-order packets were there before/after turning on Packet Order Correction?
- For any given minute, what was the percent loss?

A Sampling of Results

For tunnel **t1**, this **bar chart** displays percent of packets lost:

- before and after enabling Forward Error Correction on tunnel (upper chart), and
- before and after enabling Packet Order Correction on tunnel (lower chart)



Tables for FEC and POC are separate, as shown below.

To export this table as a .csv (comma-separated values) file that you can open as an Excel spreadsheet, click **Export**. The **File Download** dialog box displays.

Time	Pre-POC OOO (%)	Post-POC OOO (%)	Pre-POC OOO Pkts	Post-POC OOO Pkts	Wan Rx Pkts
2009/03/25 14:29:00	0.000	0.016	0	19	115,600
2009/03/25 14:30:00	0.000	0.009	0	11	118,419
2009/03/25 14:31:00	0.000	0.012	0	14	118,495
2009/03/25 14:32:00	0.000	0.010	0	12	118,521
2009/03/25 14:33:00	0.000	0.013	0	15	115,563
2009/03/25 14:34:00	0.000	0.011	0	13	117,385
2009/03/25 14:35:00	0.000	0.015	0	18	117,763
2009/03/25 14:36:00	0.000	0.011	0	13	113,646
2009/03/25 14:37:00	0.000	0.007	0	9	117,645
2009/03/25 14:38:00	0.000	0.010	0	11	118,497
2009/03/25 14:39:00	0.000	0.012	0	14	115,996
2009/03/25 14:40:00	0.000	0.013	0	15	122,622
2009/03/25 14:41:00	0.000	0.010	0	12	118,521
2009/03/25 14:42:00	0.000	0.008	0	10	115,563
2009/03/25 14:43:00	0.000	0.014	0	16	117,385
2009/03/25 14:44:00	0.000	0.012	0	14	117,645
2009/03/25 14:45:00	0.000	0.009	0	11	118,497
2009/03/25 14:46:00	0.000	0.014	0	16	115,996
2009/03/25 14:47:00	0.000	0.009	0	11	122,622
2009/03/25 14:48:00	0.000	0.011	0	13	118,521
2009/03/25 14:49:00	0.000	0.016	0	19	115,563
2009/03/25 14:50:00	0.000	0.013	0	16	117,385
2009/03/25 14:51:00	0.000	0.014	0	16	117,645
2009/03/25 14:52:00	0.000	0.011	0	13	118,497

Time	Pre-FEC Loss (%)	Post-FEC Loss (%)	Pre-FEC Pkts Loss	Post-FEC Pkts Loss	Wan Rx Pkts
2009/03/25 14:29:00	0.101	0.000	117	0	115,600
2009/03/25 14:30:00	0.094	0.000	111	0	118,419
2009/03/25 14:31:00	0.113	0.000	135	0	118,495
2009/03/25 14:32:00	0.100	0.000	118	0	118,521
2009/03/25 14:33:00	0.100	0.000	116	0	115,563
2009/03/25 14:34:00	0.100	0.000	118	0	117,385
2009/03/25 14:35:00	0.106	0.000	120	0	113,646
2009/03/25 14:36:00	0.100	0.000	118	0	117,763
2009/03/25 14:37:00	0.097	0.000	118	0	121,576
2009/03/25 14:38:00	0.121	0.000	149	0	122,622
2009/03/25 14:39:00	0.110	0.000	126	0	114,622
2009/03/25 14:40:00	0.085	0.003	99	3	116,684
2009/03/25 14:41:00	0.111	0.000	131	0	117,645
2009/03/25 14:42:00	0.123	0.000	146	0	118,497
2009/03/25 14:43:00	0.114	0.000	132	0	115,996
2009/03/25 14:44:00	0.105	0.000	127	0	121,576
2009/03/25 14:45:00	0.083	0.000	97	0	117,645
2009/03/25 14:46:00	0.110	0.000	130	0	118,521
2009/03/25 14:47:00	0.100	0.000	120	0	115,563
2009/03/25 14:48:00	0.097	0.000	117	0	117,385
2009/03/25 14:49:00	0.111	0.000	131	0	118,497
2009/03/25 14:50:00	0.111	0.000	131	0	115,996
2009/03/25 14:51:00	0.111	0.000	131	0	122,622

♦ **To view the hourly Forward Error Correction statistics**

- 1 From the **Monitoring** menu, select **Network Integrity**.
- 2 From the **Tunnel** menu, select the name of the individual tunnel.
- 3 From the **Period** menu, select from **Realtime**, **Current Hour**, and **Last 60 Minutes**.
- 4 Click **Apply** to display the chart.
- 5 To view the data displayed as a table, click **FEC Table View** or **POC Table View**, as desired.

In the **FEC Table View**, the Appliance Manager charts the following items:

Field	Definition
Pre-FEC Loss (%)	Percentage of packets received from the WAN having errors before FEC was applied
Post FEC Loss (%)	Percentage of packets received from the WAN having errors after FEC was applied
Pre-FEC Pkts Loss	Number of error packets before Forward Error Correction (FEC) applied
Post-FEC Pkts Loss	Number of error packets after Forward Error Correction (FEC) applied
WAN Rx Pkts	Total number of WAN packets received



Tip If you enable **FEC** and see an increase in out-of-order packets in this **Monitoring** report, it indicates that you need to go back to the **Configuration - Tunnels** page and increase the **Reorder Wait** time from the default of 100ms assigned at tunnel creation.

In the **POC Table View**, the Appliance Manager charts the following items:

Field	Definition
Pre-POC OOO (%)	Percentage of out-of-order packets received from the WAN before applying POC
Post-POC OOO (%)	Percentage of out-of-order packets received from the WAN after applying POC
Pre-POC OOO Pkts	Number of error packets before applying Packet Order Correction (POC)
Post-POC OOO Pkts	Number of error packets after applying Packet Order Correction (POC)
Pre-FEC Pkts Loss	Number of error packets before Forward Error Correction (FEC) applied
WAN Rx Pkts	Total number of WAN packets received

Viewing Flow Redirection Statistics

The **Monitoring - Flow Redirection** page displays the number of control packets sent and received between this appliance and the other peer(s) in the cluster, as well as how much traffic was redirected to and from the other peer(s).

It answers the following questions:

- What are the **mgmt1** IP addresses of the other peers in this cluster?
- How many control packets were exchanged between this appliance and each of its peers?
- How many redirected flows does this appliance currently own?
- How many flows is this appliance currently redirecting to peers?
- In total, how many packets/bytes have been redirected to/from any given peer?

A Sampling of Results

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

When you choose **manually**, click the browser's **Refresh** icon to view up-to-the-minute data.

This appliance owns all the **Flows Redirected From**.

Appliance Time: 2009/03/24 16:35:03 US/Pacific Version: 2.4.1.0_25504
Logged in as: admin (logout) tiger 10.0.41.72

Monitoring - Flow Redirection

Selection Criteria
Refresh: manually Actual Stats Delta Stats
Traffic Unit: Bytes/Pkts/Msgs

Stats

Peer IP	Msg Type	Tx Msgs	Tx Bytes	Rx Msgs	Rx Bytes
192.168.10.60	Hello	8,482	76,338	8,482	76,338
	Redirection	20,001	520,026	1	26
192.168.30.50	Hello	8,120	73,080	8,116	73,044
	Redirection	20,003	520,078	352	14,404
192.168.40.50	Hello	8,050	72,450	8,047	72,423
	Redirection	20,003	520,078	0	0

Flows Redirected From

Peer	Flows	Pkts	Bytes
192.168.10.60	0	877,543	825,359,474
192.168.30.50	0	0	0
192.168.40.50	0	0	0

Flows Redirected To

Peer	Flows	Pkts	Bytes
192.168.10.60	0	0	0
192.168.30.50	0	0	0
192.168.40.50	0	0	0

Copyright © 2004-2007 Silver Peak Systems, Inc. All rights reserved.

The **Flows** columns list **current flows only**.

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

The **Monitoring - Flow Redirection** page displays the following statistics:

Section	Field	Definition
Stats	Peer IP	The mgmt1 IP address of a peer appliance in the same cluster as this appliance.
	Hello	Control packets used to keep open the TCP connection between two peers' mgmt1 cluster interfaces.
	Redirection	Requests to redirect flows
	Tx Msgs	The number of messages transmitted.
	Tx Bytes	The size of the transmitted messages, in bytes.
	Rx Msgs	The number of messages received.
	Rx Bytes	The size of the received messages, in bytes.
Flows Redirected From	Peer	The mgmt1 IP address of a peer appliance in the same cluster as this appliance.
	Flows	The number of current flows redirected from the peer to this appliance.
	Pkts	To date, the total number of packets redirected from the peer to this appliance.
	Bytes	To date, the total number of bytes redirected from the peer to this appliance.
Flows Redirected To	Peer	The mgmt1 IP address of a peer appliance in the same cluster as this appliance.
	Flows	The number of current flows redirected to the peer by this appliance.
	Pkts	To date, the total number of packets redirected to the peer by this appliance.
	Bytes	To date, the total number of bytes redirected to the peer by this appliance.

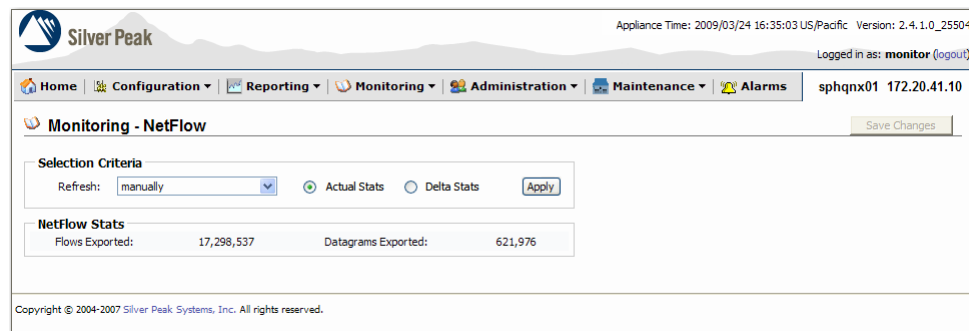
Across all other reported statistics, only the owner of a flow reports a flow's traffic statistics.

Viewing NetFlow Statistics

The **Monitoring - NetFlow** page displays the how many records were exported to the NetFlow collectors. It answers the following questions:

- How many flows were required to export the records to NetFlow?
- How many packets were required export these flows?

A Sampling of Results



The screenshot shows the Silver Peak web interface for the 'Monitoring - NetFlow' page. The top navigation bar includes links for Home, Configuration, Reporting, Monitoring (selected), Administration, Maintenance, and Alarms. The user is logged in as 'monitor' with a 'logout' link. The page title is 'Monitoring - NetFlow' and the IP address 'sphqnx01 172.20.41.10' is displayed. Below the title, there is a 'Selection Criteria' section with a 'Refresh' dropdown set to 'manually', radio buttons for 'Actual Stats' (selected) and 'Delta Stats', and an 'Apply' button. A 'Save Changes' button is also present. The 'NetFlow Stats' section displays two metrics: 'Flows Exported: 17,298,537' and 'Datagrams Exported: 621,976'. The footer contains the copyright notice: 'Copyright © 2004-2007 Silver Peak Systems, Inc. All rights reserved.'

NetFlow Stats	
Flows Exported:	17,298,537
Datagrams Exported:	621,976

Viewing Interface Statistics

The **Monitoring - Interfaces** page displays generic performance data for the actual physical LAN, WAN, and management interfaces (primary and secondary). It answers the following questions:

- How many bytes or packets is the appliance transmitting or receiving?
 - How many errors exist?
 - What types of errors exist?
- ◆ **To view the Interface Statistics**

From the **Monitoring** menu, select **Interfaces**. The page displays the actual statistics accumulated for **wan0**, **lan0**, **mgmt0**, and **mgmt1** since the appliance's last reboot.

If **Refresh** is set to **manually**, use the browser's refresh to view up-to-the minute data.

blan0 and **bwlan0** are visible when gigabit etherchannel bonding is configured.

For more information, see ["Configuring Gigabit Etherchannel Bonding" on page 99](#).

The screenshot shows the Silver Peak Monitoring - Interfaces page. The page has a navigation bar with the following tabs: Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The Monitoring - Interfaces tab is selected. Below the navigation bar, there is a section for Selection Criteria with a Refresh dropdown set to 'manually' and a Traffic Unit dropdown set to 'Bytes/Pkts'. There is an 'Apply' button next to these dropdowns. Below the Selection Criteria, there is a table of interface statistics. The table has two columns for each interface, one for Rx (Receive) and one for Tx (Transmit). The interfaces listed are blan0, bwlan0, lan0, wan0, lan1, wan1, mgmt0, and mgmt1. The statistics for each interface include Rx Bytes, Rx Pkts, Rx Discard Pkts, Rx Error Pkts, Rx Overrun Pkts, Rx MCast Pkts, Rx Frame Pkts, Tx Bytes, Tx Pkts, Tx Discard Pkts, Tx Error Pkts, Tx Overrun Pkts, Tx Carrier Pkts, and Tx Collision Pkts. The page also displays the Appliance Time (2009/03/24 16:35:03 US/Pacific), Version (2.4.1.0_25504), and Logged in as: admin (logout). There is a 'Save Changes' button in the top right corner.

Interface	Rx Bytes	Rx Pkts	Rx Discard Pkts	Rx Error Pkts	Rx Overrun Pkts	Rx MCast Pkts	Rx Frame Pkts	Tx Bytes	Tx Pkts	Tx Discard Pkts	Tx Error Pkts	Tx Overrun Pkts	Tx Carrier Pkts	Tx Collision Pkts
blan0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
bwlan0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
lan0	18,835,211,647	22,486,337	0	0	0	0	0	13,232,057,792	22,643,910	0	0	0	0	0
wan0	8,439,336,997	9,703,603	0	0	0	0	0	9,956,585,850	10,307,721	0	0	0	0	0
lan1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
wan1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
mgmt0	1,555,523	10,374	0	0	0	0	0	14,350,407	14,447	0	0	0	0	0
mgmt1	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Management connection to LAN

Management connection to PC

When the appliance is in **Bridge** mode, there are **lan0** interface statistics. When the appliance is in **Router** mode, there are not.

The **Monitoring - Interfaces** page displays the following statistics:

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

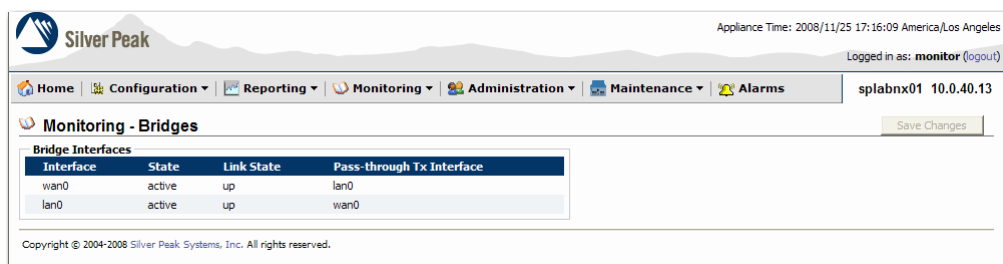
Field	Definition
Tx Bytes	Number of bytes transmitted outbound toward the WAN side
Tx Pkts	Number of packets transmitted outbound toward the WAN side, including all packets that were either discarded, contained errors, were overrun, had collisions, or were dropped because the interface detection link 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.

Viewing Bridge Mode Statistics

The **Monitoring - Bridge Mode** page displays data traffic traversing all the LAN and WAN interfaces. It answers the following questions:

- Is the appliance receiving/sending on all interfaces?
- Is the link up?

Two-Port Example

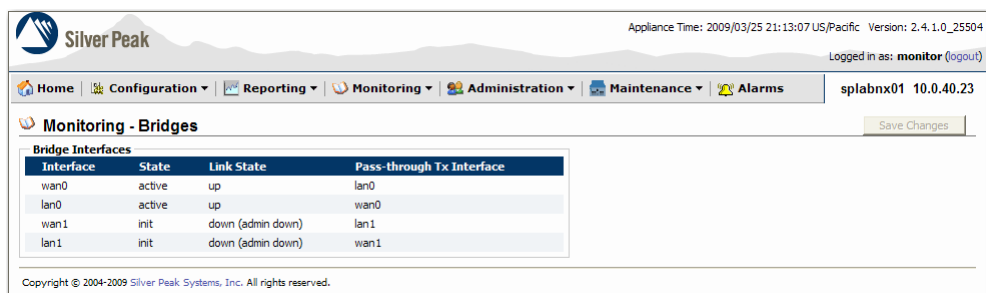


The screenshot shows the Silver Peak web interface. The top navigation bar includes links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as 'monitor'. The 'Monitoring - Bridges' section displays a table of bridge interfaces.

Interface	State	Link State	Pass-through Tx Interface
wan0	active	up	lan0
lan0	active	up	wan0

Copyright © 2004-2008 Silver Peak Systems, Inc. All rights reserved.

Four-Port Example



The screenshot shows the Silver Peak web interface for a four-port configuration. The 'Monitoring - Bridges' section displays a table of bridge interfaces.

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

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

The **Monitoring - Bridges** page displays the following 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.

Viewing IP Routes

The **Monitoring - IP Routes** page displays next-hop reachability.

It answers the following questions:

- Is the appliance receiving/sending on all interfaces?
- Is the link up?

Sampling of Results

Next-hop IP	Interface	Source	State	Uptime	WAN Configured Role	WAN Current Role
10.0.39.2	wan0	WAN	reachable	2d 13h 33m 3s	active	active
10.0.39.1	lan0	LAN	reachable	2d 13h 33m 2.999s	N/A	N/A

The **Monitoring - IP Routes** page displays the following statistics:

Field	Definition
Next-hop IP	IP address of the router to which the Silver Peak appliance sends datapath traffic
Interface	The 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	How long the next-hop router has been reachable
WAN Configured Role	Whether the next-hop router is Active or Backup . When Active , it's delivering tunneled packets.
WAN Current Role	Actual WAN role. The options are Active , Backup , Down , and N/A [not applicable].

A note about datapath connectivity

If there is either a **LAN Next-Hop Unreachable** or **WAN Next-Hop Unreachable** alarm after upgrading to Version 2.4.3.1, resolve the alarm(s) immediately by configuring the gateway(s) to respond to ICMP pings from the Silver Peak NX Appliance IP Address.



Administration Tasks

This chapter describes how to perform various administration-related tasks.

In This Chapter

- **Configuring Log Settings** See page 310.
- **Understanding the Events Log** See page 315.
- **Viewing a Log of All Alarms** See page 316.
- **Viewing the Audit Log** See page 317.
- **Managing Debug Files** See page 318.
- **Pre-Positioning Data for Enhanced Acceleration Benefits** See page 325.
- **Configuring SNMP** See page 327.
- **Managing User Accounts** See page 332.
- **Configuring Authentication, RADIUS, and TACACS+** See page 337.
- **Configuring Banners** See page 354.
- **Configuring Settings for Web Protocols and Web Users** See page 355.
- **Initial Configuration Wizard** See page 356.
- **Support** See page 358.

Configuring Log Settings

The **Administration - Log Settings** page allows you to set parameters related to the event and alarm logs. You can configure local and, optionally, remote logging parameters.

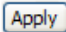
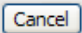
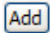
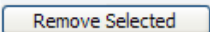
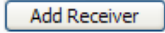
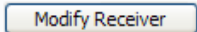
◆ **To access the Administration - Log Settings page**

From the **Administration** menu, select **Logging > Log Settings**.

The screenshot shows the Silver Peak Administration - Log Settings page. At the top, there's a navigation bar with tabs: Home, Configuration, Reporting, Monitoring, Administration (selected), Maintenance, and Alarms. The user is logged in as 'admin' and the appliance time is 2009/03/24 16:35:03 US/Pacific. The page title is 'Administration - Log Settings'. Below the title, there's a 'Local Log Settings' section with a dropdown for 'Minimum severity level' set to 'Notice', a text input for 'Start new file when log reaches' set to '50' (with a note '(1..50) MB'), and another text input for 'Keep at most' set to '30' (with a note '(1..100) log file(s)'). There are 'Apply' and 'Cancel' buttons. Below this is a 'Remote Log Receivers' section with a table. The table has two columns: 'Remote Receiver' and 'Minimum Severity'. One row is shown with the IP '172.20.2.106' and severity 'Notice'. There are 'Remove Selected' and 'Add' buttons. At the bottom, there's a copyright notice: 'Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.'

The **Administration - Log Settings** page displays the following information:

Field	Definition/Content
Minimum severity level	<p>The event level, listed here in decreasing order of severity:</p> <ul style="list-style-type: none"> • EMERGENCY The system is unusable. • ALERT Includes all alarms the appliance generates: CRITICAL, MAJOR, MINOR, and WARNING. <i>For more information, see Chapter 16, "Monitoring Alarms."</i> • CRITICAL A critical event • ERROR An error • WARNING A warning condition • NOTICE A normal, but significant, condition • INFORMATIONAL Informational. Used by Silver Peak for debugging. • DEBUG Used by Silver Peak for debugging • NONE Tells the appliance not to log anything to the Event Log. <p>[The bolded part of the level name is what displays in Appliance Manager.]</p> <p>If you select Notice, then the log records any event with a severity of NOTICE, WARNING, ERROR, CRITICAL, ALERT, and EMERGENCY.</p> <p>These are purely related to event logging levels, not alarm severities, even though some of the naming conventions overlap. Events and alarms have different sources. Alarms, once they clear, list as the ALERT level in the Event Log.</p>

Field	Definition/Content (Continued)
Start new file when log reaches	Specifies a maximum log size before a new log file is created. The default is 50 MB .
Keep at most	Sets a limit for the number of log files to store. The default is 30 .
	Apply all changes or edits to the running configuration, but not the stored configuration. To add the changes to the stored configuration, be sure to click Save Changes .
	Undo changes <i>if</i> you haven't clicked Apply .
Remote Log Receivers	These allow you to forward logs of all events at a given severity level and higher to a remote syslog server.
IP Address	The IP address of the remote syslog server.
Minimum Severity	The lowest severity level of all logged events to be forwarded.
	Displays the Add Remote Receiver area, where you can enter the remote syslog server's IP address and choose the Minimum Severity level of events to report.
	Removes the IP address of the remote syslog server from the destination list.
	Adds the IP address of the remote syslog server to the destination list of forwarded logged events.
	Applies the changes made to the IP address and/or severity level in the Modify Remote Receiver area.

Configuring Local Logging

You can configure log severity levels and rotation parameters on the **Administration - Log Settings** page.

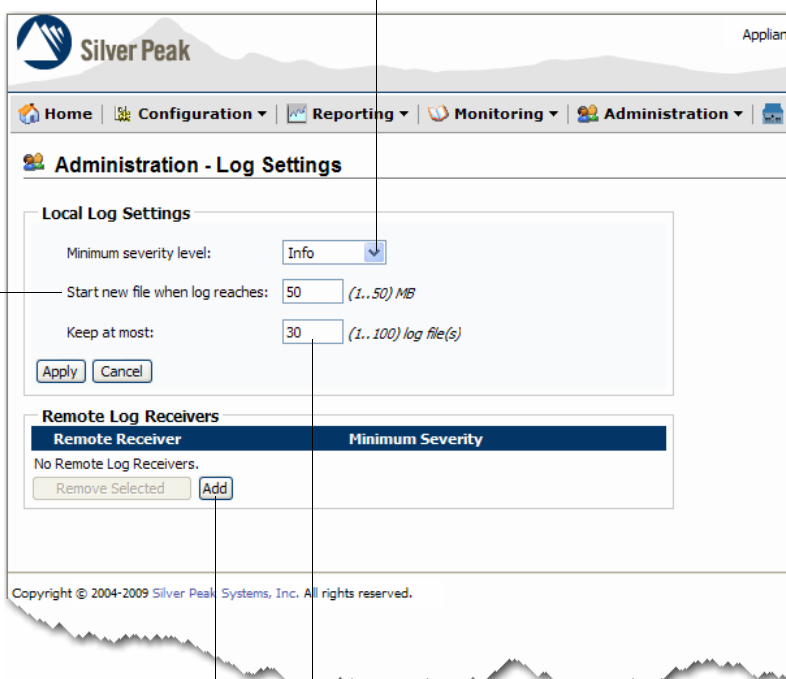
Select a **Minimum severity level** from the drop-down list.

- Silver Peak recommends that you select **NOTICE**, which is the default.
- If you select **NONE**, then no events are logged.

For severity definitions, see "Categories of Alarms" on page 402.

Choose when to begin a new log, based on maximum log size in megabytes.

The default maximum log size is **50 MB**, which means that the current log continues until it reaches that size. Of course, you can change the value.



The screenshot displays the 'Administration - Log Settings' page. Under 'Local Log Settings', the 'Minimum severity level' is set to 'Info'. The 'Start new file when log reaches' is set to '50' MB, and 'Keep at most' is set to '30' log files. The 'Remote Log Receivers' section shows no receivers are currently configured.

Click to display the **Add Remote Receiver** fields.

Enter the number of log files to retain. The default value is **30**: the current log plus 29 archives.

To save changes to the running configuration, click **Apply**.

Then, **Save** the changes to make them part of the stored configuration.

Configuring Remote Logging

You can configure the appliance to forward all events, at and above a specified severity, to a remote server that is already configured as a remote system log (**syslog**). The appliance forwards each qualified event as it occurs.

Independently, a syslog server is configured for the minimum severity level of event that it will accept. Without reconfiguring, it may not accept as low a severity level as you are forwarding.

♦ To forward events to a remote syslog server

1. In the Remote Log Sinks section of the **Administration - Log Settings** page, click **Add**.

2. In the newly displayed **Add Remote Receiver** section, enter the **IP address** and select the **Minimum Severity** level. All events at this level and above forward to the remote syslog server.

Remote Log Receivers

Remote Receiver	Minimum Severity
No Remote Log Receivers.	

Remove Selected Add

Add Remote Receiver

IP address:

Minimum Severity:

Add Receiver Cancel

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

3. Click **Add Receiver**. The new remote log server displays in the table.

Silver Peak Appliance

Home Configuration Reporting Monitoring Administration

Administration - Log Settings

Local Log Settings

Minimum severity level:

Start new file when log reaches: (1..50) MB

Keep at most: (1..100) log file(s)

Apply Cancel

Remote Log Receivers

Remote Receiver	Minimum Severity
<input type="checkbox"/> 172.20.2.106	Notice

Remove Selected Add

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

4. Click **Save Changes**.

If you want to suspend log forwarding to the remote server while leaving the IP address in the table, then change the **Minimum Severity** to **NONE**.

♦ To modify a remote receiver's Minimum Severity

You cannot modify an existing server's IP address. Instead, you must delete the old one and add a new one.

- 1 In the **Remote Log Receivers** area, click the IP address of the receiver you want to modify. The **Modify Remote Receiver (IP address)** appears.

The screenshot shows the Silver Peak Administration interface. At the top, there's a navigation bar with tabs: Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The 'Administration' tab is selected. Below the navigation bar, the page title is 'Administration - Log Settings'. There's a 'Save Changes' button in the top right corner. The main content area is divided into two sections: 'Local Log Settings' and 'Remote Log Receivers'. The 'Local Log Settings' section has a 'Minimum severity level' dropdown set to 'Notice', a 'Start new file when log reaches' field set to '50 (1..50) MB', and a 'Keep at most' field set to '30 (1..100) log file(s)'. There are 'Apply' and 'Cancel' buttons. The 'Remote Log Receivers' section contains a table with two columns: 'Remote Receiver' and 'Minimum Severity'. The table has one row with the IP address '172.20.2.106' and 'Notice' as the severity level. Below the table are 'Remove Selected' and 'Add' buttons. Below the table is a 'Modify Remote Receiver (172.20.2.106)' form. This form has an 'IP address' field with the value '172.20.2.106' and a 'Minimum Severity' dropdown set to 'Notice'. There are 'Modify Receiver' and 'Cancel' buttons at the bottom of the form.

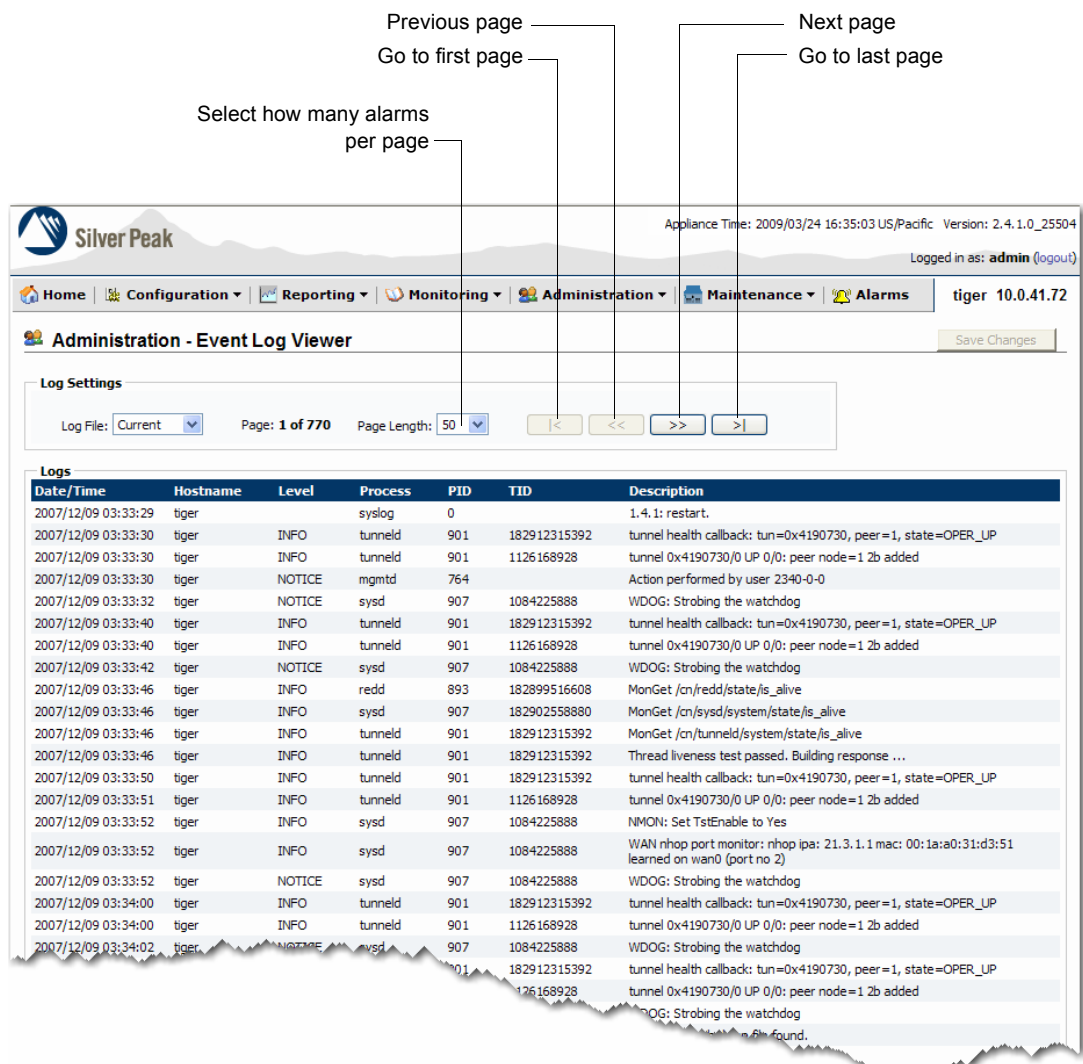
- 2 From the **Minimum Severity** drop-down menu, select the option you want.
- 3 Click **Modify Receiver**.
- 4 Click **Save Changes**.

♦ To remove a syslog server from the Remote Receiver list

- 1 Select the server in the table.
- 2 Click **Remove Selected**. The table clears the syslog server entry.
- 3 To save changes to the running configuration, click **Apply**.
- 4 **Save** the changes to the database.

Understanding the Events Log

The event log, which you access by selecting **Administration > Logging > Event Log Viewer**, contains timestamped messages for all system-level activity. It's a locally saved, read-only log:



You can configure the generic log settings on the **Administration - Log Settings** page. This includes specifying:

- The minimum severity level logged
- Whether time intervals or log size determine when a new log begins
- The maximum number of log files to keep, including the current log
- To what other remote servers to send logged events



For more information, see “Configuring Log Settings” on page 310.

Viewing a Log of All Alarms

The **Administration - Alarm Log Viewer** page displays all alarms—current and historical. It contains timestamped messages each time an alarm is raised or cleared. It is a locally saved read-only log.

To access the alarm log, select **Administration > Logging > Alarm Log Viewer**.

Previous page

Next page

Go to first page

Go to last page

Select how many alarms per page

Silver Peak

Appliance Time: 2009/03/24 16:35:03 US/Pacific Version: 2.4.1.0_25504

Logged in as: admin (logout)

Home Configuration Reporting Monitoring Administration Maintenance Alarms tiger 10.0.41.72

Administration - Alarm Log Viewer

Log Settings

Log File: Current Page: 1 of 23 Page Length: 50

Logs

Date/Time	Hostname	Level	Process	PID	Description
2007/09/07 22:36:08	silverpeak	ALERT	mgmt	630	ALARM RAISE: CRI,EQU,1, equipment_system_bypass,System is in BYPASS mode,system,2007/09/07 21:36:08,1,no,yes,no,no. NIC 2262 fail-to-wire mode - BYPASS
2007/09/07 22:36:23	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,2, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:36:23,1,no,yes,yes,yes. Network Interface (mgmt0) Link Down
2007/09/07 22:36:23	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,3, equipment_if_link_down,Network Interface Link Down,mgmt1,2007/09/07 21:36:23,1,no,yes,yes,yes. Network Interface (mgmt1) Link Down
2007/09/07 22:36:23	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,4, equipment_if_link_down,Network Interface Link Down,lan0,2007/09/07 21:36:23,1,no,yes,yes,yes. Network Interface (lan0) Link Down
2007/09/07 22:36:23	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,5, equipment_if_link_down,Network Interface Link Down,wan0,2007/09/07 21:36:23,1,no,yes,yes,yes. Network Interface (wan0) Link Down
2007/09/07 22:38:02	silverpeak	ALERT	mgmt	630	ALARM RAISE: CRI,SW,6, software_datapath_db_init,Datapath Database Loading Failed,Datapath,2007/09/07 21:38:02,1,no,yes,no,yes. Datapath Database Loading Failed
2007/09/07 22:38:02	silverpeak	ALERT	mgmt	630	ALARM CLEAR: CRI,EQU,7, equipment_system_bypass,System is in BYPASS mode,system,2007/09/07 21:36:08,1,no,yes,no,no. NIC 2262 fail-to-wire mode - NORMAL
2007/09/07 22:49:17	silverpeak	ALERT	mgmt	630	ALARM CLEAR: MAJ,EQU,8, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:36:23,2,no,yes,yes,yes. Network Interface (mgmt0) Link Up
2007/09/07 22:49:19	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,9, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:49:19,1,no,yes,yes,yes. Network Interface (mgmt0) Link Down
2007/09/07 22:49:23	silverpeak	ALERT	mgmt	630	ALARM CLEAR: MAJ,EQU,10, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:49:19,9,no,yes,yes,yes. Network Interface (mgmt0) Link Up
2007/09/07 22:49:24	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,11, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:49:24,1,no,yes,yes,yes. Network Interface (mgmt0) Link Down
2007/09/07 22:49:25	silverpeak	ALERT	mgmt	630	ALARM CLEAR: MAJ,EQU,12, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:49:24,11,no,yes,yes,yes. Network Interface (mgmt0) Link Up
2007/09/07 22:49:25	silverpeak	ALERT	mgmt	630	ALARM RAISE: MAJ,EQU,13, equipment_if_link_down,Network Interface Link Down,mgmt0,2007/09/07 21:49:24,12,no,yes,yes,yes. Network Interface (mgmt0) Link Down

You can configure the generic log settings on the **Administration - Log Settings** page. This includes specifying:

- The minimum severity level logged
- Whether time intervals or log size determine when a new log begins
- The maximum number of log files to keep, including the current log
- To what other remote servers to send logged events



For more information, see “Configuring Log Settings” on page 310.

Viewing the Audit Log

The **Administration - Audit Log Viewer** page lists all configuration changes (create, modify, delete) and all system actions such as login/logout made by any users (Command Line Interface [CLI], Appliance Manager, and/or Global Management System [GMS]).

username@IP
(/gms of GMS user)

Appliance Hostname

create /
modify /
delete /
action

SYSTEM /
INTERFACE /
ALARM /
CONFIG-DB

succeeded /
failed /
requested

Additional parameters
(context dependent)

Silver Peak

Appliance Time: 2009/08/12 16:16:33 US/Pacific Version: 3.14.159.0_27570
Logged in as: admin (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms | Tallinn2 10.0.40.36

Administration - Audit Log Viewer Save Changes

Log Settings

Log File: Current Page: 1 of 2 Page Length: 50

Logs

Date/Time	Hostname	Source	User	Operation	Object Type	Object Name	Status	Parameters
07/24 16:41:37	Tallinn2	web/1	admin@172.20.41.92	action	SYSTEM	SOAP/LOGIN	succeeded	
07/24 16:57:05	Tallinn2	web/1	admin@172.20.41.92	action	SYSTEM	SOAP/LOGOUT	succeeded	
07/28 13:19:15	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/28 13:20:00	Tallinn2	web/1	admin@172.20.41.92	action	SYSTEM	SOAP/LOGIN	succeeded	
07/28 13:20:43	Tallinn2	web/1	admin@172.20.41.92	modify	ACL	ad_1/10	succeeded	permit=true,ipaddr=10.10.10.10,mask_len=255.255.255.0,...
07/28 13:36:05	Tallinn2	web/1	admin@172.20.41.92	action	SYSTEM	SOAP/LOGOUT	succeeded	
07/28 13:54:30	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGOUT	succeeded	
07/28 13:57:17	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/28 14:22:03	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGOUT	succeeded	
07/28 14:33:58	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/28 15:33:06	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGOUT	succeeded	
07/28 16:04:53	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/28 16:18:02	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGOUT	succeeded	
07/28 16:18:08	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/28 16:52:35	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGOUT	succeeded	
07/28 16:53:34	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/28 17:28:08	Tallinn2	cli	admin@localhost	action	SYSTEM	CLI/LOGOUT	succeeded	
07/29 09:33:04	Tallinn2	cli	admin@172.20.41.21	action	SYSTEM	CLI/LOGIN	succeeded	type=interactive
07/29 09:48:04	Tallinn2	cli	admin@172.20.41.21	action	SYSTEM	CLI/LOGOUT	succeeded	
08/10 12:21:57	Tallinn2	web/1	admin@172.20.41.137	action	SYSTEM	SOAP/LOGIN	succeeded	
08/10 12:22:42	Tallinn2	web/1	admin@172.20.41.137	action	IMAGE	install	succeeded	image_name=/var/tmp/image-3.14.159.0_27570.img
08/10 12:24:15	Tallinn2	web/1	admin@172.20.41.137	action	IMAGE	boot_location	succeeded	location_id=2
08/10 12:24:20	Tallinn2	web/1	admin@172.20.41.137	action	SYSTEM	SOAP/LOGOUT	succeeded	
08/10 12:24:20	Tallinn2	web/1	admin@172.20.41.137	action	REBOOT		requested	
08/10 12:24:20	Tallinn2	web/1	admin@172.20.41.137	action	SYSTEM	flush_disk	succeeded	
08/10 12:24:20	Tallinn2	web/1	admin@172.20.41.137	action	SYSTEM	LOGIN	succeeded	

To access this page, select **Administration > Logging > View Audit Log**.

This log is only available to users with the Admin privilege level.

Managing Debug Files

This chapter describes how to manage the system files – log files, debug dump files, stat reports, and tcpdump results. It also describes how to archive them by saving them to an SCP (Secure Copy) or FTP (File Transfer Protocol) Server.

- **Types of Debug Files** See page 318.
- **Saving Files to a Remote Server** See page 320.
- **Deleting Log Files** See page 324.

Types of 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.

- The Appliance Manager's **Administration - Debug Files** page lists these files and provides a way for you to save them to another location for storage or additional handling.
- With the exception of the **[Log]** files, you cannot view these files on the Appliance Manager.
- To free up memory in the appliance, you can delete the files.

Specifically, these five file types are as follows:

Log	<p>The raw event log data, viewable on the Administration - Event Log Viewer page. This includes historical alarms, not current ones. To access this page, select Administration > Logging > Event Log Viewer.</p> <p>By default, a new file begins when the file reaches 50 MB. However, you can change the rotation criteria on the Administration - Log Settings page. To access this page, select Administration > Logging > Log Settings.</p>
Debug Dump	<p>Created as a result of any system failure.</p> <p>Can also be created on demand by clicking the Generate button next to the System & Debug Information File field.</p> <p>Transfer these files to Silver Peak's Customer Support for evaluation.</p>
Snapshot	<p>Created as a result of any system failure.</p> <p>Contains the same information as Debug Dump, and then includes additional information needed by the engineering team.</p> <p>Transfer these files to Silver Peak's Customer Support for evaluation.</p>
TCP Dump Result	<p>User-named data file generated by running using the Command Line Interface (CLI).</p> <p>Transfer these files to Silver Peak's Customer Support for evaluation.</p>
Show Tech	<p>Can also be created on demand by clicking the Generate button next to the Tech Support File field.</p> <p>Transfer these files to Silver Peak's Customer Support for evaluation.</p>

Shows the data disk space used and available, in bytes and as a percentage

To save a file to the local disk or an SCP or FTP remote server, click on the filename itself.

Silver Peak Appliance Time: 2009/04/14 17:40:06 US/Pacific Version: 2.4.2.0_25736
Logged in as: **admin** (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms | SP1 10.0.41.119

Administration - Debug Files

Save Changes

Administrative Data Information

Used: 10,969 MB Available: 10,168 MB
Percent Used: 52%

File Management

[Log] [Debug Dump] [Snapshot] [TCP Dump Result] [Show Tech]

File Name	Size (Bytes)
messages	18,876,268
alerts	54,815
alerts.1.gz	264
alerts.2.gz	
messages.2.gz	
alerts.3.gz	

Generate File

System & Debug Information File

Tech Support File

To **manually** generate Debug Dump file.

To generate the **Show Tech** file.

[Log] [Debug Dump] [Snapshot] [TCP Dump Result] [Show Tech]

File Name	Size (Bytes)
messages	19,206,676
alerts	54,815
alerts.1.gz	264
alerts.2.gz	20

Log

As a new log file is created, the earlier files increment. For example, the file **messages** will eventually be renamed **messages.3.gz**.

[Log] [Debug Dump] [Snapshot] [TCP Dump Result] [Show Tech]

File Name	Size (Bytes)
tunbug-SP1-20090414.tar	389,120
tunbug-SP1-20090410.tar.gz	437,530
tunbug-SP1-20090411.tar.gz	164,238
tunbug-SP1-20090413.tar.gz	253,394

Debug Dump

The file name format is **tunbug-[Hostname]-[YYYYMMDD-HHMMSS].tgz**

[Log] [Debug Dump] [Snapshot] [TCP Dump Result] [Show Tech]

File Name	Size (Bytes)
SP1-TR-tunneld-20081203-045339.tar.gz	404,487,421
SP1-TR-tunneld-20081204-130317.tar.gz	757,115,399

Snapshot

The file name format is **[Hostname]-[sysd|statsd|soapd|snmpd...]-[YYYYMMDD-HHMMSS].tar.gz**

[Log] [Debug Dump] [Snapshot] [TCP Dump Result] [Show Tech]

File Name	Size (Bytes)
My_File_3	328

TCP Dump Result

User-named file generated by running **tcpdump** in the Command Line Interface (CLI).

[Log] [Debug Dump] [Snapshot] [TCP Dump Result] [Show Tech]

File Name	Size (Bytes)
showtech-2009-04-14-175953.txt.gz	16,390

Show Tech

Generated when you execute the Tech Support File. The file name format is **[Hostname]-[YYYYMMDD-HHMMSS].txt.gz**

Saving Files to a Remote Server

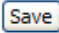
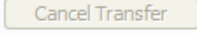
The Application Manager lets you copy non-configuration files from the appliance to a remote server. When you click to select the method, you can only edit the required fields.

Click to return to the main **Administration - Debug Files** page

The screenshot shows the Silver Peak web interface for saving debug files. The main heading is 'Administration - Debug Files - Save File'. Below this, there's a 'Save File' section where the file name is 'messages.3.gz'. There are three options for saving: 'Local File', 'SCP (Secure Copy)', and 'FTP (File Transfer Protocol)'. A 'Save' button is at the bottom of this section. Below that is the 'Transfer Save File Status' section, which shows the status as 'Ready' and a message 'The system is ready for upload'. There is a 'Cancel Transfer' button here. The top of the page has a navigation bar with links to Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as 'admin'.


The fields and options have the following definitions:

Field or Option	Definition/Content
File Name	(A read-only field) The name of the file you've chosen to save.
Save to Server with:	
Local File	For saving the software image file to your local PC.
SCP (Secure Copy)	For saving the software image file to a remote Secure Copy server.
FTP (File Transfer Protocol)	For saving the software image file to a remote File Transfer Protocol (FTP) server.
Remote Server Address	Use either the server IP address or the server name (if it's mapped to a local host table or a DNS server).
Remote User Name	The name of the user that server expects
Remote Password	The password of the user that the server expects
Remote Full Path Remote Relative Path	The type of path requested depends on which method you choose: <ul style="list-style-type: none"> If using the SCP server, enter the full path to the server. (Optional) If using the FTP server, enter the relative path to the server.
Destination File Name	(Optional) If you want to rename the file, you can do so here.
Status	If the read-only value is Ready , you may proceed with transferring the file to a remote server.

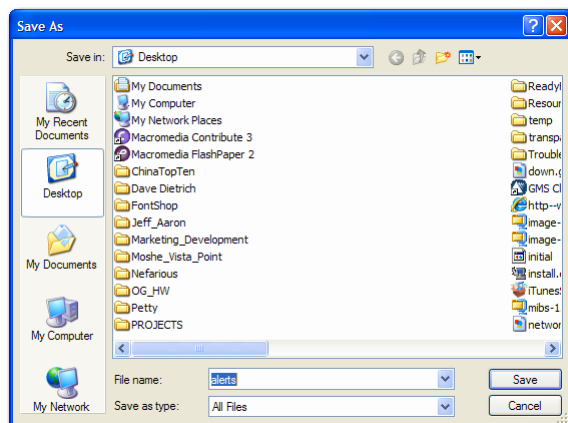
Field or Option	Definition/Content (Continued)
Last Save Status	The status at the end of the previous save operation.
Transfer Start Time	What time the file transfer began
Transfer End Time	What time the file transfer ended
	Saves the selected file/image to the remote server
	Allows you to cancel a file transfer that is in progress

♦ **To save a log file to a local PC**

- 1 Go to the **Administration - Debug Files** page.
- 2 In the **File Management** area, click the type of log you want to save: **Log**, **Debug Dump**, **Snapshot**, **Stat Report**, or **TCP Dump Result**.
- 3 In the table, click on the file name of the file you want to save.
The **Administration - Debug Files - Save File** page appears.
- 4 Click **Local File**.
- 5 Click **Save**. A security warning box displays.



- 6 Click **Save**.
- 7 When the **Save As** dialog box appears, navigate to the directory you want, accept or change the file name, and click **Save**.



♦ **To save a log file to an SCP Server**

- 1 Go to the **Administration - Debug Files** page.
- 2 In the **File Management** area, click the type of log you want to save: **Log**, **Debug Dump**, **Snapshot**, **Stat Report**, or **TCP Dump Result**.
- 3 In the table, click on the file name of the file you want to save.
The **Administration - Debug Files - Save File** page appears.
- 4 Click **SCP (Secure Copy)**.
- 5 Enter the data necessary to save the file to the SCP server.

Here, we'll use the example of saving the file, **alerts**, to the following location:

```
scp <UserName>@170.2.2.65:/home/<UserName>/work/logs/alerts
```

Save File

File Name: **alerts**

Save the selected file to the following specified:
[\[Local File\]](#) [\[SCP \(Secure Copy\)\]](#) [\[FTP \(File Transfer Protocol\)\]](#)

Remote Server Address:

Remote User Name:

Remote Password:

Remote Full Path:

Destination File Name: (Optional)

Transfer Save File Status

Status: **Ready**

- a For the **Remote Server Address** field, enter either:
 - the server IP address, as in **170.2.2.65**, or
 - the server name, if it's mapped to a local host table or a DNS server
 - b Enter the **Remote User Name** and **Remote Password** for the Secure Copy (SCP) server.
 - c For the **Remote Full Path** field, enter the *full* path.
A full pathname includes the drive (if required), starting or root directory, all attached subdirectories and ends with the file or object name. Begin the path with a forward slash (/).
 - d If you want to rename the file, enter the new file name in the **Destination File Name** field. If you leave the field blank, the Appliance Manager saves the file with its existing file name.
- 6 Click **Save**. The Appliance Manager displays the progress.

♦ **To save a log file to an FTP Server**

- 1 Go to the **Administration - Debug Files** page.
- 2 In the **File Management** area, click the type of log you want to save: **Log**, **Debug Dump**, **Snapshot**, **Stat Report**, or **TCP Dump Result**.
- 3 In the table, click on the file name of the file you want to save. The **Administration - File System - Save File** page appears.
- 4 Click **File Transfer to Protocol (FTP)**.
- 5 Enter the data necessary to save the file to the FTP server.

Here, we'll use the example of saving the file, **alerts**, to Andrew's directories on an FTP server. In the process, we'll rename the file to **alerts_HQ_NX-7500**:

The screenshot shows a web form titled "Administration - Debug Files - Save File". At the top, it says "File Name: alerts". Below that, it says "Save the selected file to the following specified:" with three links: "[Local File]", "[SCP (Secure Copy)]", and "[FTP (File Transfer Protocol)]". The "FTP" link is highlighted. Below these links are several input fields: "Remote Server Address" with the value "170.2.2.65", "Remote User Name" with the value "andrew", "Remote Password" with masked characters "••••••", "Remote Relative Path" with the value "work/logs" and "(Optional)" to its right, and "Destination File Name" with the value "alerts_HQ_NX-7500" and "(Optional)" to its right. At the bottom of the form is a "Save" button. Below the form is a section titled "Save File Status" which is partially obscured by a jagged line.

For ftp, **no slash** necessary before or after the directory name

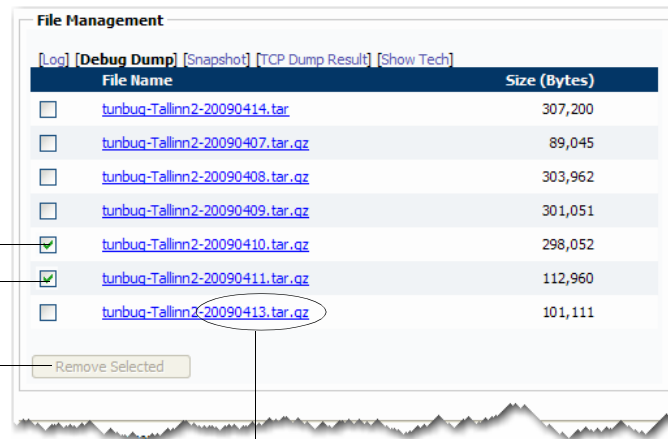
If you want to rename the original file, enter it in this field. If you leave it blank, the file saves with its existing name.

- a For the **Remote Server Address** field, enter either:
 - the server IP address, as in **170.2.2.65**, or
 - the server name, if it's mapped to a local host table or a DNS server
 - b Enter the **Remote User Name** and **Remote Password** for the FTP server.
 - c For the **Remote Relative Path** field, enter the *relative* path.
 A *relative* path is a path relative to the current working directory. Its first character can be anything but the pathname separator (here, a forward slash).
 For example, if the ftp login directory is **/home/<UserName>/**, then the *relative* path would begin at the next subdirectory, as in, **work/logs**. It is **not** necessary to begin or end the relative path with a forward slash (/).
 - d If you want to rename the file, enter the new file name in the **Destination File Name** field. If you leave the field blank, the Appliance Manager saves the file with its existing file name.
- 6 Click **Save**. The Appliance Manager displays the progress.

Deleting Log Files

All logs files are removed the same way.

Click the unchecked box(es) to select the file(s) you want to delete...



The **Debug Dump** file's creation date is encoded in the file name. Here, the file, **tunbug-Tallinn2-20090413.tar.gz**, was created on April 13, 2009.

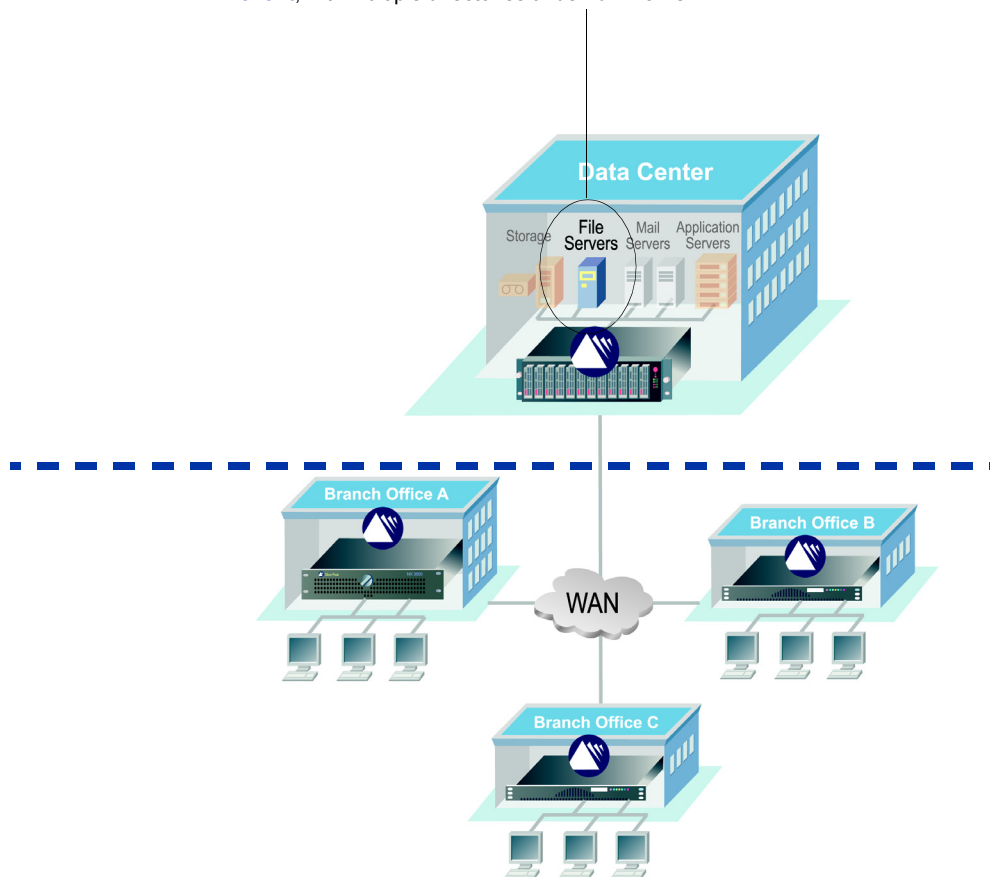
...then click **Remove Selected**. The Appliance Manager deletes the files from the appliance.

Pre-Positioning Data for Enhanced Acceleration Benefits

This section describes how to use the NX Series appliances' FTP server capability to pre-position data and enable all users to get the benefit of second-pass network performance.

The Appliance Manager allows you to pre-position data into Network Memory so that users can get the benefits of second-pass data without having to wait for Network Memory to populate.

In this scenario, the file server at the data center is the **FTP client**, with multiple directories under **/dir/home**.



With the **Administration - Pre-position** page, you can enable **FTP server** capability on each branch's NX appliance:

- Once you enable pre-positioning on the appliance, administrators can then FTP files or directories to the appliances which will, in turn, "warm" Network Memory.

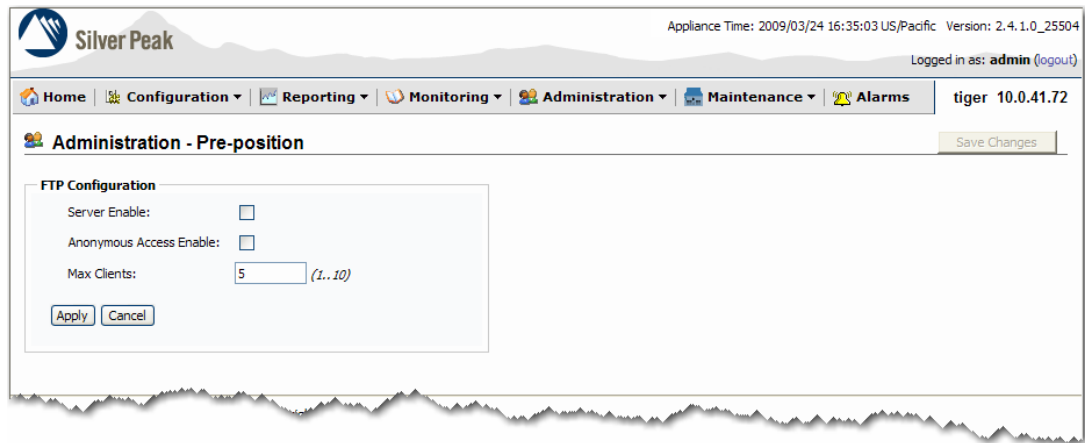
Subsequently, any user who requests data that was pre-positioned will immediately enjoy the acceleration benefits of the stored local instances.

- It's important to make sure that the relevant tunnels are admin-ed up before FTP transfer.
- If the Administrator desires, (s)he can set up a script process for prepositioning jobs that need to run automatically on an ongoing basis.

There is no down side to leaving this feature enabled by default.

♦ **To configure a branch or remote appliance to pre-position data**

- 1 From the **Administration** menu, select **Pre-position**.



The screenshot shows the Silver Peak web interface. At the top, the Silver Peak logo is on the left, and the appliance time (2009/03/24 16:35:03 US/Pacific) and version (2.4.1.0_25504) are on the right. Below the header is a navigation bar with tabs: Home, Configuration, Reporting, Monitoring, Administration (selected), Maintenance, and Alarms. The user is logged in as 'admin' with a 'logout' link. The main content area is titled 'Administration - Pre-position' and includes a 'Save Changes' button. Under the 'FTP Configuration' section, there are three settings: 'Server Enable' (checkbox), 'Anonymous Access Enable' (checkbox), and 'Max Clients' (text input with '5' and a range '(1..10)'). 'Apply' and 'Cancel' buttons are at the bottom of the configuration box.

- a Select the **Server Enable** check box. This enables the appliance to act as an FTP server.
 - b To exempt the FTP client from needing to use an existing account on the appliance, select **Anonymous Access Enable**.
 - c Enter the maximum number of clients that may access the appliance simultaneously. The default value is **5**. The range is **1** to **10**.
- 2 Click **Apply**.
 - 3 Click **Save Changes**.

Configuring SNMP

This section describes the following about Simple Network Management Protocol (SNMP):

- **Loading SNMP MIBs** See page 327.
- **Configuring SNMP Settings** See page 328.

Loading SNMP MIBs

When you took delivery of your appliance, the package also contained a CD with 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 bold blue 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

Configuring SNMP Settings

This section describes the following SNMP-related procedures:

- **To add a trap receiver** See page 329.
- **To modify a trap receiver** See page 330.
- **To remove a trap receiver** See page 330.
- **To add an SNMP v3 user** See page 331.

On the **Administration - SNMP** page, you can configure the Appliance Manager for SNMP:

- The NX Series appliance is implemented as a host that supports data in the Management Information Base (MIB) II, as described in RFC 1213, for cold start traps and warm start traps.
- The appliance issues a MIB-II trap during reset—that is, when loading a new image, recovering from a crash, or rebooting.
- You can specify trap receivers—that is, the IP address destinations.
- The appliances send traps every time an alarm is raised or cleared. The 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. Please see SILVERPEAK-MGMT-MIB.TXT in the MIBS directory of the documentation distribution for additional information.

This area and the one below it are generic — an all that's needed for SNMP v1 and SNMP v2c.

For additional security, you can enable **admin** (only) for SNMP v3, which requires authentication between the SNMP server and the user. Additionally, you can encrypt and an optional encryption component.

This area is manages trap receivers, that is, IP address destinations.

♦ **To add a trap receiver**

- 1 Click **Add**. The **Add Trap Receiver** area appears.

Copyright © 2004-2007 Silver Peak Systems, Inc. All rights reserved.

- a In the **Trap Receiver IP** field, enter the IP address where you want the traps sent.
 - b In the **Community** field, enter the same community string that is configured on the agent.
 - c From the **Trap Type** field, select either **v1** (RFC 1157) or **v2c** (RFC 1901) standards.
In SNMP v1 (RFC 1157) and v2c (RFC 1901) standards, authentication is based on a community string (text string) that represents an unencrypted username without a password.
 - d The **Enable Trap Receiver** box is checked by default.
- 2 Click **Apply**. The new host displays in the table.

The new trap receiver appears in the table.

Unless you change the setting, the server is enabled by default.

- 3 Click **Save Changes**.

♦ To modify a trap receiver

You can modify a trap receiver's **Community** and **Trap Type**, as well as enable or disable the host.

You cannot modify a **Trap Receiver IP**. Rather, you must delete the host and add a new one.

- 1 In the **Trap Receivers** table, click the desired **Host IP** address.



Host	Community	Version	Enabled
<input type="checkbox"/> 172.20.2.191	textstring	v1	yes

Remove Selected Add

The **Modify Trap Receiver** area appears, specific to that IP address.



Host	Community	Version	Enabled
<input type="checkbox"/> 172.20.2.191	textstring	v1	yes

Remove Selected Add

Modify Trap Receiver (172.20.2.191)

Trap Receiver IP: 172.20.2.191

Community: public

Trap Type: v1

Enable Trap Receiver: ☒

Apply Cancel

The default value for **Community** is **public**.
However, you can edit this text string, as needed.

- 2 Do **one** of the following, as needed:

- a Select the server and click **Enable Server**. The value in the **Enabled** column changes from **false** to **true**.
- b Select the server and click **Disable Server**. The value in the **Enabled** column changes from **true** to **false**.

♦ To remove a trap receiver

- 1 In the **Trap Receivers** table, select the desired server(s).
- 2 Click **Remove Selected**.
A dialog window displays, requesting confirmation.
- 3 Click **OK**. The Appliance Manager removes the selected hosts.

♦ **To add an SNMP v3 user**

Authentication between the user and the server acting as the SNMP agent is bilateral and **required** for SNMP v3. You can use either the MD5 or SHA-1 hash algorithm.

Using DES or AES-128 for **privacy** is optional. If you don't specify a password, it uses the same privacy password (encrypted or plain text) that you specified for authentication.

- When entering a plain text password, leave **Password Encrypted** deselected.
- When entering an encrypted password, select **Password Encrypted**.

Only the user, **admin**, can be configured to be an SNMP v3 user.

Any user with **admin** privileges (capability) can configure the SNMP v3 privileges for **admin**, the user.

1 Select **Enable User**.

If you'd prefer, you can configure and save the authorization and privacy settings without enabling the user until a later date.

2 In the **Auth Information** area:

- From the **Type** box, select either the **SHA-1** or **MD5** hash algorithm.
- Indicate whether the password you'll enter is plain text or encrypted.
 - If you're entering a plain text password, leave the **Password Encrypted** check box deselected.
 - If you're entering an encrypted password, select the **Password Encrypted** check box.

c Enter the password.

3 If you leave the **Private Information** area unchanged and blank, the appliance uses the default privacy algorithm (**AES-128**) and the same password that's specified for authentication. If you want to use the **DES** hash algorithm and/or a different password, then, in the **Private Information** area:

- From the **Type** box, select either the **DES** or **AES-128** hash algorithm.
- Indicate whether the password you'll enter is plain text or encrypted.
 - If you're entering a plain text password, leave the **Password Encrypted** check box deselected.
 - If you're entering an encrypted password, select the **Password Encrypted** check box.

c Enter the password.

4 Click **Apply**.

5 **Save** the changes.

Managing User Accounts

This section provides information about using the **Administration - Users** page for creating, deleting, and editing user accounts.

- **Guidelines for Creating Passwords** See page 332.
- **Accessing User Accounts** See page 333.
- **Creating a User Account** See page 334.
- **Modifying a User Account** See page 335.
- **Deleting a User Account** See page 336.

When you first install the Silver Peak appliance, *admin* and *monitor* are the only users: *admin* has all privileges and *monitor* has a limited subset.

However, you can add an additional layer of security by setting up authentication and authorization with the use of RADIUS and TACACS+. For more information, see *“Configuring Authentication, RADIUS, and TACACS+” on page 337*.

Guidelines for Creating Passwords

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

Accessing User Accounts

The Silver Peak appliances' **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.

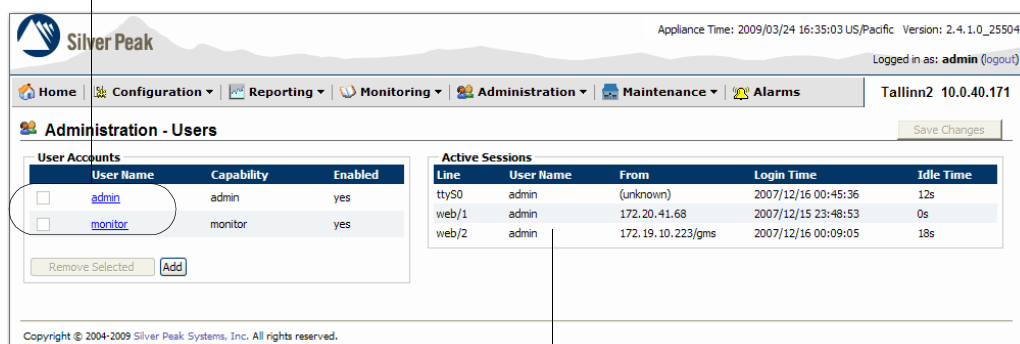
- The **monitor** group supports reading and monitoring of all data, in addition to performing all actions. This is equivalent to the Command Line Interface's (CLI) **enable** mode privileges.
- The **admin** group supports full privileges, along with permission to add, modify, and delete. This is equivalent to the CLI's **configuration** mode privileges.

♦ To access the Users page

From the **Administration** menu, select **User Management > Users**. The **Administration - Users** page displays.

The table lists all users known to this appliance, whether or not their accounts are enabled.

- The users, **admin** and **monitor**, **CANNOT** be deleted or disabled. Their check boxes are inaccessible.
- You can, however, change each one's password.



Appliance Time: 2009/03/24 16:35:03 US/Pacific Version: 2.4.1.0_25504
Logged in as: **admin** (logout)

Home Configuration Reporting Monitoring Administration Maintenance Alarms Tallinn2 10.0.40.171

Administration - Users

Save Changes

User Accounts		
User Name	Capability	Enabled
<input type="checkbox"/> admin	admin	yes
<input type="checkbox"/> monitor	monitor	yes

Remove Selected Add

Active Sessions				
Line	User Name	From	Login Time	Idle Time
tty50	admin	(unknown)	2007/12/16 00:45:36	12s
web/1	admin	172.20.41.68	2007/12/15 23:48:53	0s
web/2	admin	172.19.10.223/gms	2007/12/16 00:09:05	18s

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

The **Active Sessions** section lists all users who are logged in to the appliance. In order, these logins originate from the following locations:

- RS-232 serial port (console)
- Appliance Manager
- Global Management System (GMS)

Creating a User Account

To create a new user account, you must be logged in with **admin** group privileges.

◆ To create a new user account

- 1 On the **Administration - Users** page, click **Add**. The **Add New User** area displays.

The screenshot shows the Silver Peak web interface. The top navigation bar includes links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as 'admin'. The 'Administration - Users' section is active, displaying a 'User Accounts' table with columns for User Name, Capability, and Enabled. The table lists 'admin' and 'monitor' users, both enabled. Below the table is an 'Add New User' form with fields for User Name (Chris), Capability (admin), Password (masked), and Confirm Password (masked). There is an 'Enable User' checkbox and 'Apply' and 'Cancel' buttons. An 'Active Sessions' table is also visible on the right.

Line	User Name	From	Login Time	Idle Time
tty50	admin	(unknown)	2007/12/16 00:45:36	12s
web/1	admin	172.20.41.68	2007/12/15 23:48:53	0s
web/2	admin	172.19.10.223/gms	2007/12/16 00:09:05	18s

- 2 In the **User Name** field, enter a new user name.
- 3 In the **Capability** field, select **admin** or **monitor** from the drop-down menu.
 - The **monitor** group supports reading and monitoring of all data, in addition to performing all actions. This is equivalent to the Command Line Interface's (CLI) **enable** mode privileges.
 - The **admin** group supports full privileges, along with permission to add, modify, and delete. This is equivalent to the CLI's **configuration** mode privileges.
- 4 In the **Password** field, enter a password; the appliance has no constraints. Reenter the same sequence in the **Confirm Password** field.
- 5 Click **Apply**. The new user displays in the **User Accounts** table. By default, each new user is enabled.

New user added to the table

The screenshot shows the Silver Peak web interface after adding a new user. The 'User Accounts' table now includes a new entry for 'Chris' with 'admin' capability and 'yes' for 'Enabled'. The 'Add' button is highlighted. The 'Active Sessions' table remains on the right.

Line	User Name	From	Login Time	Idle Time
tty50	admin	(unknown)	2007/12/16 00:45:36	12s
web/1	admin	172.20.41.68	2007/12/15 23:48:53	0s
web/2	admin	172.19.10.223/gms	2007/12/16 00:09:05	18s

If you want to disable a user, you must modify the account.

- 6 To permanently save the edits, first click **Apply** and then **Save Changes**.

Modifying a User Account

To edit a user account, you must be logged in with **admin** group privileges.

When editing an account, you **can**:

- change the user's password (you can change your own if you have **admin** privileges.)
- enable or disable a user account
- remove a user account

You **cannot**:

- modify a user's name
- increase or decrease the user's level of privileges

To accomplish that, you must add a new name and delete the old name. The order in which you do it is not important.

You also **cannot** have more than one **user** with the same User Name.

♦ To edit a user account

- 1 In the **User Accounts** table, click the user's name.

We'll modify the account, **Chris**.

The screenshot shows the Silver Peak Administration interface. The 'User Accounts' table lists three users: Chris, admin, and monitor. The 'Chris' user is selected. The 'Modify User (Chris)' form is displayed below the table, showing fields for User Name, Capability, Change Password, Old Password, New Password, and Confirm Password. The 'Enable User' checkbox is checked. The 'Active Sessions' table is also visible on the right.

Line	User Name	From	Login Time	Idle Time
tty50	admin	(unknown)	2007/12/16 00:45:36	12s
web/1	admin	172.20.41.68	2007/12/15 23:48:53	0s
web/2	admin	172.19.10.223/gms	2007/12/16 00:09:05	18s

The name of the selected account displays in the header.

- 2 To change the password:
 - a Click the **Change Password** box.
 - b In the **Old Password** fields, enter the existing password.
 - c Then enter the new password in the **New Password** and **Confirm Password** fields.
- 3 To disable a user account, clear the **Enabled User** check box.
- 4 To permanently save the edits, first click **Apply** and then click **Save Changes**. The table updates its information.

Deleting a User Account

To delete a user account, you must have logged in with **admin** group privileges.

♦ To delete a user account

- 1 In the **User Accounts** table, click the check box to the left of the user's name.

We'll delete the account, **Chris**.

The screenshot shows the Silver Peak Administration - Users interface. The 'User Accounts' table has the following data:

User Name	Capability	Enabled
<input checked="" type="checkbox"/> Chris	admin	yes
<input type="checkbox"/> admin	admin	yes
<input type="checkbox"/> monitor	monitor	yes

The 'Active Sessions' table has the following data:

Line	User Name	From	Login Time	Idle Time
tty50	admin	(unknown)	2007/12/16 00:45:36	12s
web/1	admin	172.20.41.68	2007/12/15 23:48:53	0s
web/2	admin	172.19.10.223/gms	2007/12/16 00:09:05	18s

The 'Remove Selected' button is highlighted in the 'User Accounts' table.

- 2 Click **Remove Selected**, which is now accessible.

A dialog box appears, asking you to confirm the deletion(s).

- 3 Click **OK**. The table updates its information.
- 4 To permanently save the edits, click **Save Changes**.

Configuring Authentication, RADIUS, and TACACS+

This section discusses the following:

- **Authentication and Authorization** See page 337.
- **Session Idle Time-out** See page 338.
- **Configuring for RADIUS** See page 339.
- **Configuring for TACACS+** See page 346.

The Silver Peak NX appliances support user authentication and authorization before providing access rights.

- **Authentication** is the process of validating that the end user, or a device, is who they claim to be.
- **Authorization** is the action of determining what a user is allowed to do. Generally, authentication precedes authorization.

The configuration specified for authentication and authorization applies globally to all users accessing that appliance.

Additionally, if a logged-in user is inactive for an interval that exceeds the inactivity time-out, the appliance returns them to the *Login* page. You can change the value on the **Administration - Web** page.

Authentication and Authorization

To provide authentication and authorization services, Silver Peak appliances:

- support a built-in user 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 **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 CLI's *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.

TACACS+

TACACS+ uses TCP for 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.

♦ **To configure the authentication and authorization methods for all users of the appliance**


From the **Administration** menu, select **User Management > Authentication**. The page displays.

Authorization Information:

Use this section when authenticating users via RADIUS or TACACS+.

The Silver Peak appliance has a database of user names and passwords. Similarly, when you use RADIUS or TACACS+, those servers also have their own database of user names and passwords. The question then becomes how to “map” the local user information to the server’s variant of the same data, and to determine whether the appliance or the server has the “last word”.

The authorization policy, here, determines how the remote user mapping behaves.




Authentication Method List:

These settings tell the appliance to attempt authentication in the order shown. If you don't require a different second or third method, then reuse the first method in those fields.

For example, if you only want to use the Local appliance's database for authentication, then enter **Local** for the **First Method** and select **None** for the second and third methods. [This is the default setting.].

No matter how many methods you choose, at least one of them must be Local.

Map Order:




Local Only: It maps all remote users to the Default User.

Remote First: This is the default behavior. If the local user name is valid, then it maps the authenticated user to the local user. If the User Name is not valid locally, it uses the Default User.

Remote Only: If the local user name is valid, it maps only to a remote authenticated user. If the local user name is not valid, it attempts no further mapping.

Default User:



Here, you define the **Default User** (referred to by the Map Order task) as either **admin** or **monitor**.

Session Idle Time-out

The session idle time-out is the amount of time a user must be inactive before the Silver Peak appliance forces the user to login all over again. You can set this on the **Administration - Web** page; the default is **30 minutes**.

Configuring for RADIUS

If you want to use RADIUS as one of your authentication methods, make you first go to the **Administration - Authentication** page and configure RADIUS as one of the methods.

On this page, there are two places where you can potentially enter a **Server Key** — under **General Settings** and in the fields that appear when you add a new RADIUS server:

- If all (or a majority) of your servers have the same Server Key (also known as, *shared secret*), then enter that value in the **Server Key** field located in the **General Settings** section of this page. Then, when you're adding an individual server that has that same key, you can leave the specific server's **Server Key** fields blank. The software defers to the value in **General Settings**.
- If the server you're adding has a different Server Key, then enter that value to the information in the **Add RADIUS Server** section. An specifically added server key always outranks a **General Settings** Server Key value.

- 1 From the **Administration** window, select **User Management > RADIUS**. The **Administration - RADIUS** page appears.

The screenshot shows the Silver Peak Administration interface for RADIUS configuration. At the top, there's a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration (selected), Maintenance, and Alarms. The user is logged in as 'admin'. The main section is titled 'Administration - RADIUS'. It contains a 'General Settings' section with the following fields: 'Server Key' (empty), 'Confirm Server Key' (empty), 'Timeout' (3 seconds), and 'Retries' (1). There is an 'Apply' button next to the Confirm Server Key field. Below the General Settings is a table titled 'RADIUS Servers' with columns: Server IP, Port, Key, Timeout, Retries, and Enabled. The table is currently empty, with a message 'No RADIUS Servers.' and buttons for 'Remove Selected' and 'Add'.

- 2 Configure the **General Settings**, which apply to all RADIUS servers in your list. You can have a maximum of two RADIUS servers:
 - a In the **Server Key** field, enter the shared secret that was defined when setting up the RADIUS server.
 - b In the **Confirm Server Key** field, re-enter the same text string.
 - c If you wish, change the **Timeout** value. The default is 3 seconds.
 - d In the **Retries** field, specify how many retries to allow the user before locking them out.
 - e Click **Apply**.

- 3 Click **Add**. The **Add RADIUS Server** area appears.

The screenshot shows the Silver Peak Administration interface for RADIUS configuration. The top navigation bar includes links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as 'admin' with the session ID 'tiger 10.0.41.72'. The main content area is titled 'Administration - RADIUS' and features a 'Save Changes' button. The 'General Settings' section contains fields for 'Server Key', 'Confirm Server Key', 'Timeout' (set to 3 seconds), and 'Retries' (set to 1). Below this is the 'RADIUS Servers' table, which is currently empty. At the bottom is the 'Add RADIUS Server' form, which includes fields for 'Server IP', 'Authentication Port' (set to 1812), 'Server Key', 'Confirm Server Key', 'Timeout' (set to 3 seconds), 'Retries' (set to 1), and an 'Enabled' checkbox (checked). The 'Add' button is highlighted.

- a In the **Server IP** field, enter the IP address for the RADIUS server.
 - b In the **Server Key** field, do one of the following:
 - Leave it blank if this server's key is the same as the entry you made for **Server Key** in the **General Settings** section above, or
 - Assign a individual key, different from the one you entered in **General Settings**. Note that it must be the string that matches the one assigned to this RADIUS server.
 - c In the **Timeout** field, the default time period is 3 seconds. You can select from the range between 1 and 15 seconds.
 - d In the **Retries** field, enter how many retries to allow the user before locking them out. The default is 1 retry.
- 4 Click **Apply**. The server appears in the **RADIUS Servers** table.
 - 5 Click **Save Changes**.

Refer to your RADIUS documentation for additional details.

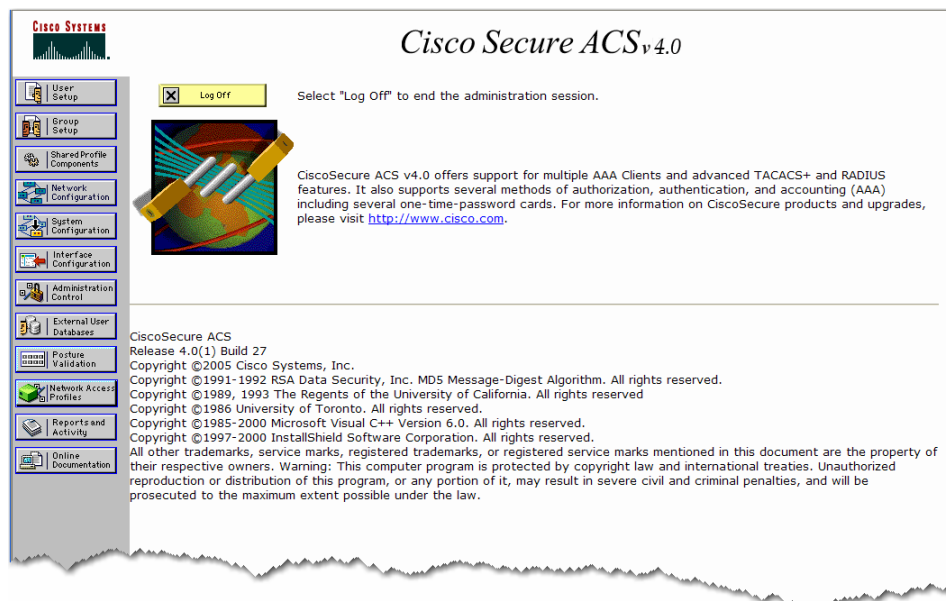
DEMO 14.1**Integrating the NX appliance into an existing RADIUS network.**

Before you begin, obtain the RADIUS shared/secret key from your network administrator. In the NX screens, we call this same item the **Server Key**.

First, we configure the RADIUS server. In this example, we're using Cisco IOS. And, we'll configure the NX appliance to use RADIUS.

Configuring the RADIUS Server

- 1 Login to your RADIUS server.



- 2 In the left column, select **Network Configuration**.

- a In the **AAA Servers** section, make sure that you have a server with the **AAA Server Type** of **RADIUS**.
- b If you haven't added a RADIUS server, do it now. For assistance, see your network administrator.

- 3 Stay on, or return to, the **Network Configuration** page.

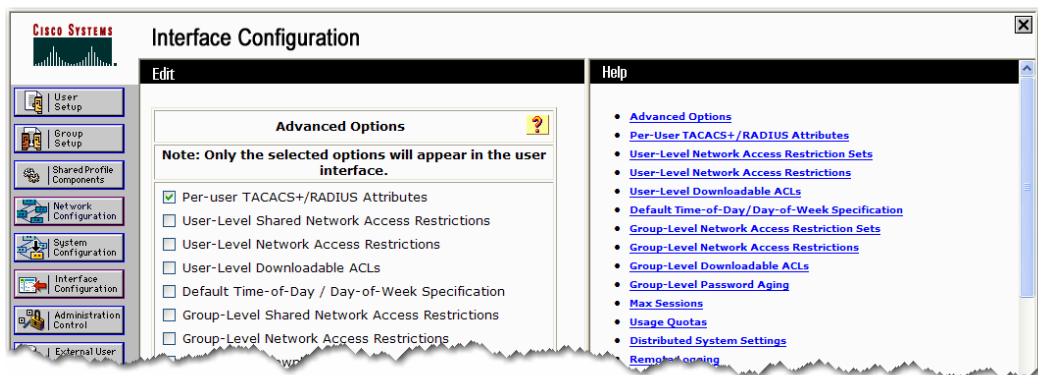
In the **AAA Clients** section, click **Add Entry**. The **Add AAA Client** page appears.

- a In the **AAA Client Hostname** field, enter the name assigned to the NX appliance. Here, we'll use **MyNX-Appliance**.
 - b In the **AAA Client IP Address** field, enter the IP address of the NX appliance. Here, we'll use **10.10.10.2**.
 - c In the **Key** field, enter the shared secret that both the AAA server (here, this RADIUS server) and the NX appliance (acting as the AAA client) will use to communicate. This entry is case sensitive. Here, we'll use **useAlsoInNXscreen**.
 - d From the **Authenticate Using** list, select **RADIUS (CISCO IOS/PIX 6.0)**.
 - e Click **Submit + Apply**.
- 4 In the left column, select **Interface Configuration**. The **Interface Configuration** page appears.

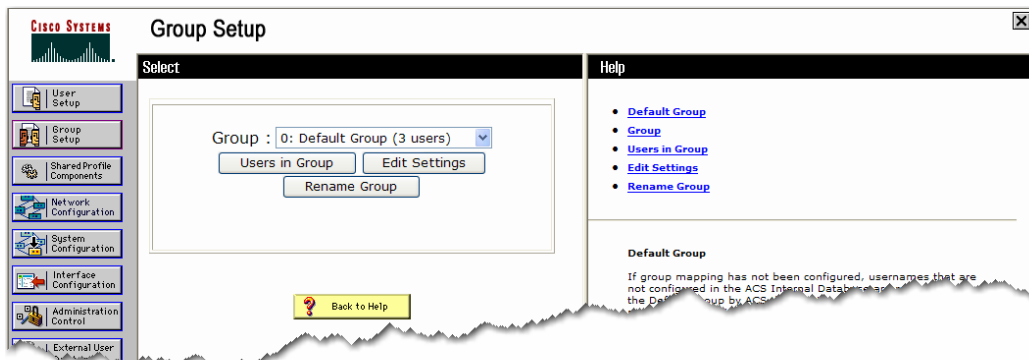
- 5 Select **Advanced Options**. The **Advanced Options** page appears.

- 6 Select **Per-user TACACS+/RADIUS Attributes** and click **Submit**.

Whether or not any other options need to be checked is a function of your own network. Therefore, check with your network administrator.

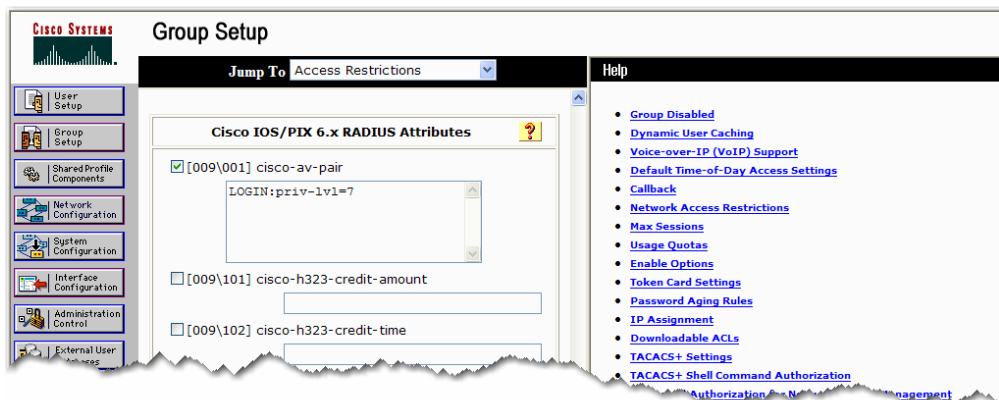


- 7 In the left column, select **Group Settings**.

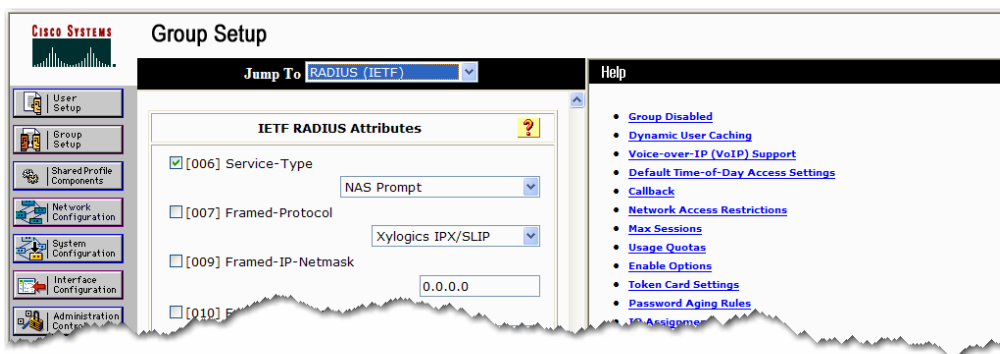


- Select your group from the **Group** list. (This assumes you've already created one, since you're already using a RADIUS server in your network.)
- Click **Edit Settings**. The **Group Settings: <Your> Group** page appears.
- Scroll down to the **Cisco IOS/PIX 6.x RADIUS Attributes** section. Select **[009\001] cisco-av-pair** and in the box below it, enter **LOGIN:priv-lvl=7**.

Level 7 and above equates to the NX's **admin** role.; Level 6 and below is **monitor**.



- d Scroll down to the **IETF RADIUS Attributes** section. Select **[006] Service-Type** and from the drop-down list, select **NAS Prompt**.



- e Click **Submit and Restart**.



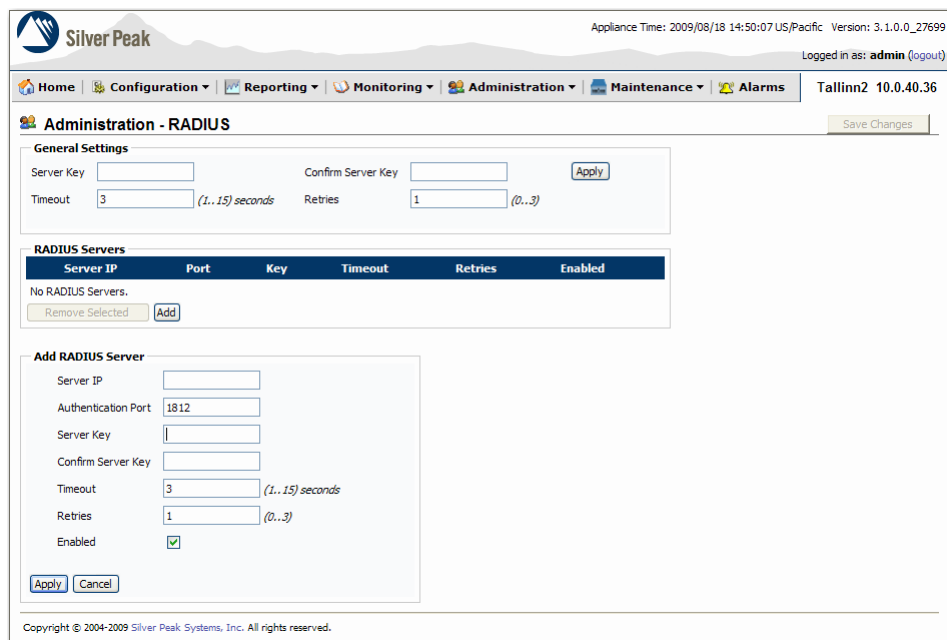
Note IMPORTANT: After making any modifications to the current ACS configuration, you must restart all ACS services for the modifications to take effect.

To do this, go to **System Configuration**, click **Service Control**, and then click **Restart**.

Configuring the NX Appliance for RADIUS

To configure RADIUS for the NX appliance, you must have **admin** privileges.

- 1 From the **Administration** menu, select **User Management > RADIUS**. The **Administration - RADIUS** page appears.
- 2 In the **RADIUS Servers** area, click **Add**.



- a In the **Server IP** field, enter the RADIUS server IP address.
 - b In the **Server Key** field, enter the RADIUS server's shared key. This is the same string entered in the Cisco Secure ACS RADIUS **Network Configuration** page for *this* NX appliance, known *there* as the **AAA client**. That key was **useAlsoInNXscreen**, so we'll enter it here.

Reenter the same key in the **Confirm Server Key** field.
 - c Accept the defaults for the **Timeout** and **Retries** fields.
 - d Make sure that the **Enabled** checkbox is selected.
- 3 Click **Apply**.
 - 4 **Save** the changes.

Now, anyone logging in to the NX appliance is subject to authentication and authorization by the RADIUS server.

Configuring for TACACS+

- 1 From the **Administration** window, select **User Management > TACACS+**. The **Administration - TACACS+** page appears.

Appliance Time: 2009/03/24 16:35:03 US/Pacific Version: 2.4.1.0_25504
Logged in as: admin (logout)

Home Configuration Reporting Monitoring Administration Maintenance Alarms tiger 10.0.41.72

Administration - TACACS+

Save Changes

General Settings

Server Key Confirm Server Key

Timeout (1..15) seconds Retries (0..3)

TACACS+ Servers

Server IP	Port	Type	Key	Timeout	Retries	Enabled
No TACACS+ Servers.						

- 2 Configure the **General Settings**, which apply to all TACACS+ servers in your list:
 - a In the **Server Key** field, enter the shared secret that was defined when setting up the TACACS+ server.
 - b In the **Confirm Server Key** field, re-enter the same text string.
 - c If you wish, change the **Timeout** value.
 - d In the **Retries** field, specify how many retries to allow the user before locking them out.
 - e Click **Apply**.
- 3 Click **Add**. The **Add TACACS+ Server** area appears.

Administration - TACACS+

Save Changes

General Settings

Server Key Confirm Server Key

Timeout (1..15) seconds Retries (0..3)

TACACS+ Servers

Server IP	Port	Type	Key	Timeout	Retries	Enabled
No TACACS+ Servers.						

Add TACACS+ Server

Server IP

Authentication Port

Authentication Type

Server Key

Confirm Server Key

Timeout (1..60) seconds

Retries (0..5)

Enabled ☒

- a In the **Server IP** field, enter the IP address for the TACACS+ server.
 - b From the Authentication Type field, select either **pap** or **ascii**, as appropriate.
 - c In the **Server Key** field, enter the password [shared secret] that was set up for the TACACS+ server.
 - d In the **Confirm Server Key** field, re-enter the same text string.
 - e In the **Timeout** field, select a timeout period between 1 and 15 seconds.
 - f In the **Retries** field, enter how many retries to allow the user before locking them out.
- 4 Click **Apply**. The server appears in the **TACACS+ Servers** table.
 - 5 Click **Save Changes**.

Refer to your TACACS+ documentation for additional details.

DEMO 14.2



Configuring a TACACS+ server for authenticating NX appliance users.

In this example, we'll configure an AAA server running TACACS+ on Cisco Secure ACS software. Then, we'll configure an NX appliance to be an AAA client to that AAA server and discuss the possible error messages.



Tip If you're adding the NX appliance to a network already being served by TACACS+, then just skip ahead to **Step 5**.

- 1 Install the Cisco Secure ACS (Access Control Server) software. Select everything and accept the defaults.
- 2 If you don't already have it, install the latest Java Runtime Environment (JRE) from java.sun.com.

Configuring the TACACS+ Server

- 3 In a web browser, enter **http://<tacacs-server-ip>:2002**. The user interface for ACS appears.
- 4 Select **Administration Control**.
 - a Create the user, **admin**, with the password, **admin**.
 - b Under **Administrator Privileges**, click **Grant All**.
 - c Click **Submit**.
- 5 Select **Interface Configuration**.
 - a Click **TACACS+ (Cisco IOS)**.
 - b Define a new service, **silverpeak**, with protocol, **ip**.

- c Enable configuration for **User**, **Group**, or both. In reality, you only **need** to configure one—either the user or the group.

Cisco Systems Interface Configuration

TACACS+ (Cisco)

TACACS+ Services

User	Group	Protocol
<input type="checkbox"/>	<input type="checkbox"/>	PPP IP
<input type="checkbox"/>	<input type="checkbox"/>	PPP IPX
<input type="checkbox"/>	<input type="checkbox"/>	PPP Multilink
<input type="checkbox"/>	<input type="checkbox"/>	PPP Apple Talk
<input type="checkbox"/>	<input type="checkbox"/>	PPP VPDN
<input type="checkbox"/>	<input type="checkbox"/>	PPP LCP
<input type="checkbox"/>	<input type="checkbox"/>	ARAP
<input type="checkbox"/>	<input type="checkbox"/>	Shell (exec)
<input type="checkbox"/>	<input type="checkbox"/>	PIX Shell (pixshell)
<input type="checkbox"/>	<input type="checkbox"/>	SLIP

New Services

Service	Protocol
<input checked="" type="checkbox"/> silverpeak	<input type="text" value="ip"/>
<input type="checkbox"/>	<input type="text"/>
<input type="checkbox"/>	<input type="text"/>

Advanced Configuration Options

☒ Advanced TACACS+ Features

☐ Display a Time-of-Day access grid for every TACACS+ service where you can override the default Time-of-Day settings

☒ Display a window for each service selected in which you can enter customized TACACS+ attributes

☒ Display enable default (Undefined) service configuration

[Back to Help](#)

Help

- [TACACS+ \(Cisco\)](#)
- [Advanced Configuration Options](#)

TACACS+ (Cisco)

Select the check box for either **User** or **Group** supported by the NAS. When you have

It is unlikely that you will use every sen a user or group cumbersome. To simpl are displayed.

This list has two sections:

- TACACS+ Services.** This section include
- New Services.** Enter the new services or Setup and/or Group Setup.

Notes: If you have configured ACS to int Center for PIX firewalls, ACS may displa function of ACS, the management appli applications, do not change or delete a

For more information about each attrib

[\[Back to Top\]](#)

Advanced Configuration Options

The Advanced Configuration Options se configurations. Select the applicable ch

- Advanced TACACS+ Features.** This opt Group Setup windows. These options ini and SENDAUTH clients, such as routers.
- Display a Time-of-Day access grid for** this option is selected, a grid appears li Group Setup window.
- Display a window for each service sele** an area appears in the User Setup and
- Display enable Default (Undefined) Ser** and Group Setup> windows that enable

Notes: This option should be used by ac

[\[Back to Top\]](#)

- d Scroll down and click **Submit**.

- 6 Select **Group Setup**. The **Group Setup** page appears.

Cisco Systems Group Setup

Select

Group : 0: Default Group (2 users) ▼

[Back to Help](#)

Help

- [Default Group](#)
- [Group](#)
- [Users in Group](#)
- [Edit Settings](#)
- [Rename Group](#)

Default Group

If group mapping has not been configured, usernames that are not configured in the ACS Internal Database are assigned to the Default Group by ACS the first time they log in. The privileges and restrictions for the default group are applied to first-time users. If you have upgraded from a previous version of ACS and kept your database information, users will map as configured in the previous version.

[\[Back to Top\]](#)

- a From the **Group** drop-down menu, select the desired group and click **Edit Settings**. The **Group Settings : <Group Name>** page appears.

- b** Scroll down the left page to **TACACS+ Settings**, and select the **silverpeak ip** checkbox.

Group Setup

Jump To: Access Restrictions

☒ No Enable Privilege
☐ Max Privilege for any AAA Client
 Level 0

IP Assignment

☐ No IP address assignment
☒ Assigned by dialup client
☐ Assigned from AAA Client pool

TACACS+ Settings

☒ **silverpeak ip**
☒ Custom attributes
 role=admin

Checking this option will PERMIT all UNKNOWN Services

☐ Default (Undefined) Services

Back to Help

Submit Submit + Restart Cancel

Help

- Group Disabled
- Dynamic User Caching
- Voice-over-IP (VoIP) Support
- Default Time-of-Day Access Settings
- Callback
- Network Access Restrictions
- Max Sessions
- Usage Quotas
- Enable Options
- Token Card Settings
- Password Aging Rules
- IP Assignment
- Downloadable ACLs
- TACACS+ Settings
- TACACS+ Shell Command Authorization
- Command Authorization for Network Device Management
- Applications
- TACACS+ Unknown Services
- IEEE RADIUS Attributes
- RADIUS Vendor-Specific Attributes

To enable administrators to tailor what authorizations are displayed for a configuration and to simplify the interface, ACS displays only the information for the current configuration. Specific Group Setup configuration options and security protocol attributes are displayed in Group Setup only in the following circumstances:

- A AAA client that uses the specified protocol has been configured in the Network Configuration section. For example, RADIUS settings appear only if you have configured a AAA client that uses RADIUS.
- The specific services, protocols, and attributes have been selected for display for the appropriate protocol in the Interface Configuration section.
- A Token Card Server has been configured in the External User Databases section.

Group Setup is used to enable and configure the particular authorizations assigned to an entire group of users. The group a user is assigned to is configured in the User Setup section. User Setup overrides Group Setup.

- c** Select the **Custom attributes** checkbox and, in the window below, enter **role=<authorization level>**.

TACACS+ Settings

☒ **silverpeak ip**
☒ Custom attributes
 role=admin

The two valid role values are as follows:

This Appliance Manager user role...	...equates to this TACACS+ Custom Attribute
admin	role=admin
monitor	role=monitor

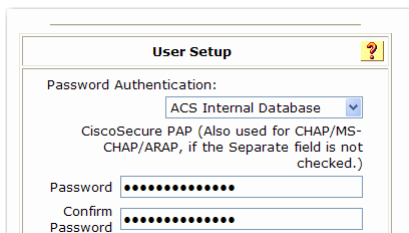


Note If you define the **Custom attributes** here, at the group level, you don't need to do it again at the user level. Thus, when you get to **Steps 7f** and **7g**, you can skip them.

- d** Click **Submit**.

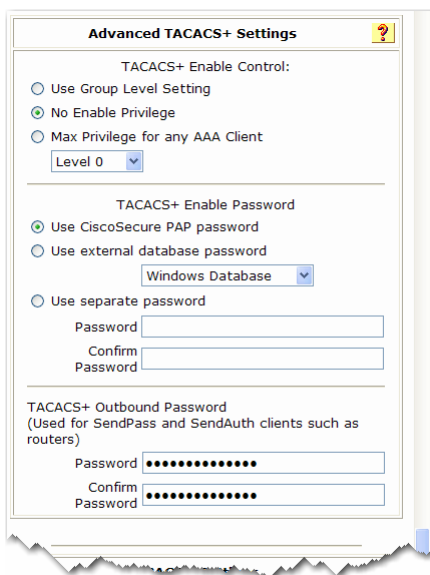
7 Select User Setup.

- a** Create a user, and set a password.
- b** In the **User Setup** section, from the **Password Authentication** field, select **ACS Internal Database**.



The **User Setup** window shows the **Password Authentication** dropdown menu set to **ACS Internal Database**. Below this, a note states: "CiscoSecure PAP (Also used for CHAP/MS-CHAP/ARAP, if the Separate field is not checked.)". There are two password fields: **Password** and **Confirm Password**, both containing masked characters (dots).

- c** Under **Advanced TACACS+ Settings**, for **TACACS+ Enable Password**, select **Use CiscoSecure PAP Password**.




The **Advanced TACACS+ Settings** window has three sections:

- TACACS+ Enable Control:** Three radio buttons are present:
☐ Use Group Level Setting
☒ No Enable Privilege
☐ Max Privilege for any AAA Client
Below the radio buttons is a **Level 0** dropdown menu.
- TACACS+ Enable Password:** Three radio buttons are present:
☒ Use CiscoSecure PAP password
☐ Use external database password (with a **Windows Database** dropdown menu)
☐ Use separate password (with **Password** and **Confirm Password** text boxes)
- TACACS+ Outbound Password:** A note says "(Used for SendPass and SendAuth clients such as routers)". It has **Password** and **Confirm Password** text boxes, both containing masked characters (dots).

Remember this. You'll need to know it later when you configure the NX appliance.

- d** Also, for the **TACACS+ Outbound Password**, enter the same password you set in Step 7a.
- e** Select the **silverpeak ip** checkbox.

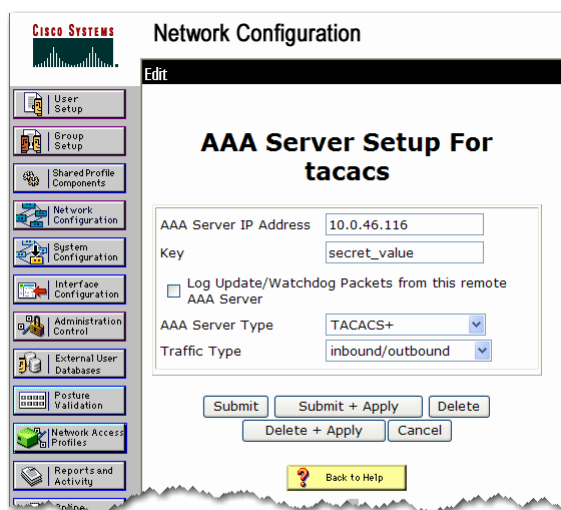
- f If you didn't already do this at the group level, then select the **Custom attributes** checkbox and, in the window below, enter **role=<authorization level>**.



The two valid role values are as follows:

This Appliance Manager user role...	...equates to this TACACS+ Custom Attribute
admin	role=admin
monitor	role=monitor

- g Click **Submit**.
- 8 Select **Network Configuration**.
- a In the **AAA Servers** section, click the AAA server name, **tacacs**. The **AAA Server Setup For tacacs** page appears.
- b If the **Key** field is blank, then enter a string of your choosing. This key, here, is for the AAA server to use in contacting the ACS server. It *is not* the key shared and used with the NX appliance(s). For our purposes here, we don't care what this password is.



- c In the **AAA Server Type** field, select **TACACS+**.
- d Click **Submit + Apply**.

- 9 Select **Network Configuration** again.
 - a In the **AAA Clients** section, click **Add Entry**. The **Add AAA Client** page appears.

The screenshot shows the Cisco Systems Network Configuration web interface. On the left is a navigation menu with icons and labels for: User Setup, Group Setup, Shared Profile Components, Network Configuration (highlighted), System Configuration, Interface Configuration, Administration Control, External User Databases, Posture Validation, Network Access Profiles, Reports and Activity, and Online Documentation. The main content area is titled 'Network Configuration' and has a sub-header 'Edit'. Below this is the 'Add AAA Client' form. The form contains the following fields and options:

- AAA Client Hostname: Text box containing 'MyNX-Appliance'
- AAA Client IP Address: Text box containing '172.10.10.223'
- Key: Text box containing 'morph'
- Authenticate Using: Dropdown menu with 'TACACS+ (Cisco IOS)' selected
- Four checkboxes:
 - ☐ Single Connect TACACS+ AAA Client (Record stop in accounting on failure).
 - ☐ Log Update/Watchdog Packets from this AAA Client
 - ☐ Log RADIUS Tunneling Packets from this AAA Client
 - ☐ Replace RADIUS Port info with Username from this AAA Client
- Buttons at the bottom: 'Submit', 'Submit + Apply', and 'Cancel'

- b In the **AAA Client Hostname** field, enter the name assigned to the NX appliance. Here, we'll use **MyNX-Appliance**.
- c In the **AAA Client IP Address** field, enter the IP address of the NX appliance. Here, we'll use **172.10.10.223**.
- d In the **Key** field, enter the shared secret that both the AAA server (here, this TACACS+ server) and the NX appliance (acting as the AAA client) will use to communicate. This entry is case sensitive. Here, we'll use **morph**.
- e From the **Authenticate Using** list, select **TACACS+ (Cisco IOS)**.
- f Click **Submit + Apply**.

Configuring the NX Appliance for TACACS+

To configure authentication and authorization for the NX appliance, you must have **admin** privileges.

- 1 From the **Administration** menu, select **User Management > TACACS+**. The **Administration - TACACS+** page appears.
- 1 In the **TACACS+ Servers** area, click **Add**.

The screenshot shows the Silver Peak Administration interface for TACACS+ configuration. The top navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The 'Administration - TACACS+' page is active, showing a 'General Settings' section with input fields for 'Server Key', 'Confirm Server Key', 'Timeout' (set to 3 seconds), and 'Retries' (set to 1). Below this is a table titled 'TACACS+ Servers' which is currently empty. At the bottom, there is an 'Add TACACS+ Server' form with fields for 'Server IP', 'Authentication Port' (set to 49), 'Authentication Type' (set to pap), 'Server Key', 'Confirm Server Key', 'Timeout' (set to 3 seconds), 'Retries' (set to 1), and an 'Enabled' checkbox which is checked. The page also includes a 'Save Changes' button and a copyright notice at the bottom.

- a In the **Server IP** field, enter the TACACS+ server IP address.
 - b Accept the default value in **Authentication Port**.
 - c In the **Authentication Type** field, select **pap**, based on what's already configured on the TACACS+ server in this example. The other supported option is **ascii**.
 - d In the **Server Key** field, enter the TACACS+ server's key. This is the same string entered in the Cisco Secure ACS TACACS+ **Network Configuration** page for *this* NX appliance, known *there* as the **AAA client**. That key was **morph**, so we'll enter it here.
Reenter the same key in the **Confirm Server Key** field.
 - e Accept the defaults for the **Timeout** and **Retries** fields.
 - f Make sure that the **Enabled** checkbox is selected.
- 2 Click **Apply**.
 - 3 **Save** the changes.

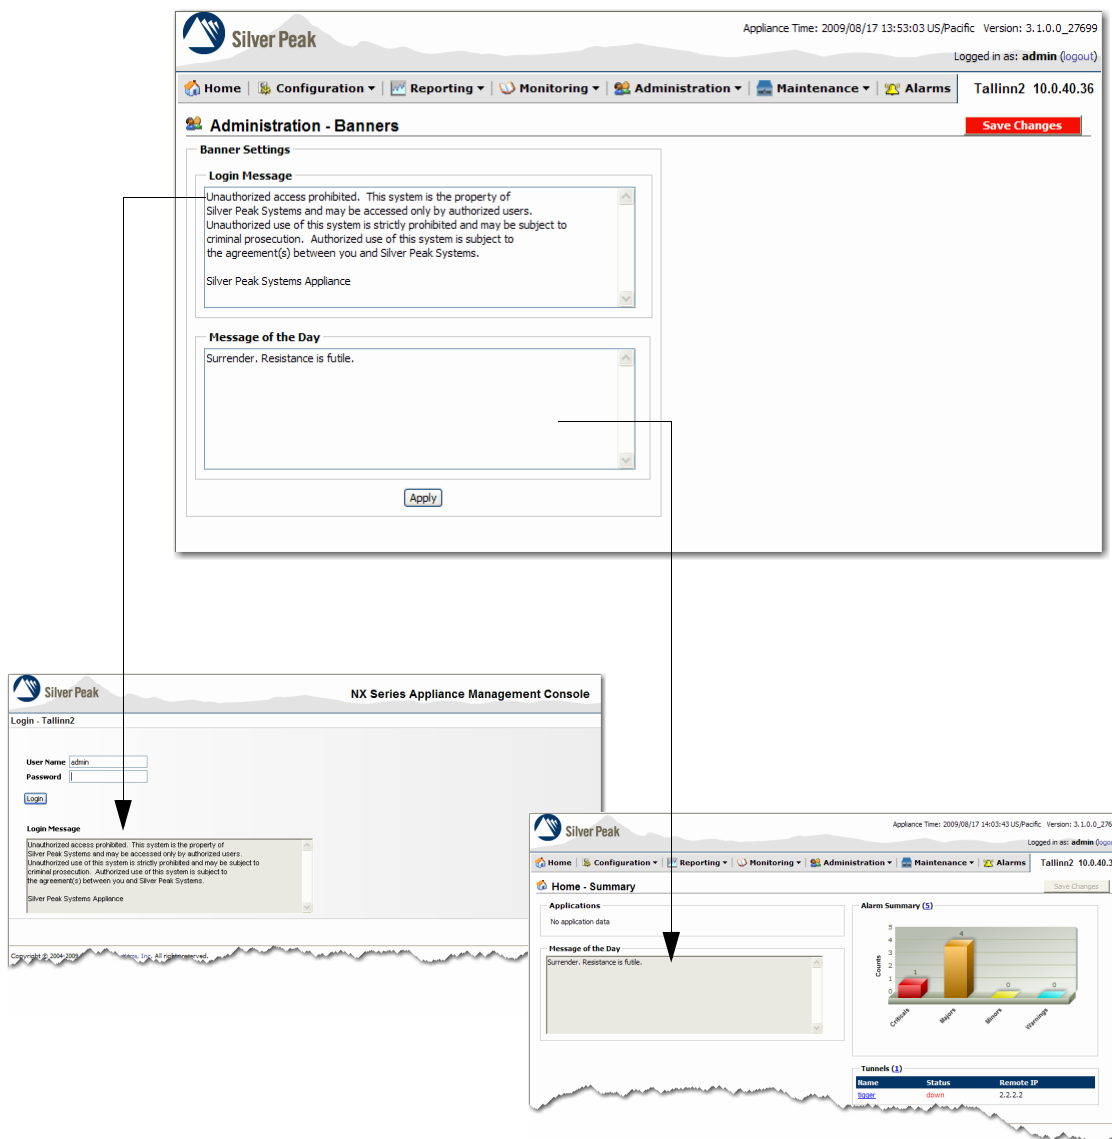
Now, anyone logging in to the Appliance Manager is subject to authentication and authorization by the TACACS+ server.

Configuring Banners

With the Appliance Manager, you can configure two different types of banners:

- The **Login Message** appears on the **Login** page.
- The **Message of the Day** appears on the **Home** page after you log in.

You can configure either, neither, or both.



Enter whatever text you want, and click **Apply**.

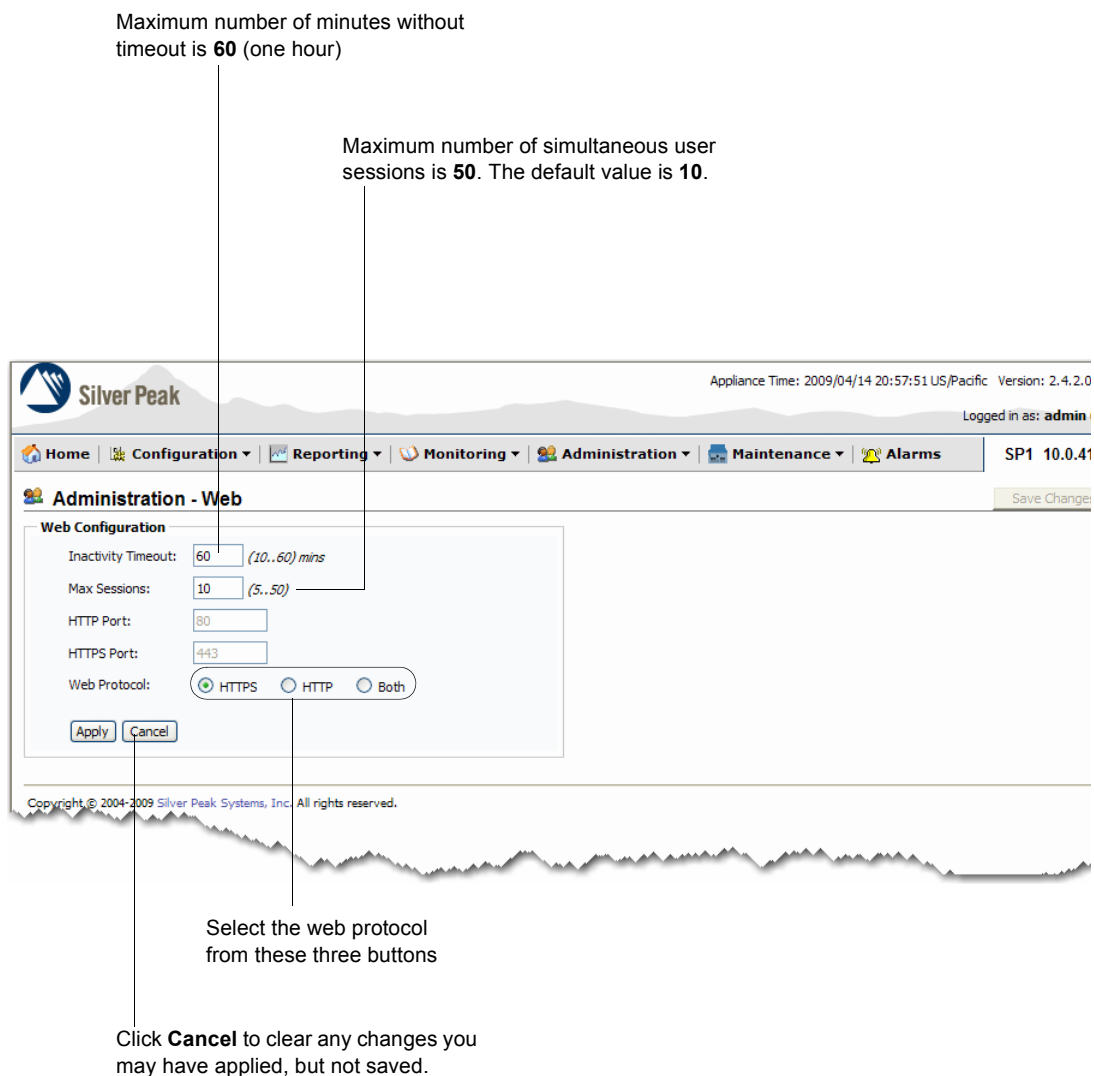
If you leave the field blank, no text displays.

Configuring Settings for Web Protocols and Web Users

This section describes how to configure the web protocol settings and web user settings.

Using the Appliance Manager, you can configure the following on the **Administration - Web** page:

- whether to enable HTTP, HTTPS, or both protocols
- how long you can go without using the Application Manager before it times out and you're forced to log on again
- the maximum number of simultaneous user sessions allowed on an appliance



Apply the changes, and then click **Save Changes** to save them to the configuration.

Initial Configuration Wizard

You can return to the **Administration - Initial Configuration Wizard** page to change appliance addressing information.

- 1 From the **Administration** Menu, select **Initial Config Wizard**.

Unlike when you first configured the appliance, the menu bar displays,

After you finish configuring, you can also see the results of this field's entry on the **Configuration - IP Route** page. To access it, go to **Configuration > Network > IP Route**.

Clicking **Skip Wizard** takes you out of the wizard and into the Home Page, without saving any configuration information. This is not recommended unless Support advises you to do so.



Note If you want to access this page in the future, just go to the **Administration** menu, and select **Initial Config Wizard**.

- 2 Based on the information you collected in the worksheet earlier in the chapter, complete the first page of this two-page wizard.

- 3 Click **Apply**. The last page of the wizard appears.

The page content depends on whether you selected **Bridge** or **Router** for **Mode**, on the first page.

How the page looks if you selected **Bridge** mode...

lan1 and wan1 are inaccessible if you have a 2-port appliance.

In Bridge mode, you can see this field's entry on the **Configuration - IP Datapath Route** page. To access it, go to **Configuration > Network > IP Datapath Route**.

The screenshot shows the 'Administration - Wizard (Page 2/2: Bridge)' interface. At the top is a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration, and Maintenance. Below the title bar is a diagram illustrating the network topology for Bridge mode: a central appliance (Silver Peak logo) is connected to a LAN (green circle) and a WAN (cloud icon). The LAN is labeled 'LAN Next-hop IP (optional)' and the WAN is labeled 'WAN Next-hop IP'. The appliance has two ports labeled 'lan0' and 'wan0', and an 'Appliance IP' is indicated. Below the diagram are configuration fields for 'Speed/Duplex' (wan0: auto/auto, lan0: auto/auto, wan1: auto/auto, lan1: auto/auto), 'Appliance IP' (Appliance IP / Netmask: 10.10.5.11 / 24, WAN Next-hop IP: 10.10.5.1, LAN Next-hop IP: (optional)), and 'Bridge Mode' (Propagate Link Down checkbox). At the bottom are buttons: < Back, Apply, Cancel, and Skip Wizard. A copyright notice at the very bottom reads: Copyright © 2004-2007 Silver Peak Systems, Inc. All rights reserved.

How the page looks if you selected **Router** mode...

The screenshot shows the 'Administration - Wizard (Page 2/2: Router)' interface. At the top is a navigation bar with links: Home, Configuration, Reporting, Monitoring, Administration, and Maintenance. Below the title bar is a diagram illustrating the network topology for Router mode: a central appliance (Silver Peak logo) is connected to a Next-hop IP (green circle) and a WAN (cloud icon). The appliance has a port labeled 'Appliance IP (wan0)'. Below the diagram are configuration fields for 'Speed/Duplex' (wan0: auto/auto), 'Appliance IP' (Appliance IP / Netmask: 10.10.5.11 / 24, Next-hop IP: 10.10.5.1, Second IP / Netmask: / 1, Second Next-hop IP: (optional)), and 'Bridge Mode' (Propagate Link Down checkbox). At the bottom are buttons: < Back, Apply, Cancel, and Skip Wizard.

- 4 Complete the fields, and click **Apply**. When the appliance asks permission to reboot, allow it.

Support

The **Administration - Support** page lists the appliance-specific information you need when calling Technical Support. It also tells you how to contact Support via web, e-mail, and phone.

The screenshot displays the Silver Peak web interface. At the top, the Silver Peak logo is on the left, and the appliance time and version are on the right. A navigation bar contains links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The 'Administration' link is selected, and the page title is 'Administration - Support'. A 'Save Changes' button is visible. The main content area is titled 'Technical Support' and contains the following information:

Model:	NX-2600 200178-001 Rev C
Serial Number:	00-e0-81-78-58-f8
Version:	3.3.0.0_32424 2010-07-12 19:21:47

Below the table, the contact information is listed:

Website: www.silver-peak.com/support

Email: support@silver-peak.com

Phone: (877) 210-7325 (408) 935-1850

At the bottom, a copyright notice reads: Copyright © 2004-2010 Silver Peak Systems, Inc. All rights reserved.



System Maintenance

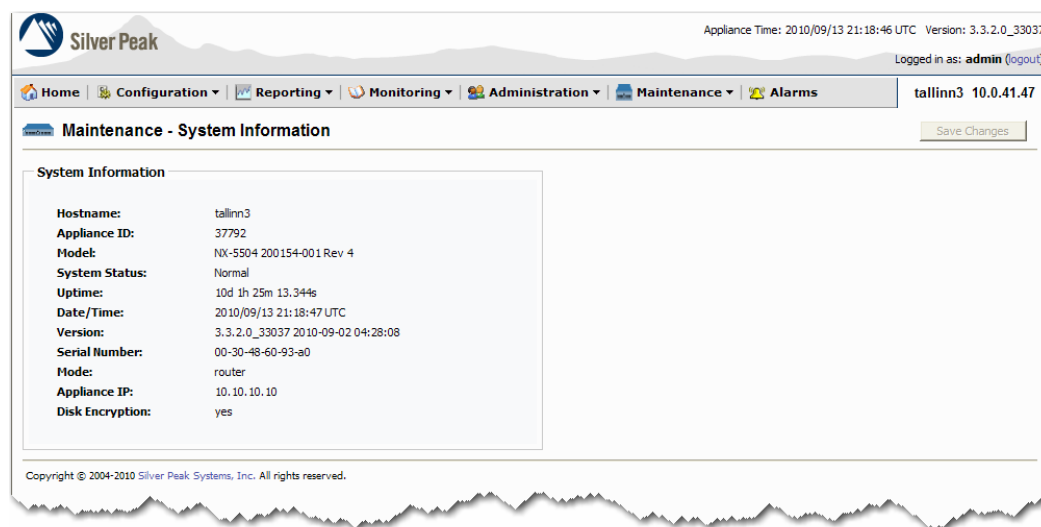
This chapter describes how to perform various system maintenance tasks.

In This Chapter

- **Viewing System Information** See page 360.
- **Upgrading the Appliance Manager Software** See page 361.
- **Managing the Appliance Configuration File** See page 372.
- **Testing Network Connectivity** See page 386.
- **Erasing Network Memory** See page 399.
- **Restarting the Appliance** See page 400.

Viewing System Information

The **Maintenance - System Information** page displays system information specific to this appliance.



The **Maintenance - System Information** page summarizes the following information:

Field	Definition/Content
Hostname	The name assigned to the appliance when using the initial configuration wizard. To edit it later, you can use the Command Line Interface (CLI) and the hostname command. The hostname is limited to a maximum of 24 characters.
Appliance ID	A network-wide unique number between 1 and 65534, assigned automatically during initial configuration.
Model	The appliance's model number. For example, NX-7600 or NX-5600.
System Status	<p>The options are Normal and Bypass.</p> <ul style="list-style-type: none"> When the status is Normal, traffic goes through tunnels, as configured. Bypass refers to hardware bypass. If there is a major problem with the appliance hardware, software, or power, all traffic goes through the appliance without any processing. Additionally, you can manually put the appliance into Bypass as an aid to troubleshooting. <p><i>To see how the GUI otherwise highlights system status, see "Banners" on page 80.</i></p>
Uptime	The time elapsed since the last reboot. For example, 3d 3h 28m 28s means "3 days, 3 hours, 28 minutes, and 28 seconds".
Date/Time	The local date and time at the appliance's location, specified by time zone.
Version	The currently running software version of the Appliance Manager.
Serial Number	The serial number of the appliance hardware.
Mode	Whether the appliance is configured for Bridge (in-line) or Router (out-of-path) mode.
Appliance IP	The IP address of this Silver Peak appliance
Disk Encryption	Yes means that Network Memory is encrypted; No means that it's not. Selecting either enforces the choice from that moment until you change it, in which case you'd have some network memory encrypted and some not. Although we don't recommend that you do this, the Appliance Manager manages both seamlessly.

Upgrading the Appliance Manager Software

This section consists of the following topics:

- **Overview** See page 361.
- **Installing a New Software Image into a Partition** See page 364.
- **Installing the Software Image from the Local Disk** See page 365.
- **Installing the Software Image from a URL** See page 366.
- **Installing the Software Image from an SCP Server** See page 367.
- **Installing the Software Image from an FTP Server** See page 369.
- **Switching to the Other Software Load** See page 371.

Overview

The Appliance Manager provides multiple options for managing appliance software. You can:

- Store two software images on the appliance
- Select which software version to run from the installed images
- Set up to switch to the other partition at the next reboot
- Install a software image from a local file, URL, Secure Copy (SCP) server, or a File Transfer Protocol (FTP) server, and install it into the appliance's inactive partition.
- Choose to reboot and begin running a newly installed software image either immediately, or at the next reboot.

When appliances within a network are operating at different software release level, the higher numbered software release determines interoperability. For more information, see *“Tunnel Compatibility Mode” on page 126*, and check the *Release Notes* to verify software version compatibility.

A software image *always* installs automatically into the inactive partition, no matter what the selection is. For all these software management tasks, use the **Maintenance - Software Upgrade** page.

Reboots from the partition that has **yes** in the **Next Boot** column

For details, see “Switching to the Other Software Load” on page 371.

When you select a file source, only the appropriate fields display.

To execute the **Install Options** choice you made with the radio buttons.

This section displays the download progress and results.

Choose what, if anything, you want the Appliance Manager to do after installing the new software image into the inactive partition.

Maintenance - Software Upgrade

Installed Images

Partition ID	Active	Next Boot	Version
1	yes	yes	2.4.3.0_25935 2009-04-15 23:07:54
2	no	no	2.4.2.0_25777 2009-04-06 23:07:54

Switch Boot Partition Reboot

Install Image

Install Options: ☒ Install ☐ Install and set next boot partition ☐ Install and reboot

[Local File] [URL] [SCP (Secure Copy)] [FTP (File Transfer Protocol)]

Remote Server Address:

Remote User Name:

Remote Password:

Remote Relative Path: Optional

Source File Name:

Install

Install Status

Status: **Ready**

Last Install Status: **The system is ready for the upgrade**

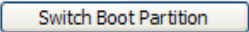
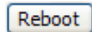
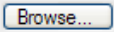

Install Start Time:

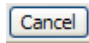
Install End Time:

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

The fields have the following definitions:

Field	Definition/Content
Partition ID	The number of the partition
Active	The partition is active if the appliance booted from that partition; in that case, the value of the field is yes . If the partition contains the inactive image, or software, then the value is no .
Next Boot	Yes specifies the partition that the appliance accesses at the next reboot.

Field	Definition/Content (Continued)
Version	<p>The version number, in the following format:</p> <p>MM.mm.Mn.PP_<internal number> <date></p> <ul style="list-style-type: none"> MM Major release version (0-99) mm Minor release version (0-99) Mn Maintenance release version (0-99) PP Patch release version (0-99) <internal number> Silver Peak's internal engineering build number
	Tells the appliance to point to the other image for the next reboot. The inactive image's Next Boot value changes from no to yes .
	Reboots the appliance now, from the partition that has yes in the Next Boot column.
Install Options:	
Install	Downloads and installs the image file into the inactive partition.
Install and set next boot partition	Downloads and installs the image file into the inactive partition. Then, designates the appliance to boot from this partition at the next reboot.
Install and reboot	Downloads the image, installs it into the inactive partition, and reboots the appliance. While the appliance reboots, it goes into the hardware bypass state, allowing all traffic to pass through the appliance without intervention. Once the reboot is complete, the appliance returns to its previous state (not in hardware bypass state). Finally, the appliance requires you to log in again.
[Local File]	When you select this item, you can Browse to the image file's location on your local hard drive.
[URL]	When you select this item, enter the image location address. The entry must begin with http://
[SCP (Secure Copy)]	When you choose to download the image from a Secure Copy server, the following fields highlight to indicate that they need entries: Remote Server Address , Remote User Name , Remote Password , Remote Full Path , and Source File Name .
FTP (File Transfer Protocol)	When you choose to download the image from a File Transfer Protocol (FTP) server, the following fields highlight to indicate that they need entries: Remote Server Address , Remote User Name , Remote Password , Remote Relative Path (optional), and Source File Name .
	Lets you find and select the image file when you've placed it on your local hard drive.
	Executes the Install Options choice you made with the radio buttons.
Remote Server Address	Use either the server IP address or the server name (if it's mapped to a local host table or a DNS server).
Remote User Name	The name of the user that the server expects
Remote Password	The password of the user that the server expects
Remote [Full/Relative] Path	<ul style="list-style-type: none"> If using the SCP server, enter the full path to the server. If using the FTP server, enter the relative path to the server.
Source File Name	The name of the software image file, ending with the suffix, .zip .
Status	If the read-only value is Ready , you may proceed with the download, installation, and reboot.
Last Install Status	The status at the end of the previous download.
Install Start Time	The time the installation began

Field	Definition/Content (Continued)
Install End Time	The time the installation was complete
	Allows you to cancel a download in progress, but not an installation once it has begun.



Note All software upgrade files end with **.zip**.

Installing a New Software Image into a Partition

When you install a new software image, the Appliance Manager automatically downloads it into the inactive partition. Depending on the option you choose, you can install the software image and:

- Store it there indefinitely
- Set it as the image to use at the next reboot
- Reboot immediately to begin running the newly installed software.

While the Silver Peak appliance reboots, it goes into the hardware bypass state, allowing all traffic to pass through the appliance without intervention. Once the reboot is complete, the appliance comes out of the hardware bypass state and requires you to log in again.

Since some traffic may drop while the appliance goes into and out of the hardware bypass state, Silver Peak suggests that you perform upgrades when traffic volume is lower, preferably after hours.



Tip Best practices recommend that before upgrading (or switching to the other partition), you preserve a copy of the running configuration file by saving it to a server. For this, use the **Maintenance - Configuration Management** page's **Save Configuration** link.



Tip Best practices also recommend scheduling a maintenance window to best accommodate the appliance reboot when installing an image.

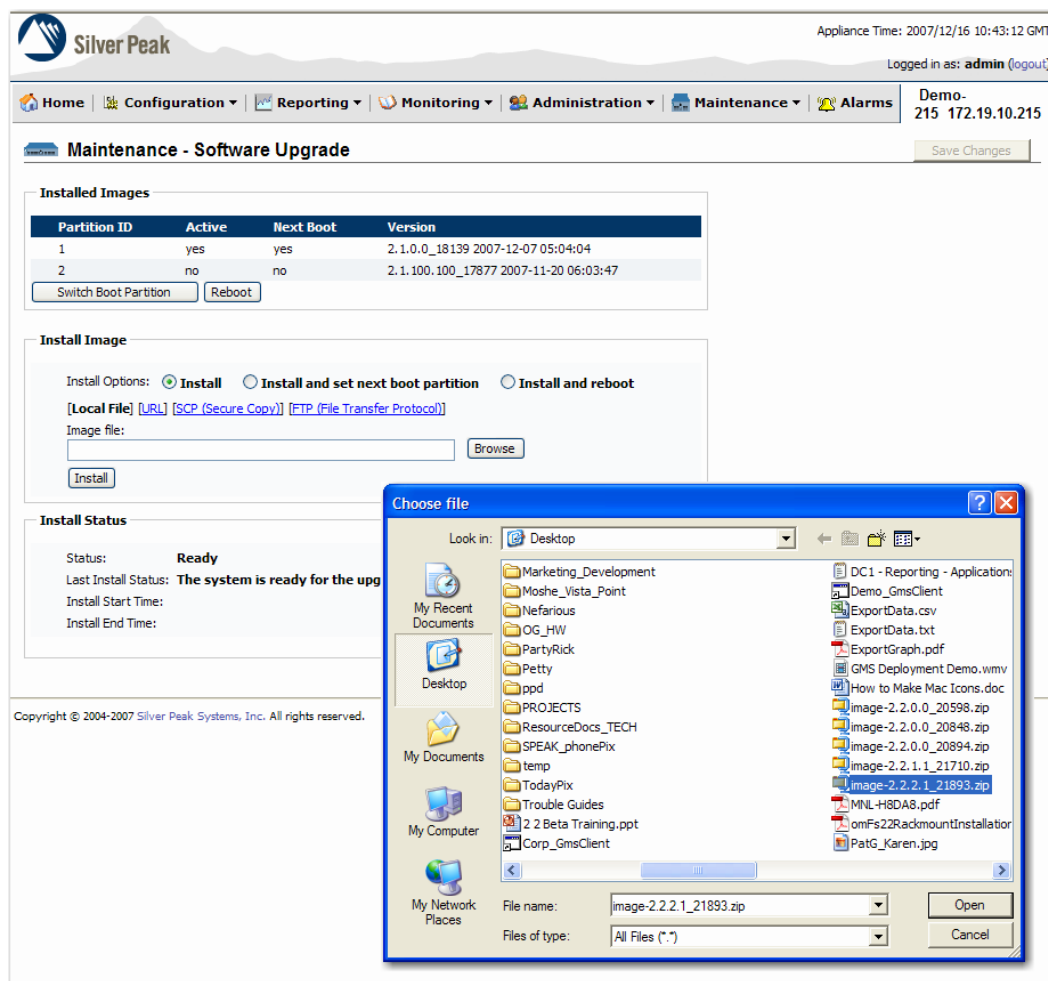


CAUTION The database schema may change with software image upgrades. Because the database is in the same partition as its associated software version, going “backwards” is not an issue. However, be aware that configuration changes made since you last ran the earlier version will be lost. To verify the feasibility, consult first with Silver Peak Customer Support.

Installing the Software Image from the Local Disk

♦ To install the software image from your computer's hard disk

- 1 On the **Maintenance - Software Upgrade** page, click **Local File**.
- 2 To use your computer's directory structure to access the file, click **Browse** and navigate to the image file.



- 3 From the **Install Options**, select one of the following:
 - **Install** – to install the image into the inactive partition
 - **Install and set next boot partition** – to install the image into the inactive partition and designate it as the partition to be used during the next boot
 - **Install and Reboot.** – to install the image into the inactive partition, switch to that partition, and then reboot to begin running it immediately. While the Silver Peak appliance reboots, it goes into the hardware bypass state, allowing all traffic to pass through the appliance without intervention. Once the reboot is complete, the appliance comes out of the hardware bypass state and requires you to log in again.
- 4 Click **Install**. The browser reports your progress during download and installation.

Installing the Software Image from a URL

♦ To install the software image from a URL

- 1 On the **Maintenance - Software Upgrade** page, click **URL**. The source-specific fields appear.
- 2 From the **Install Options**, select one of the following:
 - **Install** – to install the image into the inactive partition
 - **Install and set next boot partition** – to install the image into the inactive partition and designate it as the partition to be used during the next boot
 - **Install and Reboot** – to install the image into the inactive partition, switch to that partition, and then reboot to begin running it immediately. While the Silver Peak appliance reboots, it goes into the hardware bypass state, allowing all traffic to pass through the appliance without intervention. Once the reboot is complete, the appliance comes out of the hardware bypass state and requires you to log in again.
- 3 In the field, enter the URL (ending with the filename) after **http://** .
For example: **http://server/directory/image-2.2.1.0_21705.zip**.

Appliance Time: 2008/07/22 20:34:00 US/Pacific
Logged in as: admin (logout)

Home Configuration Reporting Monitoring Administration Maintenance Alarms Tallinn2 10.0.40.171

Maintenance - Software Upgrade

Save Changes

Installed Images

Partition ID	Active	Next Boot	Version
1	yes	yes	2.2.2.1_21893 2008-07-22 14:12:06
2	no	no	2.2.1.1_21710 2008-07-10 07:55:41

Switch Boot Partition Reboot

Install Image

Install Options: ☒ **Install** ☐ Install and set next boot partition ☐ Install and reboot

[\[Local File\]](#) [\[URL\]](#) [\[SCP \(Secure Copy\)\]](#) [\[FTP \(File Transfer Protocol\)\]](#)

Image File URL:

Install

Install Status

Status: **Ready**

Last Install Status: **The system is ready for the upgrade**

Install Start Time:

Install End Time:

Copyright © 2004-2008 Silver Peak Systems, Inc. All rights reserved.

- 4 Click **Install**. The browser reports your progress during download and installation.

Installing the Software Image from an SCP Server

♦ To install the software image from an SCP server

- 1 Access the **Maintenance - Software Upgrade** page.
- 2 From the **Install Options**, select one of the following:
 - **Install** – to install the image into the inactive partition
 - **Install and set next boot partition** – to install the image into the inactive partition and designate it as the partition to be used during the next boot
 - **Install and Reboot** – to install the image into the inactive partition, switch to that partition, and then reboot to begin running it immediately. While the Silver Peak appliance reboots, it goes into the hardware bypass state, allowing all traffic to pass through the appliance without intervention. Once the reboot is complete, the appliance comes out of the hardware bypass state and requires you to log in again.
- 3 Click **SCP (Secure Copy)**. The appropriate fields display.

Silver Peak Appliance Time: 2007/12/16 10:49:06 GMT
 Logged in as: **admin** (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms

Maintenance - Software Upgrade Save Changes

Installed Images

Partition ID	Active	Next Boot	Version
1	yes	yes	2.1.0.0_18139 2007-12-07 05:04:04
2	no	no	2.1.100.100_17877 2007-11-20 06:03:47

Switch Boot Partition Reboot

Install Image

Install Options: ☒ **Install** ☐ Install and set next boot partition ☐ Install and reboot

[\[Local File\]](#) [\[URL\]](#) [\[SCP \(Secure Copy\)\]](#) [\[FTP \(File Transfer Protocol\)\]](#)

Remote Server Address:

Remote User Name:

Remote Password:

Remote Full Path:

Source File Name:

Install

Install Status

- 4 Enter the data necessary to download the file from the SCP server.

Here, we'll use the example of downloading the file, **image-2.0.0.0_14936.zip**, from the following location:

```
scp <UserName>@10.10.10.11:/home/<UserName>/SWimages/image-2.0.0.0_14936.zip
```

Install Image

Install Options: ☒ **Install** ☐ Install and set next boot partition ☐ Install and reboot

[\[Local File\]](#) [\[URL\]](#) [\[SCP \(Secure Copy\)\]](#) [\[FTP \(File Transfer Protocol\)\]](#)

Remote Server Address:

Remote User Name:

Remote Password:

Remote Full Path:

Source File Name:

Install

Install Status

- a** For the **Remote Server Address** field, enter either:
 - the server IP address, as in **10.10.10.11**, or
 - the server name, if it's mapped to a local host table or a DNS server
 - b** Enter the **Remote User Name** and **Remote Password** for the Secure Copy (SCP) server.
 - c** For your **Remote Full Path** field, enter the *full* path, without the file name.

A full pathname includes the drive (if required), starting or root directory, and all attached subdirectories.
 - d** In the **Source File Name**, enter the image's filename. This is a **.zip** file.
- 5** Click **Install**. The browser reports your progress during download and installation.

Installing the Software Image from an FTP Server

♦ To install the software image from an FTP server

- 1 On the **Maintenance - Software Upgrade** page, click **FTP (File Transfer Protocol)**.

Appliance Time: 2007/12/16 10:52:34 GMT
Logged in as: admin (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms

Demo-215 172.19.10.215

Maintenance - Software Upgrade

Save Changes

Installed Images

Partition ID	Active	Next Boot	Version
1	yes	yes	2.1.0.0_18139 2007-12-07 05:04:04
2	no	no	2.1.100.100_17877 2007-11-20 06:03:47

Switch Boot Partition Reboot

Install Image

Install Options: ☒ Install ☐ Install and set next boot partition ☐ Install and reboot

[\[Local File\]](#) [\[URL\]](#) [\[SCP \(Secure Copy\)\]](#) [\[FTP \(File Transfer Protocol\)\]](#)

Remote Server Address:

Remote User Name:

Remote Password:

Remote Relative Path: (Optional)

Source File Name:

Install

- 2 From the **Install Options**, select one of the following:

- **Install** – to install the image into the inactive partition
- **Install and set next boot partition** – to install the image into the inactive partition and designate it as the partition to be used during the next boot
- **Install and Reboot** – to install the image into the inactive partition, switch to that partition, and then reboot to begin running it immediately. While the Silver Peak appliance reboots, it goes into the hardware bypass state, allowing all traffic to pass through the appliance without intervention. Once the reboot is complete, the appliance comes out of the hardware bypass state and requires you to log in again.

- 3 Enter the data necessary to download the file from the FTP server.

Here, we'll use the example of downloading the file, **image-2.0.0.0_14936.zip**, from Vince's directories on an FTP server:

Install Image

Install Options: ☒ Install ☐ Install and set next boot partition ☐ Install and reboot

[\[Local File\]](#) [\[URL\]](#) [\[SCP \(Secure Copy\)\]](#) [\[FTP \(File Transfer Protocol\)\]](#)

Remote Server Address: 10.10.10.11

Remote User Name: vince

Remote Password: ••••••••

Remote Relative Path: /SWimages (Optional)

Source File Name: image-2.0.0.0_14936.zip

Install

Install Status

- a** For the **Remote Server Address** field, enter either:
 - the server IP address, as in **10.10.10.11**, or
 - the server name, if it's mapped to a local host table or a DNS server
 - b** Enter the **Remote User Name** and **Remote Password** for the FTP server.
 - c** *[Optional]* For your **Remote Relative Path** field, enter the *relative* path.

A *relative* path is a path relative to the current working directory. Its first character can be anything but the pathname separator (here, a forward slash).

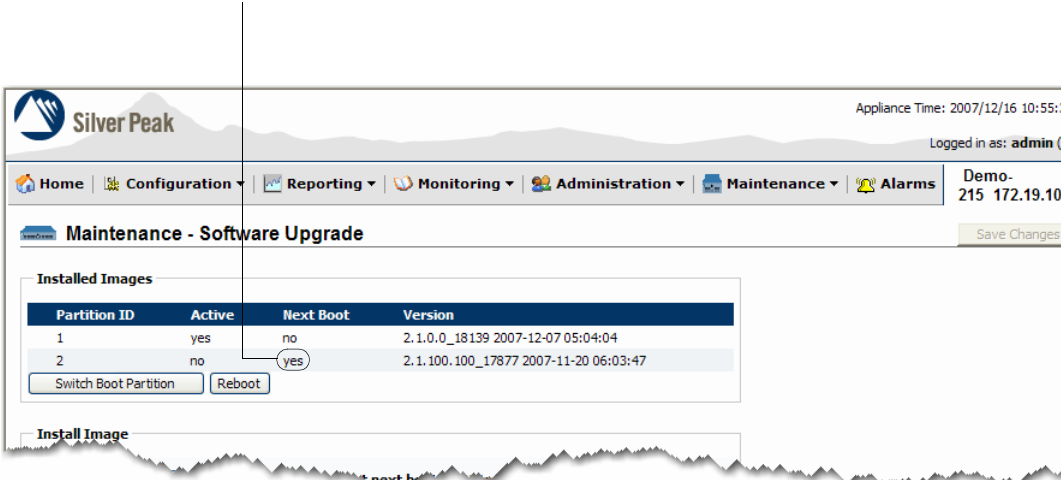
For example, if the ftp login directory is **/home/<UserName>/**, then the *relative* path would begin at the next subdirectory, as in, **/SWimages**.
 - d** In the **Source File Name**, enter the image's filename. This is a **.zip** file.
 - 4** Click **Install**. The browser reports your progress during download and installation.

Switching to the Other Software Load

You can specify that you want to switch to the other, inactive partition for the next reboot:

- To select the partition now, for a later reboot, click **Switch Boot Partition**. The inactive image's **Next Boot** value changes from **no** to **yes**.
- To select the partition now and reboot immediately, click **Reboot**.

In this example, the user clicked **Switch Boot Partition**. As a result, Partition 2, which is not currently active, has **yes** in the **Next Boot** column.



The screenshot shows the Silver Peak Maintenance - Software Upgrade interface. The 'Installed Images' table is displayed with the following data:

Partition ID	Active	Next Boot	Version
1	yes	no	2.1.0.0_18139 2007-12-07 05:04:04
2	no	yes	2.1.100.100_17877 2007-11-20 06:03:47

Below the table are buttons for 'Switch Boot Partition' and 'Reboot'. The 'Next Boot' value for Partition 2 is circled, and a line points from the text above to it.

Managing the Appliance Configuration File

This section consists of the following topics:

- **Viewing the Appliance Configuration File** See page 372.
- **Saving the Appliance Configuration File** See page 375.
- **Downloading the Appliance Configuration File** See page 380.

To protect the Appliance Manager database against loss or corruption, you can store a backup of the configuration database file, either locally on the appliance or on a local hard drive, an SCP (Secure Copy) server, or an FTP (File Transfer Protocol) server.

You can also restore or load a configuration database file from a local disk, SCP server, FTP server, or web-based location (URL).

The Appliance Manager provides multiple options for managing the active and inactive configuration files. Functionally, the Appliance Manager's **Maintenance - Configuration Management** page is divided into four distinct areas:

For removing or activating an [inactive configuration file](#) that's stored on the appliance.

To view the contents of a configuration file, click its name and a separate window opens.

For managing the [running configuration file](#). That is, an active file with unsaved changes.

For copying a [saved configuration file](#) (active or inactive) to/from the local disk or a remote location:

- **Save Configuration** copies the file from the appliance to the local disk or to a remote server.
- **Load Configuration** copies the file to the appliance from the local disk, a website's URL or a remote server.

Displays the status of any configuration file that you're saving to, or downloading from, another location.

Silver Peak

Home | Configuration | Reporting | Monitoring | Administration

Maintenance - Configuration Management

Configuration Files

Configuration	Active	Last Save Time
<input type="checkbox"/> initial	yes	200
<input type="checkbox"/> initial.bak	no	200

Remove Selected | Activate Selected

Active Configuration

Revert | Discard the running configuration and apply the contents of the active configuration file.

Save As | Save the running configuration to a new file. New filename:

[Save Configuration] | [Load Configuration]

Save the selected configuration file

[Local File] [SCP (Secure Copy)] [FTP (File Transfer Protocol)]

Save the selected configuration to the local file system

Configuration File Transfer Status

Status: **Ready**

Last Save Status: **The system is ready for upload**

Transfer Start Time:

Transfer End Time:

Copyright © 2004-2009 Silver Peak Systems, Inc. All rights reserved.

Viewing the Appliance Configuration File

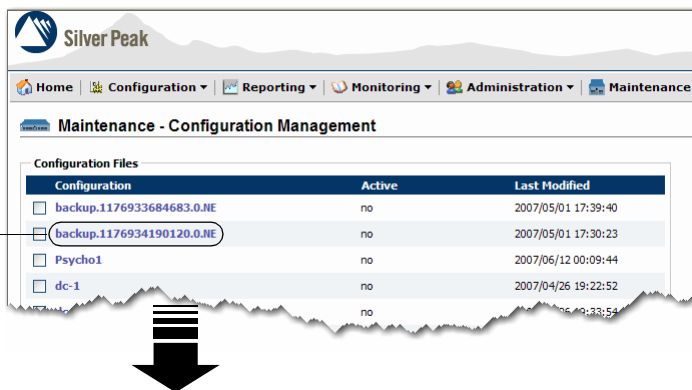
This section consists of the following topics:

- **To view the last saved version of a configuration file** See page 373.
- **To view the running configuration file** See page 374.

◆ **To view the last saved version of a configuration file**

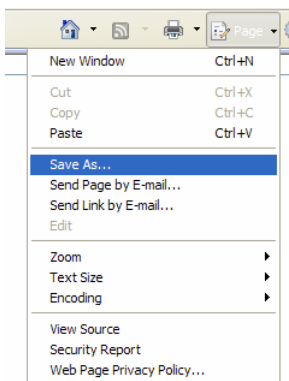
- 1 Go to the **Maintenance - Configuration Management** page.
- 2 In the **Configuration Files** table at the top of the page, click on the name of the file you want to view. A separate window opens.

To view the text files that display the Command Line Interface (CLI) commands underlying the configuration, click the file name itself. The contents open in a separate window.



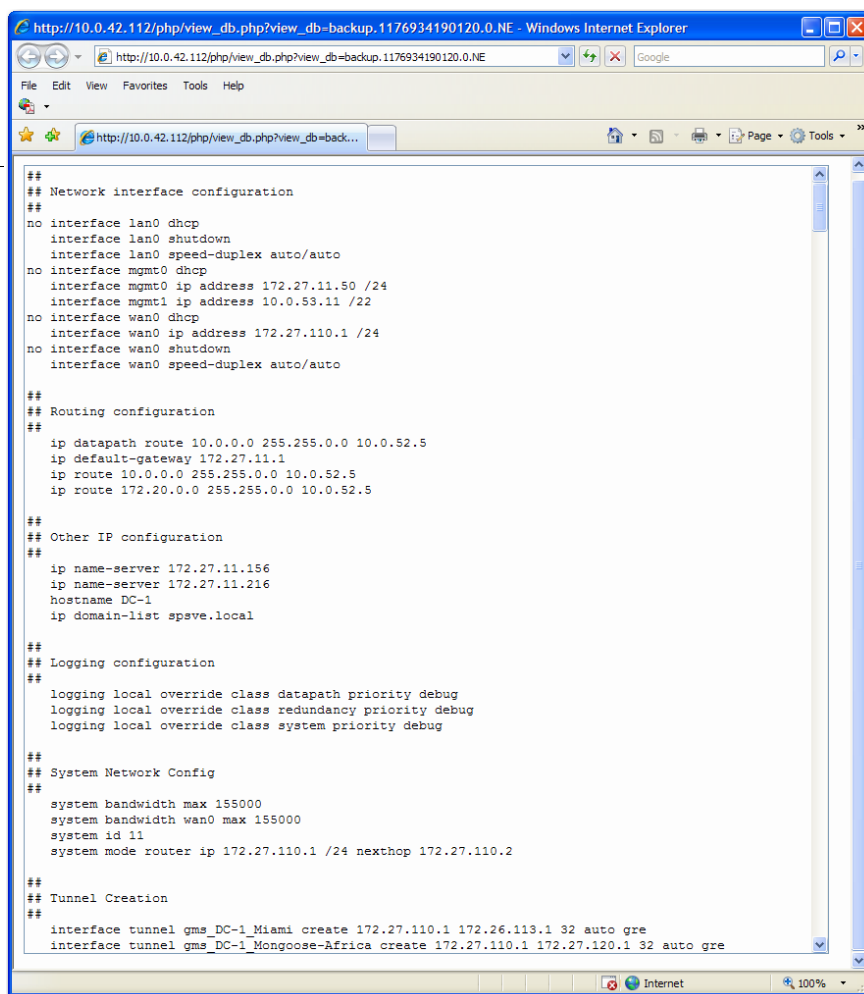
The content display **does not** change after you open the file. This view of the saved configuration is static.

Browser's menu...



IMPORTANT: You can only **save** this **text** file to your computer's **local** hard disk.

To apply this configuration to another appliance, you must first open an SSH shell to the target appliance and then copy and paste these configuration commands into the shell.



- ◆ **To view the running configuration file**

If you've made configuration changes but haven't saved them yet, then that information displays after the file name.

Indicates unsaved changes in the configuration file

☐ **initial (not saved: view running config)**

Appliance Time: 2008/07/24 14:09:20 US/Pacific
Logged in as: **admin** (logout)

Home | Configuration | Reporting | Monitoring | Administration | Maintenance | Alarms | Tallinn2 10.0.40.171

Maintenance - Configuration Management Save Changes

Configuration Files

Configuration	Active	Last Modified
<input type="checkbox"/> initial.bak.bak	no	2008/07/16 17:23:20
<input type="checkbox"/> initial (not saved: view running config)	yes	2008/07/24 14:08:45
<input type="checkbox"/> initial.bak	no	2008/07/24 14:08:34

Remove Selected Activate Selected

Active Configuration

Revert Discard the running configuration and apply the contents of the active configuration file.

Save As Save the running configuration as a new file named:

- 1 Go to the **Maintenance - Configuration Management** page.
- 2 In the **Configuration Files** table at the top of the page, click on the name of the file you want to view. A separate window opens.

Saving the Appliance Configuration File

The Appliance Manager supports saving a configuration file to three destinations external to the appliance, as follows:

- **To save the configuration file to a local disk** See page 377.
- **To save the configuration file to an SCP Server** See page 378.
- **To save the configuration file to an FTP Server** See page 379.

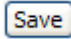

When you click to select the method, the appropriate fields appear.

Also, you can rename a file when saving it to the local disk, but you cannot rename it in the process of saving it to an SCP or FTP server.

The screenshot shows the Silver Peak Appliance Manager web interface. At the top, it displays the Silver Peak logo, the appliance time (2009/03/24 16:35:03 US/Pacific), and the version (2.4.1.0_25504). The user is logged in as 'admin'. The navigation bar includes links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The 'Maintenance - Configuration Management' section is active, showing a 'Save Changes' button. Below this, the 'Configuration Files' table lists three files: 'initial' (not saved, view running config), 'initial.bak', and 'backup.1158658595322.287.NE'. The 'Active Configuration' section has 'Revert' and 'Save As' buttons. The 'Save Configuration' dialog is open, showing a dropdown menu set to 'initial' and radio buttons for '[Local File]', '[SCP (Secure Copy)]', and '[FTP (File Transfer Protocol)]'. The 'Configuration File Transfer Status' section shows 'Status: Ready' and 'The system is ready for upload'.

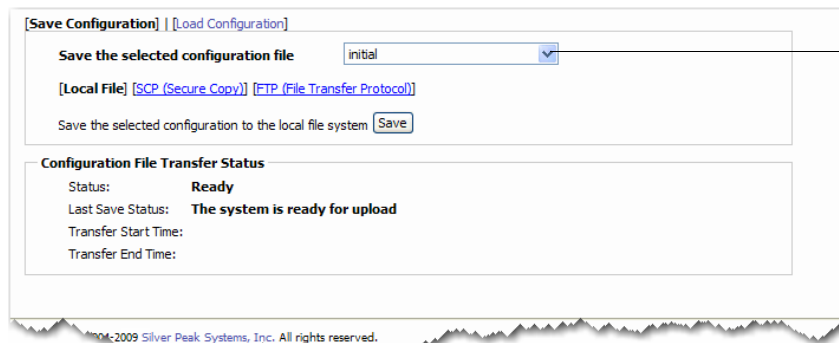
The fields and options in the **Save Configuration** area have the following definitions:

Field or Option	Definition/Content
Save the selected configuration file	This drop-down menu allows you to select any active or inactive that's saved on the appliance. These are the same files that appear in the Configuration Files table at the top of the page.
[Local File]	For saving the configuration file to your computer's local hard disk.
[SCP (Secure Copy)]	For saving the configuration file to a remote Secure Copy server.
[FTP (File Transfer Protocol)]	For saving the configuration file to a remote File Transfer Protocol (FTP) server.
Remote Server Address	Enter either the server IP address, or the server name (if it's mapped to a local host table or a DNS server).

Field or Option	Definition/Content (Continued)
Remote User Name	The name of the user that server expects
Remote Password	The password of the user that the server expects
Remote Full Path	When using the SCP server, enter the full path to the server. Initial slashes are required for full path; end slashes are not required at all.
Remote Relative Path	<i>(Optional)</i> When using the FTP server, enter the relative path to the server. Relative paths do not require initial or end slashes.
Destination File Name	<i>(Optional)</i> If you want to rename the file, do it in this field.
Status	If the read-only value is Ready , you may proceed with transferring the file to a remote server.
Last Save Status	The status at the end of the previous download.
Transfer Start Time	What time the file transfer began
Transfer End Time	What time the file transfer ended
	Saves the selected file/image to the remote server
	Allows you to cancel a file transfer that is in progress.


♦ **To save the configuration file to a local disk**

- 1 Go to the **Maintenance - Configuration Management** page, and select the file you want to save.
- 2 Click **Save Configuration**, if its fields aren't already visible.
- 3 Click **Local File**. The relevant fields appear.

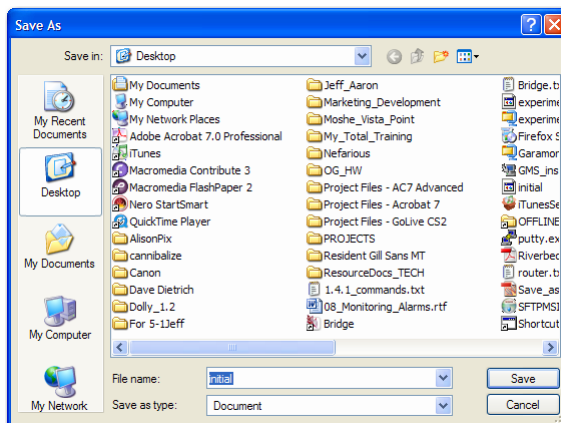


You can access any of the appliance's **saved** active or inactive configuration files from this menu.

- 4 From the menu following **Save the selected configuration file**, select the file you want to save.
- 5 Click **Save**. A dialog box appears, prompting you for confirmation.



- 6 Click **Save**. A **Save As** dialog box appears.



- 7 Browse to the target directory, rename the file if you wish, and click **Save**. A Windows-based progress bar displays as the (renamed) file saves to the computer's local disk.

If no progress bar displays, you're probably trying to save a running configuration file. Save your changes and then try again.

♦ **To save the configuration file to an SCP Server**

- 1 Go to the **Maintenance - Configuration Management** page. The **Save Configuration** area displays by default.
- 2 Select the file you want to save.
- 3 Click **SCP (Secure Copy)**. The relevant fields appear.

- 4 Enter the data necessary to save the file to the SCP server.

Here, we'll use the example of saving the file, **initial**, to the following location:

```
scp <UserName>@180.6.7.243:/home/<UserName>/work/image/initial
```

You can access any of the appliance's **saved** active or inactive configuration files from this menu.

You can begin the pathname with or without an initial slash (/).

You can save with either the existing, or a new, destination filename.

- For the **Remote Server Address** field, enter either:
 - the server IP address, as in **180.6.7.243**, or
 - the server name, if it's mapped to a local host table or a DNS server
 - Enter the **Remote User Name** and **Remote Password** for the Secure Copy (SCP) server.
 - For your **Remote Full Path** field, enter the **full** path. A full pathname includes the drive (if required), starting or root directory, and all attached subdirectories.
A full pathname requires an initial slash; no end slashes are required.
 - In the **Destination File Name**, enter the image's filename. You can use the existing file name, or save with a new file name.
- 5 Click **Save**. The Appliance Manager displays the progress.

♦ **To save the configuration file to an FTP Server**

- 1 Go to the **Maintenance - Configuration Management** page. The **Save Configuration** area displays by default.
- 2 Select the file you want to save.
- 3 Click **File Transfer Protocol (FTP) to Server**. The appropriate fields display.

[Save Configuration] | [Load Configuration]

Save the selected configuration file: initial

[Local File] [SCP (Secure Copy)] [FTP (File Transfer Protocol)]

Remote Server Address:

Remote User Name:

Remote Password:

Remote Full Path:

Destination File Name: (Optional)

Save

- 4 Enter the data necessary to save the file to the FTP server.

Here, we'll use the example of saving the file, **initial**, to Roger's directories on an FTP server:

You can access any of the appliance's **saved** active or inactive configuration files from this menu.

No slash necessary before the directory name.

[Save Configuration] | [Load Configuration]

Save the selected configuration file: initial

[Local File] [SCP (Secure Copy)] [FTP (File Transfer Protocol)]

Remote Server Address: 180.6.7.243

Remote User Name: roger

Remote Password: ••••••••

Remote Relative Path: work/image (Optional)

Destination File Name: last_1.5_backup (Optional)

Save

You can save with either the existing, or a new, destination file name. Here, we've renamed the file, **initial**, to **last_1.5_backup**.

- a For the **Remote Server Address** field, enter either:
 - the server IP address, as in **180.6.7.243**, or
 - the server name, if it's mapped to a local host table or a DNS server, as in **<myserver>**
- b Enter the **Remote User Name** and **Remote Password** for the FTP server.
- c For your **Remote Relative Path** field, enter the *relative* path and the file name.

A *relative* path is a path relative to the current working directory. Its first character can be anything but the pathname separator (here, a forward slash).

For example, if the ftp login directory is **/home/<UserName>/**, then the *relative* path would begin at the next subdirectory, as in, **work/image/**. The end slash isn't required, but is accepted.

- 5 Click **Save**. The Appliance Manager displays the progress.

Downloading the Appliance Configuration File

The Appliance Manager supports downloading a configuration file from four sources external to the appliance, as follows:

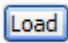
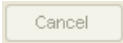
- **To load the configuration file from a local disk** See page 382.
- **To load the configuration file from a URL** See page 383.
- **To load the configuration file from an SCP Server** See page 384.
- **To load the configuration file from an FTP Server** See page 385.

When you click to select the method, the appropriate fields appear.

The screenshot shows the Silver Peak Appliance Manager web interface. The top navigation bar includes links for Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The user is logged in as 'admin'. The main content area is titled 'Maintenance - Configuration Management'. It features a 'Configuration Files' table with columns for 'Configuration' and 'Active'. The table lists three files: 'initial' (Active: yes), 'initial.bak' (Active: no), and 'backup.1158658595322.287.IIE' (Active: no). Below the table are 'Remove Selected' and 'Activate Selected' buttons. The 'Active Configuration' section includes a 'Revert' button and a 'Save As' button with a text field for a new filename. The 'Load Configuration' dialog is open, showing options to load the configuration file from a local file, URL, SCP, or FTP. The 'Local File' option is selected, and the 'Load' button is visible. The 'Configuration File Transfer Status' section shows the status as 'Ready' and the last load status as 'The system is ready for download'.

The fields and options in the **Load Configuration** area have the following definitions:

Field or Option	Definition/Content
Destination File Name	What name you want the configuration file to be assigned on the appliance.
[Local File]	For saving the configuration file to your computer's local hard disk.
[URL]	The web address to which you want to save the file. When you select this item, enter the image location address. The entry must begin with http://
[SCP (Secure Copy)]	For saving the configuration file from a remote Secure Copy server.
[FTP (File Transfer Protocol)]	For saving the configuration file from a remote File Transfer Protocol (FTP) server.
Remote Server Address	Use either the server IP address or the server name (if it's mapped to a local host table or a DNS server).

Field or Option	Definition/Content (Continued)
Remote User Name	The name of the user that server expects
Remote Password	The password of the user that the server expects
Remote [Full/Relative] Path	<p>The type of path requested depends on which method you choose:</p> <ul style="list-style-type: none">• If using the SCP server, enter the full path to the server.• If using the FTP server, enter the relative path to the server. <p>Full paths require initial slashes; end slashes are not required at all. Relative paths do not require initial or end slashes.</p>
Status	If the read-only value is Ready , you may proceed with transferring the file to a remote server.
Last Upgrade Status	The status at the end of the previous download.
Transfer Start Time	What time the file transfer began
Transfer End Time	What time the file transfer ended
	Loads the selected file/image from the remote location
	Allows you to cancel a file transfer that is in progress.


◆ **To load the configuration file from a local disk**

- 1 Go to the **Maintenance - Configuration Management** page.
- 2 Click **Load Configuration**. The **Local File** fields display by default.




If you want the file you're downloading to the appliance to have a new name, enter it here.

- 3 If you want the file that you're restoring to have a different file name on the appliance, enter a new name in the **Destination File Name** field.
- 4 Click **Browse**. A **Choose file** dialog box appears.



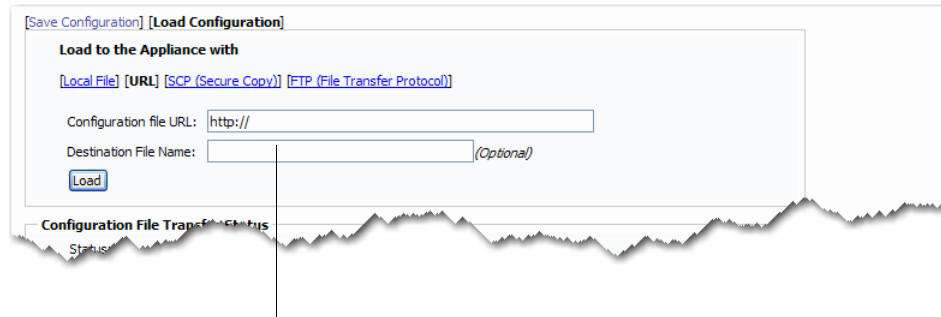
- 5 Select your file from its directory and click **Open**. The filename and path appear in the **Configuration file** field.



- 6 Click **Load**. The Appliance Manager saves the file to the appliance and lists it at the top of the page, under **Configuration Files**.

♦ **To load the configuration file from a URL**

- 1 Go to the **Maintenance - Configuration Management** page.
- 2 Click **Load Configuration**, and click **URL**.



If you want to rename the file you're downloading to the appliance, enter it here.

- 3 If you want to save the file with a different filename, enter the new name in the **Destination File Name** field.
- 4 In the field, enter the complete URL after **http://**
- 5 If you want to save the file with a new name, enter the new name in the **Destination File Name** field.
- 6 Click **Load**. The Appliance Manager reports your progress during the download and lists it at the top of the page, under **Configuration Files**.

◆ **To load the configuration file from an SCP Server**

- 1 Go to the **Maintenance - Configuration Management** page, and click **Load Configuration**.
- 2 Click **SCP (Secure Copy)**. The appropriate fields appear.

[Save Configuration] [Load Configuration]

Load to the Appliance with

[Local File] [URL] [SCP (Secure Copy)] [FTP (File Transfer Protocol)]

Remote Server Address:

Remote User Name:

Remote Password:

Remote Full Path:

Source File Name:

Destination File Name: (Optional)

Configuration File Transfer Status

Ready

- 3 Enter the data necessary to save the file from the SCP server to the appliance.
- Here, we'll use the example of renaming and restoring the file, **testfile**, from the following location:
- ```
scp <UserName>@180.6.7.243:/home/<UserName>/work/configfiles/testfile
```

[Save Configuration] [Load Configuration]

**Load to the Appliance with**

[Local File] [URL] [SCP (Secure Copy)] [FTP (File Transfer Protocol)]

Remote Server Address:

Remote User Name:

Remote Password:

Remote Full Path:

Source File Name:

Destination File Name:  (Optional)

**Configuration File Transfer Status**

Status: Ready

If you want to rename the file during the download, enter the new file name here.

- a For the **Remote Server Address** field, enter either:
    - the server IP address, as in **180.6.7.243**, or
    - the server name, if it's mapped to a local host table or a DNS server
  - b Enter the **Remote User Name** and **Remote Password** for the Secure Copy (SCP) server.
  - c For your **Remote Full Path** field, enter the full path and the file name.  
 A full pathname includes the drive (if required), starting or root directory, all attached subdirectories and ends with the file or object name.  
 A full pathname requires an initial slash; no end slashes are required.
  - d In the **Destination File Name** field, enter the existing file name or rename the file.
- 4 Click **Load**. The Appliance Manager displays the progress.

♦ **To load the configuration file from an FTP Server**

- 1 Go to the **Maintenance - Configuration Management** page, and click **Load Configuration**.
- 2 Click **FTP (File Transfer Protocol)**. The appropriate fields appear.

- 3 Enter the data necessary to save the file from the FTP server.

Here, we'll use the example of loading the file, **testfile**, from Roger's directory on an FTP server. In the process, we'll rename it to **newfilename**:

If you want to rename the file during the download, enter the new file name here.

- a For the **Remote Server Address** field, enter either:
    - the server IP address, as in **180.6.7.243**, or
    - the server name, if it's mapped to a local host table or a DNS server, as in **<myserver>**
  - b Enter the **Remote User Name** and **Remote Password** for the FTP server.
  - c For your **Remote Relative Path** field, enter the *relative* path.  
 A *relative* path is a path relative to the current working directory. Its first character can be anything but the pathname separator (here, a forward slash).  
 For example, if the ftp login directory is **/home/<UserName>/**, then the *relative* path would begin at the next subdirectory, as in, **work/configfiles**. The end slash isn't required, but is accepted.
  - d In the **Destination File Name** field, enter the existing file name or rename the file.
- 4 Click **Load**. The Appliance Manager reports your progress during the download and lists it at the top of the page, under **Configuration Files**.

## Testing Network Connectivity

The Appliance Manager enables you to test network connectivity, using three commands: **ping**, **traceroute**, and **tcpdump**.

- There can only be one connectivity test session per appliance at any time, regardless of which command you're using.
- Click **Stop** to terminate a test.
- If a user logs on to an appliance while a testing session is in progress, only the **Abort** button is accessible to them. Otherwise, that button is not visible.

### ♦ To run a Network Connectivity test

- 1 From the **Maintenance** menu, select **Network Connectivity**. The **Maintenance - Network Connectivity** page appears.

- The **ping** and **traceroute** tests provide an **IP/Hostname** field.
- The **tcpdump** test displays a **File Name** field instead, and automatically enters a name in the format, **tcpdump\_<hostname>**. After running a tcpdump test, you can locate the captured results on the **Administration - Debug Files** page, via the **TCP Dump Result** link. You can download the resulting file to your PC for viewing and analyzing via Wireshark® or Ethereal®.

When a test begins, **Start** changes to **Stop**. Use as needed.

After a test has begun, this area displays its status and start/end times.

To view the arguments that the selected command can take

**Abort** allows a user with admin privileges to terminate another user's connectivity test session. It's available whenever there's a session in progress.

The screenshot shows the Silver Peak Appliance Manager interface. The top navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The 'Maintenance - Network Connectivity' page is active. It features a 'Settings' section with a 'Type' dropdown set to 'ping', an 'IP / Hostname' field with '10.0.41.119', and an 'Option' field. A 'Start' button and a 'Help' link are present. To the right, the 'Status of Execution' box shows 'Status: Ready', 'Start Time: 2009/04/20 14:10:11', and 'End Time: 2009/04/20 14:10:40'. Below this is a 'Network Connectivity Result' section with a graph showing network activity over time. A 'Save Changes' button is in the top right corner.

- 2 Complete the fields as follows:
  - a From the **Type** field, click to select the test you want.
  - b If the **IP / Hostname** field is present, enter the IP address or hostname of the destination device. If the **File Name** field displays, its field is populated by a default name.
  - c In the **Option** field, enter the command option you want. For example, for **ping**, you could enter **-c 3** to stop after sending three ECHO\_REQUEST packets. For available arguments, click **Help**.  
Options for each command are listed on pages 389 thru 396.
- 3 Click **Start**. The **Network Connectivity Result** area displays intermediate results every few seconds. To stop the test and see the complete results, click **Stop**. For example:

### ping

The screenshot shows the Silver Peak web interface for Network Connectivity testing. The top navigation bar includes Home, Configuration, Reporting, Monitoring, Administration, Maintenance, and Alarms. The main section is titled "Maintenance - Network Connectivity".

**Network Connectivity Settings:**

- Type: ping (selected from a dropdown)
- IP / Hostname: 10.0.41.119
- Option: (empty text field)
- Buttons: Start, Help

**Status of Execution:**

- Status: Ready
- Start Time: 2009/04/20 15:02:37
- End Time: 2009/04/20 15:02:50
- Buttons: Abort

**Network Connectivity Result:**

```

PING 10.0.41.119 (10.0.41.119) 56(84) bytes of data.
64 bytes from 10.0.41.119: icmp_seq=1 ttl=63 time=0.153 ms
64 bytes from 10.0.41.119: icmp_seq=2 ttl=63 time=0.107 ms
64 bytes from 10.0.41.119: icmp_seq=3 ttl=63 time=0.099 ms
64 bytes from 10.0.41.119: icmp_seq=4 ttl=63 time=0.092 ms
64 bytes from 10.0.41.119: icmp_seq=5 ttl=63 time=0.092 ms
64 bytes from 10.0.41.119: icmp_seq=6 ttl=63 time=0.112 ms
64 bytes from 10.0.41.119: icmp_seq=7 ttl=63 time=0.121 ms
64 bytes from 10.0.41.119: icmp_seq=8 ttl=63 time=0.114 ms
64 bytes from 10.0.41.119: icmp_seq=9 ttl=63 time=0.092 ms
64 bytes from 10.0.41.119: icmp_seq=10 ttl=63 time=0.095 ms
64 bytes from 10.0.41.119: icmp_seq=11 ttl=63 time=0.094 ms
64 bytes from 10.0.41.119: icmp_seq=12 ttl=63 time=0.118 ms
64 bytes from 10.0.41.119: icmp_seq=13 ttl=63 time=0.127 ms

--- 10.0.41.119 ping statistics ---
13 packets transmitted, 13 received, 0% packet loss, time 11997ms
rtt min/avg/max/mdev = 0.092/0.108/0.153/0.022 ms

```

## traceroute

The screenshot shows the Silver Peak web interface with the 'Maintenance - Network Connectivity' section. The 'Network Connectivity' settings are configured with Type: traceroute, IP / Hostname: 10.0.41.119, and Option: (empty). The 'Status of Execution' shows the test is Ready, with Start Time: 2009/04/20 15:06:29 and End Time: 2009/04/20 15:06:35. The 'Network Connectivity Result' section displays the following output:

```

traceroute to 10.0.41.119 (10.0.41.119), 30 hops max, 40 byte packets
1 10.0.40.33 (10.0.40.33) 0.189 ms 0.130 ms 1.117 ms
2 ny-sp1-s.speak.local (10.0.41.119) -0.903 ms 0.082 ms 0.068 ms

```

## tcpdump

The **Option** field populates with **-n** by default. This flag results in **tcpdump** not trying to resolve IP addresses, so that the process doesn't try to perform a DNS lookup on every new IP address it encounters.

The screenshot shows the Silver Peak web interface with the 'Maintenance - Network Connectivity' section. The 'Network Connectivity' settings are configured with Type: tcpdump, File Name: tcpdump\_Tallinn2, and Option: -n. The 'Status of Execution' shows the test is Ready, with Start Time: 2009/04/20 15:12:38 and End Time: 2009/04/20 15:13:03. The 'Network Connectivity Result' section displays the following output:

```

File name: tcpdump_Tallinn2, size: 0.00 KB
File name: tcpdump_Tallinn2, size: 0.08 KB
File name: tcpdump_Tallinn2, size: 0.08 KB
File name: tcpdump_Tallinn2, size: 0.08 KB
File name: tcpdump_Tallinn2, size: 0.08 KB
File name: tcpdump_Tallinn2, size: 0.08 KB
File name: tcpdump_Tallinn2, size: 0.15 KB


```

Access this file on the **Administration - Debug Files** page, under the **TCP Dump Result** link.

## 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 datapath interfaces, use the **-I** option with the local appliance IP address. For example:



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

The following **ping** options are supported:

| Option    | Explanation                                                                                                                                                                                                                                                                                             |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-A</b> | 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>-b</b> | Allow pinging a broadcast address.                                                                                                                                                                                                                                                                      |
| <b>-B</b> | Do not allow ping to change source address of probes. The address is bound to one selected when ping starts.                                                                                                                                                                                            |
| <b>-c</b> | <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.                                                                                                                                                |
| <b>-d</b> | Set the SO_DEBUG option on the socket being used. Essentially, this socket option is not used by Linux kernel.                                                                                                                                                                                          |
| <b>-F</b> | <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.                                                                                                                                                     |
| <b>-i</b> | <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.                                                                     |
| <b>-I</b> | <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.                                                                                                       |
| <b>-l</b> | <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.                                                                                                                                                       |
| <b>-L</b> | Suppress loopback of multicast packets. This flag only applies if the ping destination is a multicast address.                                                                                                                                                                                          |
| <b>-M</b> | <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).                                                                     |
| <b>-n</b> | Numeric output only. No attempt will be made to lookup symbolic names for host addresses.                                                                                                                                                                                                               |

| Option        | Explanation (Continued)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-p</b>     | <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, <i>-p ff</i> will cause the sent packet to be filled with all ones.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>-Q</b>     | <p><i>tos</i>: Set Quality of Service -related bits in ICMP datagrams. <i>tos</i> can be either decimal or hex number.</p> <p>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.</p> <p>Possible settings for Type of Service are: minimal cost: 0x02, reliability: 0x04, throughput: 0x08, low delay: 0x10.</p> <p>Multiple TOS bits should not be set simultaneously.</p> <p>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.</p> <p>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).</p> |
| <b>-q</b>     | Quiet output. Nothing is displayed except the summary lines at startup time and when finished.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>-R</b>     | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>-r</b>     | 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 <b>-I</b> is also used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>-s</b>     | <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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>-S</b>     | <i>sndbuf</i> : Set socket sndbuf. If not specified, it is selected to buffer not more than one packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-t ttl</b> | Set the IP Time to Live.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>-T</b>     | <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).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>-U</b>     | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>-v</b>     | Verbose output.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-V</b>     | Show version and exit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>-w</b>     | <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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

When you click the **Help** button, the following displays in the **Network Connectivity Result** area.

```
Usage: ping [-LRUbdnqrvVaA] [-c count] [-i interval] [-w deadline]
 [-p pattern] [-s packetsize] [-t ttl] [-I interface or address]
 [-M mtu discovery hint] [-S sndbuf]
 [-T timestamp option] [-Q tos] [hop1 ...] destination
```



## Using traceroute

Use the **traceroute** command to trace the route that packets take to a destination.

The following **traceroute** options are supported:

| Option    | Explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-d</b> | Enable socket level debugging.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>-f</b> | Set the initial time-to-live used in the first outgoing probe packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>-F</b> | Set the “don’t fragment” bit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>-g</b> | Specify a loose source route gateway (8 maximum).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>-i</b> | 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 <b>-s</b> flag for another way to do this.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-I</b> | Use ICMP ECHO instead of UDP datagrams.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>-m</b> | 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).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>-n</b> | Print hop addresses numerically rather than symbolically and numerically (saves a nameserver address-to-name lookup for each gateway found on the path).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-p</b> | 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.                                                                                                                                                                                                                                                                                             |
| <b>-q</b> | nqueries                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-r</b> | 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)).                                                                                                                                                                                                                                                                                                                                                 |
| <b>-s</b> | 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 <b>-i</b> flag for another way to do this.)                                                                                                                                                           |
| <b>-t</b> | 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. |
| <b>-v</b> | Verbose output. Received ICMP packets other than TIME_EXCEEDED and UNREACHABLEs are listed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>-w</b> | Set the time (in seconds) to wait for a response to a probe (default 5 sec.).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Option | Explanation (Continued)                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -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).                                                                                                                                                                                                                                             |

When you click the **Help** button, the following displays in the **Network Connectivity Result** area.

```
Version 1.4a12
Usage: traceroute [-dFIrvx] [-g gateway] [-i iface] [-f first_ttl]
[-m max_ttl] [-p port] [-q nqueries] [-s src_addr] [-t tos]
[-w waittime] [-z pausesecs] host [packetlen]
```

## Using tcpdump

Use the **tcpdump** command to display packets on a network.

For example, to capture 100 packets on the wan0 interface, use the command, **-n -i wan0 -c 100**.

The following **tcpdump** options are supported:

| Option      | Explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-A</b>   | Print each packet (minus its link level header) in ASCII. Handy for capturing web pages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>-c</b>   | Exit after receiving count packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>-C</b>   | Before writing a raw packet to a savefile, check whether the file is currently larger than file_size and, if so, close the current savefile and open a new one. Savefiles after the first savefile will have the name specified with the <b>-w</b> flag, with a number after it, starting at 1 and continuing upward. The units of file_size are millions of bytes (1,000,000 bytes, not 1,048,576 bytes).                                                                                                                                                                                                                                                                                                                                                                      |
| <b>-d</b>   | Dump the compiled packet-matching code in a human readable form to standard output and stop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-dd</b>  | Dump packet-matching code as a C program fragment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>-ddd</b> | Dump packet-matching code as decimal numbers (preceded with a count).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>-D</b>   | Print the list of the network interfaces available on the system and on which tcpdump can capture packets. For each network interface, a number and an interface name, possibly followed by a text description of the interface, is printed. The interface name or the number can be supplied to the <b>-i</b> flag to specify an interface on which to capture.<br><br>This can be useful on systems that don't have a command to list them (e.g., Windows systems, or UNIX systems lacking ifconfig -a); the number can be useful on Windows 2000 and later systems, where the interface name is a somewhat complex string.<br><br>The <b>-D</b> flag will not be supported if tcpdump was built with an older version of libpcap that lacks the pcap_findalldevs() function. |
| <b>-e</b>   | Print the link-level header on each dump line.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Option     | Explanation (Continued)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-E</b>  | <p>Use <code>spi@ipaddr algo:secret</code> for decrypting IPsec ESP packets that are addressed to <code>addr</code> and contain Security Parameter Index value <code>spi</code>. This combination may be repeated with comma or newline separation.</p> <p>Note that setting the secret for IPv4 ESP packets is supported at this time.</p> <p>Algorithms may be <code>des-cbc</code>, <code>3des-cbc</code>, <code>blowfish-cbc</code>, <code>rc3-cbc</code>, <code>cast128-cbc</code>, or <code>none</code>. The default is <code>des-cbc</code>. The ability to decrypt packets is only present if <code>tcpdump</code> was compiled with cryptography enabled.</p> <p><code>secret</code> is the ASCII text for ESP secret key. If preceded by <code>0x</code>, then a hex value will be read.</p> <p>The option assumes RFC2406 ESP, not RFC1827 ESP. The option is only for debugging purposes, and the use of this option with a true 'secret' key is discouraged. By presenting IPsec secret key onto command line you make it visible to others, via <code>ps(1)</code> and other occasions.</p> <p>In addition to the above syntax, the syntax <code>file name</code> may be used to have <code>tcpdump</code> read the provided file in. The file is opened upon receiving the first ESP packet, so any special permissions that <code>tcpdump</code> may have been given should already have been given up.</p> |
| <b>-f</b>  | <p>Print 'foreign' IPv4 addresses numerically rather than symbolically (this option is intended to get around serious brain damage in Sun's NIS server â usually it hangs forever translating non-local internet numbers).</p> <p>The test for 'foreign' IPv4 addresses is done using the IPv4 address and netmask of the interface on which capture is being done. If that address or netmask are not available either because the interface on which capture is being done has no address or netmask or because the capture is being done on the Linux "any" interface, which can capture on more than one interface, this option will not work correctly.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-F</b>  | Use <code>file</code> as input for the filter expression. An additional expression given on the command line is ignored.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-i</b>  | <p>Listen on interface. If unspecified, <code>tcpdump</code> searches the system interface list for the lowest numbered, configured up interface (excluding loopback). Ties are broken by choosing the earliest match.</p> <p>On Linux systems with 2.2 or later kernels, an interface argument of "any" can be used to capture packets from all interfaces. Note that captures on the "any" device will not be done in promiscuous mode.</p> <p>If the <b>-D</b> flag is supported, an interface number as printed by that flag can be used as the interface argument.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>-l</b>  | <p>Make stdout line buffered. Useful if you want to see the data while capturing it. For example,</p> <pre>tcpdump -l   tee dat <u>or</u> tcpdump -l &gt; dat &amp; tail -f dat</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>-L</b>  | List the known data link types for the interface and exit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>-m</b>  | Load SMI MIB module definitions from file module. This option can be used several times to load several MIB modules into <code>tcp-dump</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-M</b>  | Use <code>secret</code> as a shared secret for validating the digests found in TCP segments with the TCP-MD5 option (RFC 2385), if present.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>-n</b>  | Don't convert host addresses to names. This can be used to avoid DNS lookups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>-nn</b> | Don't convert protocol and port numbers etc. to names either.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>-N</b>  | Don't print domain name qualification of host names. For example, if you give this flag then <code>tcpdump</code> will print <code>nic</code> instead of <code>nic.ddn.mil</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Option       | Explanation (Continued)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-O</b>    | Do not run the packet-matching code optimizer. This is useful only if you suspect a bug in the optimizer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>-p</b>    | Don't put the interface into promiscuous mode. Note that the interface might be in promiscuous mode for some other reason; hence, <b>-p</b> cannot be used as an abbreviation for 'ether host {local-hw-addr} or ether broadcast'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-q</b>    | Quick output. Print less protocol information so output lines are shorter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-R</b>    | Assume ESP/AH packets to be based on old specification (RFC1825 to RFC1829). If specified, tcpdump will not print replay prevention field. Since there is no protocol version field in ESP/AH specification, tcpdump cannot deduce the version of ESP/AH protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>-r</b>    | Read packets from file (which was created with the <b>-w</b> option). Standard input is used if file is '-'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>-S</b>    | Print absolute, rather than relative, TCP sequence numbers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>-s</b>    | Snarf snaplen bytes of data from each packet rather than the default of 68 (with SunOS's NIT, the minimum is actually 96). 68 bytes is adequate for IP, ICMP, TCP and UDP but may truncate protocol information from name server and NFS packets. Packets truncated because of a limited snapshot are indicated in the output with <b>[ proto]</b> , where <b>proto</b> is the name of the protocol level at which the truncation has occurred.                                                                                                                                                                                                                                                               |
|              | Note that taking larger snapshots both increases the amount of time it takes to process packets and, effectively, decreases the amount of packet buffering. This may cause packets to be lost. You should limit snaplen to the smallest number that will capture the protocol information you're interested in. Setting snaplen to 0 means use the required length to catch whole packets.                                                                                                                                                                                                                                                                                                                    |
| <b>-t</b>    | Don't print a timestamp on each dump line.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>-tt</b>   | Print an unformatted timestamp on each dump line.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>-ttt</b>  | Print a delta (in micro-seconds) between current and previous line on each dump line.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>-tttt</b> | Print a timestamp in default format proceeded by date on each dump line.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>-T</b>    | Force packets selected by "expression" to be interpreted the specified type. Currently known types are: <div> <div><b>aodv</b></div> <div>Ad-hoc On-demand Distance Vector protocol</div> <div><b>cnfp</b></div> <div>Cisco NetFlow protocol</div> <div><b>rpc</b></div> <div>Remote Procedure Call</div> <div><b>rtp</b></div> <div>Real-Time Applications protocol</div> <div><b>rtcp</b></div> <div>Real-Time Applications control protocol</div> <div><b>snmp</b></div> <div>Simple Network Management Protocol</div> <div><b>tftp</b></div> <div>Trivial File Transfer Protocol</div> <div><b>vat</b></div> <div>Visual Audio Tool)</div> <div><b>wb</b></div> <div>distributed White Board</div> </div> |
| <b>-u</b>    | Print undecoded NFS handles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>-U</b>    | Make output saved via the <b>-w</b> option "packet-buffered"; that is, as each packet is saved, it will be written to the output file, rather than being written only when the output buffer fills.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|              | The <b>-U</b> flag will not be supported if tcpdump was built with an older version of libpcap that lacks the pcap_dump_flush() function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>-v</b>    | When parsing and printing, produce (slightly more) verbose output. For example, the time to live, identification, total length and options in an IP packet are printed. Also enables additional packet integrity checks such as verifying the IP and ICMP header checksum.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|              | When writing to a file with the <b>-w</b> option, report, every 10 seconds, the number of packets captured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

| Option                                                        | Explanation (Continued)                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-vv</b>                                                    | Even more verbose output. For example, additional fields are printed from NFS reply packets, and SMB packets are fully decoded.                                                                                                                                                                                                            |
| <b>-vvv</b>                                                   | Even more verbose output. For example, telnet SB... SE options are printed in full. With <b>-X</b> Telnet options are printed in hex as well.                                                                                                                                                                                              |
| <b>-w</b>                                                     | Write the raw packets to file rather than parsing and printing them out. They can later be printed with the <b>-r</b> option. Standard output is used if file is "-".                                                                                                                                                                      |
| <b>-W</b>                                                     | Used in conjunction with the <b>-C</b> option, this will limit the number of files created to the specified number, and begin overwriting files from the beginning, thus creating a 'rotating' buffer. In addition, it will name the files with enough leading 0s to support the maximum number of files, allowing them to sort correctly. |
| <b>-x</b>                                                     | Print each packet (minus its link level header) in hex. The smaller of the entire packet or snaplen bytes will be printed. Note that this is the entire link-layer packet, so for link layers that pad (e.g. Ethernet), the padding bytes will also be printed when the higher layer packet is shorter than the required padding.          |
| <b>-xx</b>                                                    | Print each packet, including its link level header, in hex.                                                                                                                                                                                                                                                                                |
| <b>-X</b>                                                     | Print each packet (minus its link level header) in hex and ASCII. This is very handy for analyzing new protocols.                                                                                                                                                                                                                          |
| <b>-XX</b>                                                    | Print each packet, including its link level header, in hex and ASCII.                                                                                                                                                                                                                                                                      |
| <b>-y</b>                                                     | Set the data link type to use while capturing packets to datalinktype.                                                                                                                                                                                                                                                                     |
| <b>-Z</b>                                                     | Drops privileges (if root) and changes user ID to user and the group ID to the primary group of user.                                                                                                                                                                                                                                      |
| This behavior can also be enabled by default at compile time. |                                                                                                                                                                                                                                                                                                                                            |

When you click the **Help** button, the following displays in the **Network Connectivity Result** area.

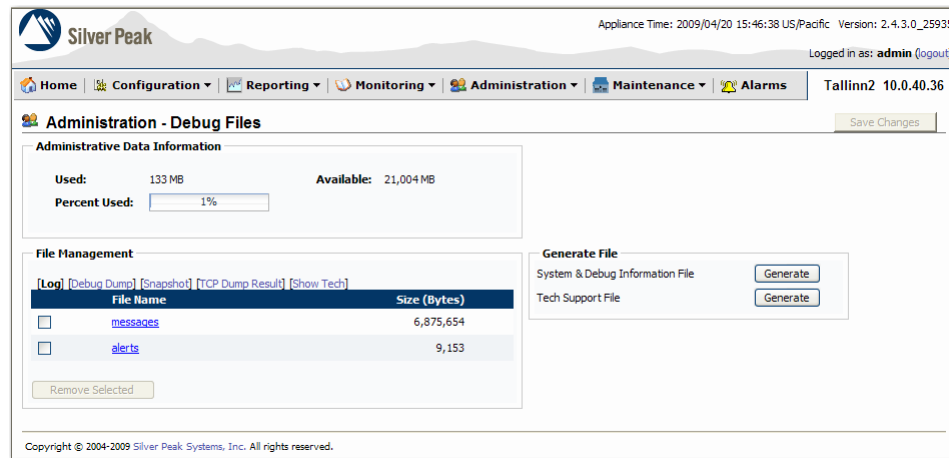
```

tcpdump version 3.8
libpcap version 0.8.3
Usage: tcpdump [-aAdDeflLnNOpqRStuUvxX] [-c count] [-C file_size]
 [-E algo:secret] [-F file] [-i interface] [-M secret]
 [-r file] [-s snaplen] [-T type] [-w file] [-W filecount]
 [-y datalinktype] [-Z user]
 [expression]

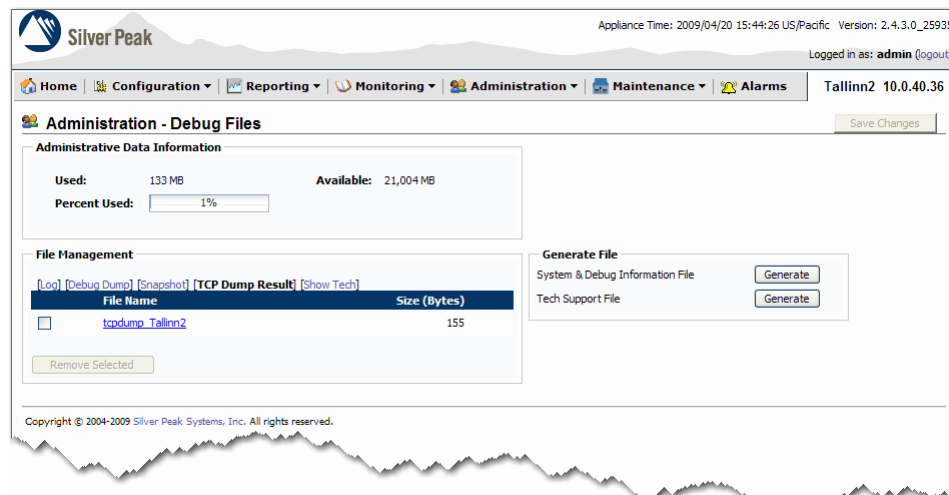
```

- ◆ **To retrieve tcpdump results**

- 1 From the **Administration** menu, select **Debug Files**. The **Administration - Debug Files** page appears.



- 2 In the **File Management** area, click **TCP Dump Result**. Any saved tcpdump files display.



- 3 To access the tcpdump file, click on its name link. The **Administration - Debug Files - Save File** page appears.



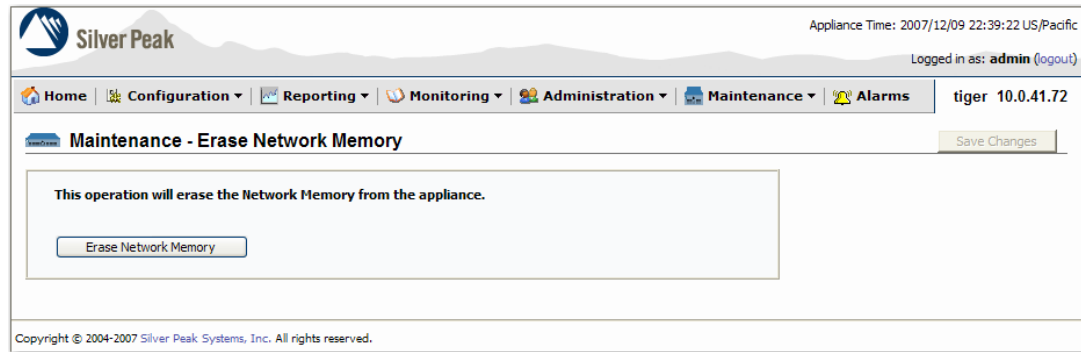
- 4 Select whether you want to save the file to your PC, an SCP server, or an FTP server, and click **Save**.
- 5 Complete the fields for the method you've chosen.



## Erasing Network Memory

The **Maintenance - Erase Network Memory** page is useful in lab and evaluation environments, when you need first-pass numbers to establish a baseline before Network Memory is applied.

You can use this page to clear Network Memory, without having to reboot.



For this to succeed, erase Network Memory from the appliances at each end of the tunnel.

## Restarting the Appliance

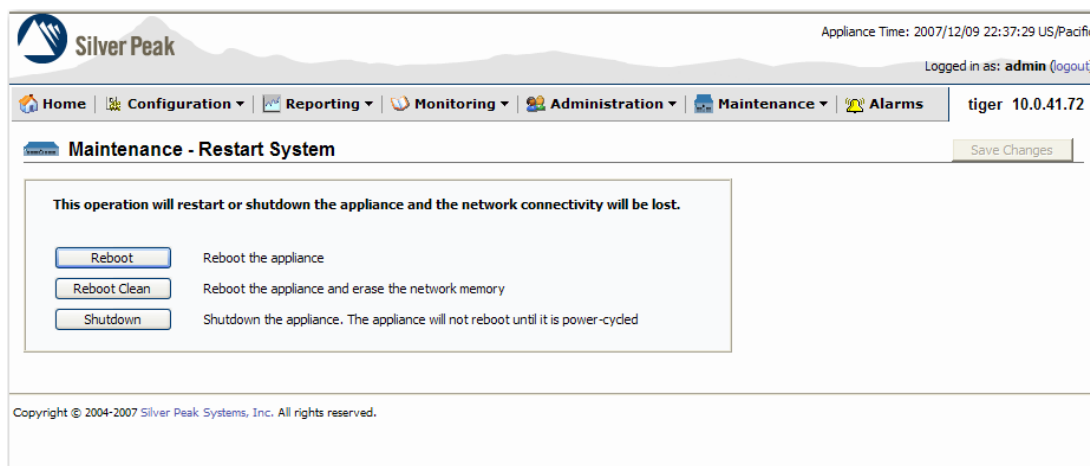
This section describes the types of reboots available for restarting the appliance, the possible reasons for choosing a particular method, and the consequences of each.

The Silver Peak appliance supports four different types of reboot, appropriate to your various circumstances and needs at the time.

Regardless of the reboot type, while it's rebooting, the appliance goes into:

- the hardware bypass state if the appliance is deployed in-line (Bridge mode), or
- the open-port state if the appliance is deployed out-of-path (Router mode).

This allows traffic to pass, but without the benefits of compression, acceleration, or Network Memory™.



| Restart Type        | What it does...                                                                                                                                          | You might need to use it if...                                                                                                                                                                                                                  |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Reboot</b>       | Reboots the appliance gracefully. This is your typical, "vanilla" restart.                                                                               | <ul style="list-style-type: none"> <li>• You're changing the deployment mode and other configuration parameters that require a reboot.</li> </ul>                                                                                               |
| <b>Reboot Clean</b> | Reboots the appliance and cleans out the Network Memory™.                                                                                                | <ul style="list-style-type: none"> <li>• You need to restart the appliance with clean Network Memory™ data.</li> </ul>                                                                                                                          |
| <b>Shutdown</b>     | Shuts down the appliance and turns the power off. To restart, you'll need to go to the appliance and physically turn the power on with the power switch. | <ul style="list-style-type: none"> <li>• You're decommissioning the appliance.</li> <li>• You need to physically move the appliance to another location.</li> <li>• You need to recable the appliance for another type of deployment</li> </ul> |

### ♦ To restart the appliance

- 1 From **Maintenance** menu, select **Restart System**.
- 2 Click the type of reboot you want. The appliance asks you to confirm your decision.
- 3 Click **Yes**. The appliance reboots.



# 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 402.
- **Types of Alarms** See page 403.
- **Handling Current Alarms** See page 408.

## Understanding Alarms

This section defines the four alarm severity categories and provides tables of all Silver Peak appliance alarms. It also describes the difference between viewing current alarms and reviewing a history of alarms.

The **Alarms - Current Alarms** page lists currently existing alarm conditions for 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 itself.

### 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.

An alarm, whether acknowledged or unacknowledged, is only accessible via the **Alarms - Current Alarms** page until it clears.

## Types of Alarms

The appliance can raise alarms based on problems with tunnels, traffic class, and equipment. 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 16-1 NX Series Alarms

| Subsystem | Alarm Type ID (Hex) | Alarm Severity | Alarm Text                                                                                           |
|-----------|---------------------|----------------|------------------------------------------------------------------------------------------------------|
| Tunnel    | 00010003            | CRITICAL       | Tunnel keepalive version mismatch                                                                    |
|           | 00010001            | CRITICAL       | Tunnel state is Down                                                                                 |
|           | 00010000            | MAJOR          | Tunnel Remote ID is Misconfigured                                                                    |
|           | 00010005            | MINOR          | Tunnel software version mismatch                                                                     |
|           | 00010002            | MINOR          | Tunnel Vrrp IP Misconfigured                                                                         |
|           | 00010004            | MINOR          | Tunnel wccp gid is misconfigured                                                                     |
| Software  | 00040003            | CRITICAL       | The licensing for this virtual appliance has expired. [VX-5000 only] <sup>a</sup>                    |
|           | 00040004            | CRITICAL       | There is no license installed on this virtual appliance. [VX-5000 only]                              |
|           | 00040002            | MAJOR          | Significant change in time of day has occurred, and might compromise statistics. Please contact TAC. |
|           | 00040001            | MAJOR          | System is Low on resources                                                                           |
| Equipment | 00030007            | CRITICAL       | Encryption card H/W Failure                                                                          |
|           | 00030005            | CRITICAL       | LAN/WAN Fail-to-Wire Card Failure                                                                    |
|           | 00030006            | CRITICAL       | LAN/WAN Fail-to-Wire Card Relay Failure                                                              |
|           | 00030021            | CRITICAL       | NIC interface failure                                                                                |
|           | 00030000            | CRITICAL       | RAID Array is Degraded                                                                               |
|           | 00030004            | CRITICAL       | System is in BYPASS mode                                                                             |
|           | 0003001d            | MAJOR          | Bonding members have different speed/duplex                                                          |
|           | 0003001c            | MAJOR          | Cluster peer is down                                                                                 |
|           | 00030010            | MAJOR          | Datapath Internal loopback Test Failed                                                               |
|           | 00030001            | MAJOR          | Disk is failed                                                                                       |
|           | 00030015            | MAJOR          | Disk is not in service                                                                               |
|           | 0003000b            | MAJOR          | Interface is Half Duplex                                                                             |
|           | 0003000c            | MAJOR          | Interface Speed is 10 Mbps                                                                           |
|           | 00030022            | MAJOR          | LAN Next-Hop unreachable <sup>b</sup>                                                                |
|           | 0003001a            | MAJOR          | LAN/WAN interface has been shut down due to link propagation of paired interface                     |
|           | 00030018            | MAJOR          | LAN/WAN interfaces have different admin states                                                       |
|           | 00030019            | MAJOR          | LAN/WAN interfaces have different link carrier states                                                |
|           | 0003000a            | MAJOR          | Management Interface Link Down                                                                       |

Table 16-1 NX Series Alarms (continued)

| Subsystem                        | Alarm Type ID (Hex) | Alarm Severity | Alarm Text                                                                                                             |
|----------------------------------|---------------------|----------------|------------------------------------------------------------------------------------------------------------------------|
|                                  | 00030009            | MAJOR          | Network Interface Link Down                                                                                            |
|                                  | 00030020            | MAJOR          | Power supply not connected, not powered or failed                                                                      |
|                                  | 00030012            | MAJOR          | VRRP instance is down                                                                                                  |
|                                  | 00030014            | MAJOR          | WAN next hop router is learned on a LAN port (box is in backwards)                                                     |
|                                  | 00030011            | MAJOR          | WAN Next-Hop unreachable <sup>a</sup>                                                                                  |
|                                  | 0003001e            | MAJOR          | WCCP adjacency(ies) down                                                                                               |
|                                  | 0003001f            | MAJOR          | WCCP assignment table mismatch                                                                                         |
|                                  | 00030002            | MINOR          | Disk is Degraded                                                                                                       |
|                                  | 00030016            | MINOR          | Disk is rebuilding                                                                                                     |
|                                  | 00030017            | MINOR          | Disk has been removed by operator                                                                                      |
|                                  | 0003001b            | MINOR          | Disk SMART threshold exceeded                                                                                          |
|                                  | 00030008            | WARNING        | Network Interface Admin Down                                                                                           |
|                                  | 00050001            | WARNING        | The average WAN-side transmit throughput of X Mbps over the last minute [exceeded, fell below] the threshold of Y Mbps |
| Threshold Crossing Alerts (TCAs) | 00050002            | WARNING        | The average LAN-side receive throughput of X Mbps over the last minute [exceeded, fell below] the threshold of Y Mbps  |
|                                  | 00050003            | WARNING        | The total number of X optimized flows at the end of the last minute [exceeded, fell below] the threshold of Y          |
|                                  | 00050004            | WARNING        | The total number of X flows at the end of the last minute [exceeded, fell below] the threshold of Y                    |
|                                  | 00050005            | WARNING        | The file system utilization of X% at the end of the last minute [exceeded, fell below] the threshold of Y              |
|                                  | 00030013            | WARNING        | VRRP State Master --> Backup                                                                                           |
|                                  | 00050006            | WARNING        | The peak latency of X during the last minute [exceeded, fell below] the threshold of Y                                 |
|                                  | 00050007            | WARNING        | The average pre-FEC loss of X% over the last minute [exceeded, fell below] the threshold of Y%                         |
|                                  | 00050008            | WARNING        | The average post-FEC loss of X% over the last minute [exceeded, fell below] the threshold of Y%                        |
|                                  | 00050009            | WARNING        | The average pre-POC out-of-order packets of X% over the last minute [exceeded, fell below] the threshold of Y%         |
|                                  | 0005000a            | WARNING        | The average post-POC out-of-order packets of X% over the last minute [exceeded, fell below] the threshold of Y%        |
|                                  | 0005000b            | WARNING        | The average tunnel utilization of X% over the last minute [exceeded, fell below] the threshold of Y%                   |
|                                  | 0005000c            | WARNING        | The average tunnel reduction of X% over the last minute [exceeded, fell below] the threshold of Y%                     |

a. The VX-5000 is a virtual appliance, comprised of the VX-5000 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 after upgrading to Version 2.4.3.1, resolve the alarm(s) immediately by configuring the gateway(s) to respond to ICMP pings from the Silver Peak NX Appliance IP Address.

## Viewing Current Alarms

You can view current alarms as follows:

To view the **Alarms - Current Alarms** page, either:

- Click the hyperlinked number after **Alarm Summary**, or
- Click **Alarms**.

This appliance is in **System Bypass**.

Appliance Time: 2007/12/16 14:52:31 US/Pacific  
Logged in as: admin (logout)

Tallinn 10.0.55.90 (BYPASS)

Home - Summary

Applications

Top 5 Applications (LAN to WAN - Last 24 Hours)

Bandwidth (Bytes)

Reduction

unassigned, leap, topnet-2153, cifs\_smb, ms\_rpc

Alarm Summary (9)

Counts

Criticals: 2, Majors: 7, Minors: 0, Warnings: 0

Tunnels (1)

| Name                | Status | Remote IP   |
|---------------------|--------|-------------|
| Tallinn_to_Helsinki | down   | 172.56.6.11 |

Copyright © 2004-2007 Silver Peak Systems, Inc. All rights reserved.

To disable System Bypass:

- Go to the **Configuration - System** page
- Click to deselect **System Bypass**
- Click **Apply**.

☒ System Bypass

☐ System Bypass

...and the **Alarms - Current Alarms** page displays.

Appliance Time: 2007/12/16 14:52:31 US/Pacific  
Logged in as: admin (logout)

Tallinn 10.0.40.109 (BYPASS)

Alarms - Current Alarms

Alarm Summary (9)

Criticals: 2, Majors: 7, Minors: 0, Warnings: 0

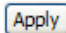
Alarms

| Seq No. | Date/Time           | Type | Severity | Source              | Description                               | Clear                    | Ack                      |
|---------|---------------------|------|----------|---------------------|-------------------------------------------|--------------------------|--------------------------|
| 15      | 2007/12/16 14:52:21 | EQU  | Critical | System              | System BYPASS mode                        | <input type="checkbox"/> | <input type="checkbox"/> |
| 14      | 2007/12/16 11:06:57 | EQU  | Major    | system              | Datapath Gateway Connectivity Test Failed | <input type="checkbox"/> | <input type="checkbox"/> |
| 11      | 2007/12/13 09:24:07 | TUN  | Critical | Tallinn_to_Helsinki | Tunnel state is Down                      | <input type="checkbox"/> | <input type="checkbox"/> |
| 9       | 2007/12/13 09:23:57 | EQU  | Major    | 10.10.10.4          | Cluster peer is down                      | <input type="checkbox"/> | <input type="checkbox"/> |
| 8       | 2007/12/13 09:23:57 | EQU  | Major    | 10.10.10.3          | Cluster peer is down                      | <input type="checkbox"/> | <input type="checkbox"/> |
| 7       | 2007/12/13 09:23:57 | EQU  | Major    | 10.10.10.2          | Cluster peer is down                      | <input type="checkbox"/> | <input type="checkbox"/> |
| 6       | 2007/12/13 09:23:52 | EQU  | Major    | 1                   | Disk has been removed by operator         | <input type="checkbox"/> | <input type="checkbox"/> |
| 5       | 2007/12/13 09:23:51 | EQU  | Major    | 1                   | Disk is not-in-service                    | <input type="checkbox"/> | <input type="checkbox"/> |
| 3       | 2007/12/13 09:23:29 | EQU  | Major    | wan0                | Network Interface Link Down               | <input type="checkbox"/> | <input type="checkbox"/> |

Apply Clear All Ack All Unack All



The **Alarm - Current Alarms** page displays the following information:

| Field                                                                             | Definition/Content                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Applies any changes or edits to the running configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Seq No.</b>                                                                    | The sequential number of the alarm, based on the time the alarm raised.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Date/Time</b>                                                                  | The local date and time at the appliance's location, specified by a 24-hour clock.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Type</b>                                                                       | <p>The type of alarm:</p> <ul style="list-style-type: none"> <li>• <b>Tunnel</b> A tunnel-based alarm</li> <li>• <b>TC</b> A traffic class-based alarm</li> <li>• <b>EQU</b> An equipment-based alarm</li> <li>• <b>SW</b> A code- or software-based alarm</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Severity</b>                                                                   | <p>The severity of the alarm, listed here in decreasing order of severity:</p> <ul style="list-style-type: none"> <li>• <b>Critical</b> A critical alarm, such as "Tunnel Down"</li> <li>• <b>Major</b> A major alarm, such as "Disk out of Service"</li> <li>• <b>Minor</b> A minor alarm, such as "Disk Degraded"</li> <li>• <b>Warning</b> A warning, such as "Software Process Restart"</li> <li>• <b>Info</b> For Silver Peak debugging purposes.</li> </ul> <p>These are purely related to alarms severities, <b>not</b> event logging levels, even though some of the naming conventions overlap. Events and alarms have different sources. Alarms, once they clear, list as the ALERT level in the <b>Alarms - Log Viewer</b> page.</p> |
| <b>Source</b>                                                                     | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                | A brief description of the alarm.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Clear</b>                                                                      | <p>Check boxes with a black border are user-clearable.</p> <p>To clear the alarm, click the <b>Clear</b> box, and click <b>Apply</b>. Once cleared, the row is removed and the content is viewable in the read-only page, <b>Alarms - Log Viewer</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Ack</b>                                                                        | Select <b>Yes</b> to acknowledge the alarm; Select <b>No</b> to remove acknowledgement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Handling Current Alarms

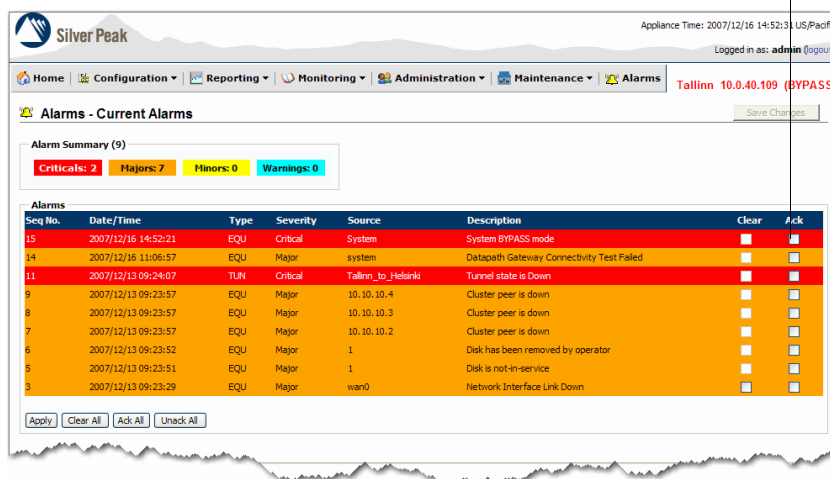
On the **Alarms - Current Alarms** page, you can:

- Acknowledge alarms
- Clear alarms, if they are clearable

### Acknowledging Alarms

Acknowledge or unacknowledge alarms as follows:

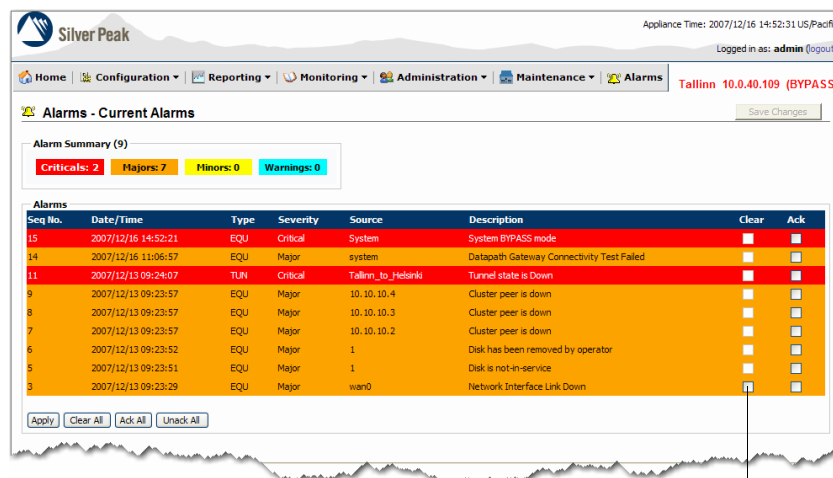
- To acknowledge an individual alarm, select the **Ack** column. Click **Apply**.
- To unacknowledge an alarm, deselect the **Ack** column. Click **Apply**.



### Clearing 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.

Clear alarms as follows:



You can only access the check boxes in the **Clear** column that have a black outline. To clear an alarm, click its box. Then click **Apply**.

# Hardware Maintenance

This chapter describes how to replace a hard disk and how to replace a power supply for those appliances for which the customer is authorized to make the replacement.



**CAUTION** The NX-2500's hard disk and power supply are NOT customer replaceable. Replacing either voids the warranty.



**CAUTION** The NX-2610's power supply is NOT customer replaceable. Replacing it voids the warranty.



**WARNING Battery Handling:** There is the danger of explosion if the battery is replaced incorrectly. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

## In This Chapter

- **Replacing a Hard Disk Drive** See page 410.
- **Replacing a Power Supply** See page 421.

## Replacing a Hard Disk Drive

The appliances feature RAID arrays with encrypted disk drives and redundant power supplies. Some models also use SSDs (Solid State Disks).

**RAID** stands for *Redundant Array of Independent (or Inexpensive) Disks*, a category of disk drives that employs two or more drives in combination for fault tolerance and performance.

With the exception of the NX-1700, NX-2500 and NX-2600 appliances (which only have one hard drive), Silver Peak ensures recoverability by mirroring data on paired hard drives. If a disk fails, the Appliance Manager displays a critical alarm, and the specific disk's LED stops illuminating on the appliance.

Before physically extracting any hard drive, you need to first “remove” it from the array by using the **Maintenance - Disk Management** page.

Displays the progress of a new disk that's being rebuilt from its array partner. Otherwise, this field displays **100** percent.

If a disk has been physically removed, the **Status** is **NOT-IN-SERVICE** and no **Serial Number** displays.

If a disk's **Status** is **DEGRADED**, you need to **Remove** it from the database,

**Maintenance - Disk Management**

| ID                          | Pairing Disk ID | Status         | Percent Complete (%) | Size (GB) | Serial Number      | Removable |
|-----------------------------|-----------------|----------------|----------------------|-----------|--------------------|-----------|
| <input type="checkbox"/> 0  | 1               | OK             | 100                  | 465       | 9SP19FXL           | yes       |
| <input type="checkbox"/> 1  | 0               | NOT-IN-SERVICE | 0                    | 465       |                    | yes       |
| <input type="checkbox"/> 2  |                 | OK             | 100                  | 59        | CVEM003300RV064KGN | yes       |
| <input type="checkbox"/> 3  |                 | OK             | 100                  | 59        | CVEM0041004W064KGN | yes       |
| <input type="checkbox"/> 4  |                 | OK             | 100                  | 59        | CVEM00410012064KGN | yes       |
| <input type="checkbox"/> 5  |                 | DEGRADED       | 100                  | 59        | CVEM0041004A064KGN | yes       |
| <input type="checkbox"/> 6  |                 | ---            | 0                    | 0         | ---                | no        |
| <input type="checkbox"/> 7  |                 | ---            | 0                    | 0         | ---                | no        |
| <input type="checkbox"/> 8  |                 | ---            | 0                    | 0         | ---                | no        |
| <input type="checkbox"/> 9  |                 | ---            | 0                    | 0         | ---                | no        |
| <input type="checkbox"/> 10 |                 | ---            | 0                    | 0         | ---                | no        |
| <input type="checkbox"/> 11 |                 | ---            | 0                    | 0         | ---                | no        |
| <input type="checkbox"/> 12 | 13              | OK             | 100                  | 465       | 9SP13JV2           | yes       |
| <input type="checkbox"/> 13 | 12              | OK             | 100                  | 465       | 9SP12MXM           | yes       |
| <input type="checkbox"/> 14 | 15              | OK             | 100                  | 465       | 9SP11HZC           | yes       |
| <input type="checkbox"/> 15 | 14              | OK             | 100                  | 465       | 9SP11HZC           | yes       |
| <input type="checkbox"/> 16 | 17              | OK             | 100                  | 465       | 9SP11HZC           | yes       |
| <input type="checkbox"/> 17 | 16              | OK             | 100                  | 465       | 9SP11HZC           | yes       |
| <input type="checkbox"/> 18 | 19              | OK             | 100                  | 465       | 9SP11HZC           | yes       |
| <input type="checkbox"/> 19 | 18              | OK             | 100                  | 465       | 9SP11HZC           | yes       |

Remove Insert

**Disk Layout**

Copyright © 2004-2010 Silver Peak Systems, Inc. All rights reserved.

For a pop-up showing the disk numbers and positions, roll the cursor over **Disk Layout**.

The following table summarizes information about replacing hard disks in specific appliance models:

| Appliance Model | #of hard drives | User can replace? | Hot-swappable? | For detailed directions for disk replacement, see ...                                          |
|-----------------|-----------------|-------------------|----------------|------------------------------------------------------------------------------------------------|
| NX-9700         | 14              | yes               | yes            | "To replace a disk in the NX-9700, NX-8700, NX-7700, NX-5700, NX-3700, or NX-2700" on page 412 |
| NX-9610         | 16              | yes               | yes            | "To replace a disk in the NX-9610, NX-8600, NX-7600, or NX-3600" on page 414                   |
| NX-8700         | 14              | yes               | yes            | page 412                                                                                       |
| NX-8600         | 16              | yes               | yes            | page 414                                                                                       |
| NX-8504         | 14              | yes               | yes            | "To replace a disk in the NX-8504, NX-7500, NX-7504, NX-5500, or NX-5504" on page 417          |
| NX-7700         | 10              | yes               | yes            | page 412                                                                                       |
| NX-7600         | 12              | yes               | yes            | page 414                                                                                       |
| NX-7504         | 8               | yes               | yes            | page 417                                                                                       |
| NX-7500         | 8               | yes               | yes            | page 417                                                                                       |
| NX-5700         | 8               | yes               | yes            | page 412                                                                                       |
| NX-5600         | 8               | yes               | yes            | page 414                                                                                       |
| NX-5504         | 8               | yes               | yes            | page 417                                                                                       |
| NX-5500         | 8               | yes               | yes            | page 417                                                                                       |
| NX-3700         | 2               | yes               | yes            | page 412                                                                                       |
| NX-3600         | 2               | yes               | no             | page 414                                                                                       |
| NX-3500         | 2               | yes               | no             | "To replace a disk in the NX-3500" on page 419                                                 |
| NX-2700         | 2               | yes               | yes            | page 412                                                                                       |
| NX-2610         | 2               | yes               | no             | "To replace a disk in the NX-2610" on page 420                                                 |
| NX-2600         | 1               | no                | --             | Contact Customer Support for return and repair instructions.                                   |
| NX-2500         | 1               | no                | --             | Contact Customer Support for return and repair instructions.                                   |
| NX-1700         | 1               | no                | --             | Contact Customer Support for return and repair instructions.                                   |


♦ **To replace a disk in the NX-9700, NX-8700, NX-7700, NX-5700, NX-3700, or NX-2700**

The first disk on the left is **Disk 0**. The numbers increment by one from left to right. These appliances' hard disks are hot-swappable.

**Note that the NX-9700 and NX-8700 appliances contain a mix of SATA hard disk drives and SSDs (solid-state drives).**


**NX-9700  
NX-8700**

Hard drives: 14




**NX-7700**

Hard drives: 10




**NX-5700**

Hard drives: 8



**NX-3700  
NX-2700**

Hard drives: 2



These are the two types of hard disks:

Solid-state disk (SSD) with spacer



SATA hard disk drive




- 1 On the **Maintenance - Disk Management** page, select the disk and click **Remove**.
- 2 Click **Save Changes**.
- 3 Unlatch the hard drive by pinching the latch together and then pulling the tab towards yourself.

Pinch the latch together.




Grasp the tab and pull forward to release.




- 4 Pull the disk out of its slot.
- 5 Insert the new disk and push until it clicks into place.

Push the top of the disk inward until it clicks into place.



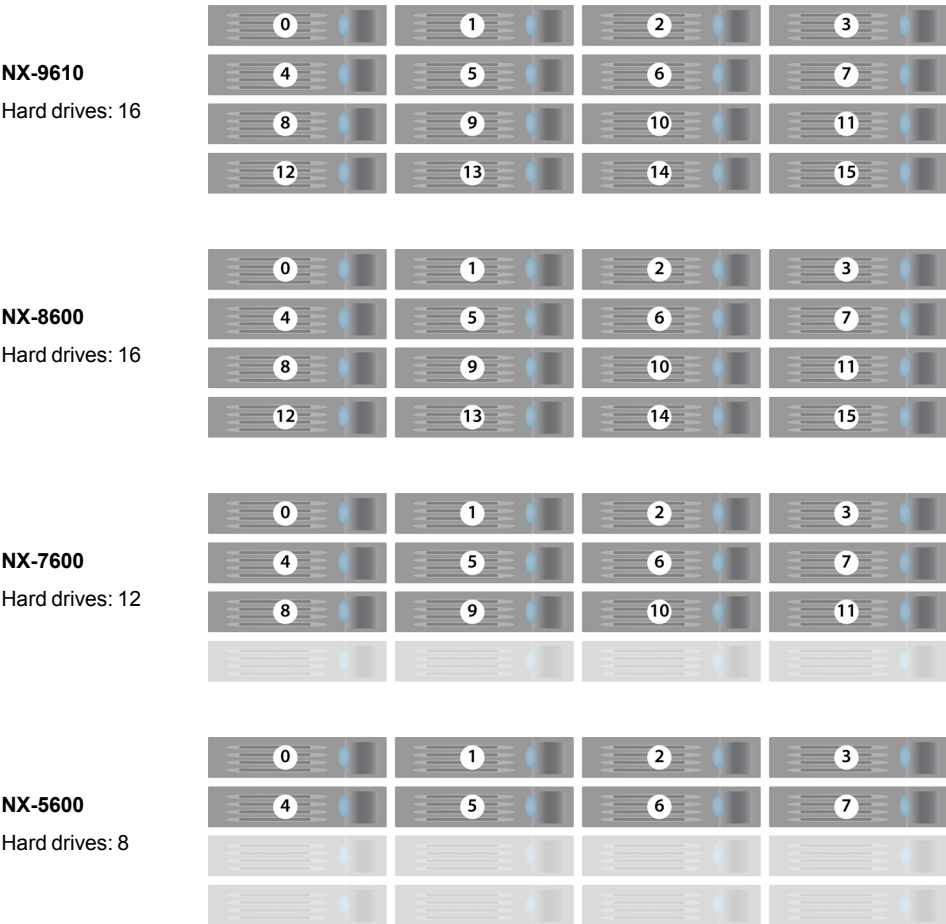
Push the latch against the tray to secure it.



- 6 On the **Maintenance - Disk Management** page, select the new disk and click **Insert**.  
The drive powers up.

◆ **To replace a disk in the NX-9610, NX-8600, NX-7600, or NX-3600**

The first disk on the left is **Disk 0**. The numbers increment by one from left to right.





These appliances' hard disks are hot-swappable.

- 1 On the **Maintenance - Disk Management** page, select the disk and click **Remove**.
- 2 Click **Save Changes**.
- 3 Unlatch the hard drive by pressing the end of the blue button toward the left and then pulling the tab towards yourself.

Depress the blue button leftward, into the tab.



Slip your finger behind the tab and pull forward to release.



- 4 Pull the disk out of its slot.
- 5 Insert the new disk and push until it clicks into place.

Push the tray inward until it clicks into place.



Push the tab against the tray to secure it.



- 6 On the **Maintenance - Disk Management** page, select the new disk and click **Insert**. The drive powers up.

♦ **To replace a disk in the NX-3600**



**CAUTION** The NX-3600's hard disks are **NOT** hot-swappable.



These two slots house the hard disks you can remove and replace.



- 1 On the **Maintenance - Disk Management** page, select the disk and click **Remove**.
- 2 Click **Save Changes**.
- 3 To power down the appliance, go to the **Maintenance - Restart System** page.
- 4 Click **Shutdown**.
- 5 After the drive powers down, unlatch the hard drive by pressing the end of the blue button toward the left and then pulling the tab towards yourself.

Although these photos show the NX-7600, the physical motions required to remove and re-insert the disks are accurate for the NX-3600.

Depress the blue button leftward, into the tab.



Slip your finger behind the tab and pull forward to release.



- 6 Pull the disk out of its slot.
- 7 Insert the new disk and push until it clicks into place.

Push the tray inward until it clicks into place.



Push the tab against the tray to secure it.




- 8 Power up the appliance.
- 9 On the **Maintenance - Disk Management** page, select the new disk and click **Insert**. The drive powers up.

♦ **To replace a disk in the NX-8504, NX-7500, NX-7504, NX-5500, or NX-5504**

The first disk on the left is **Disk 0**. The numbers increment by one from left to right.


**NX-8500**

Hard drives: 14




**NX-7500**

Hard drives: 8




**NX-5500**

Hard drives: 8




These appliances' hard disks are hot-swappable.

- 1 On the **Maintenance - Disk Management** page, select the disk and click **Remove**.
- 2 Click **Save Changes**.
- 3 To install the hard disk drive into the chassis, first remove the drive tray from the chassis.
- 4 Press the release tab located on the drive tray door to release the drive tray from its locking position, as shown below.



- 5 Pull the drive tray door upward and then pull the drive try out from the chassis.



- 6 Insert the new hard disk.
- 7 In the **Disk Information** area, select the new disk and click **Insert**.


♦ **To replace a disk in the NX-3500**




**CAUTION** The NX-3500's hard disks are **NOT** hot-swappable.

- 1 On the **Maintenance - Disk Management** page, select the disk and click **Remove**.
- 2 Click **Save Changes**.
- 3 To power down the appliance, go to the **Maintenance - Restart System** page.
- 4 Click **Shutdown**.
- 5 After the appliance has powered down, open the front cover of the NX-3500 by turning the knob one quarter-turn clockwise and pulling forward.


Move the vertical button from left to right.



Slip your finger behind the tab and pull forward to release.



Pull the disk out of its slot.




- 6 Insert the new disk by reversing the order of the steps above.
- 7 Power up the appliance.
- 8 In the **Disk Information** area, select the new disk and click **Insert**.



♦ **To replace a disk in the NX-2610**



**CAUTION** The NX-2610's hard disks are **NOT** hot-swappable.



These two slots house the hard disks you can remove and replace.

- 1 On the **Maintenance - Disk Management** page, select the disk and click **Remove**.
- 2 Click **Save Changes**.
- 3 To power down the appliance, go to the **Maintenance - Restart System** page.
- 4 Click **Shutdown**.
- 5 After the drive powers down, unlatch the hard drive by pressing the end of the blue button toward the left and then pulling the tab towards yourself.

Depress the blue button leftward, into the tab.



Slip your finger behind the tab and pull forward to release.



- 6 Pull the disk out of its slot.
- 7 Insert the new disk and push until it clicks into place.

Push the tray inward until it clicks into place.



Push the tab against the tray to secure it.



- 8 Power up the appliance.
- 9 On the **Maintenance - Disk Management** page, select the new disk and click **Insert**. The drive powers up.

## Replacing a Power Supply

The following table summarizes information about replacing power supplies in specific appliance models:

| Appliance Model | # of power supplies                                                                 | Hot-swappable when used redundantly? | Replacement instructions on ... |
|-----------------|-------------------------------------------------------------------------------------|--------------------------------------|---------------------------------|
| NX-9700         | 2                                                                                   | Yes                                  | page 422                        |
| NX-9610         | 3                                                                                   | Yes                                  | page 422                        |
| NX-8700         | 2                                                                                   | Yes                                  | page 422                        |
| NX-8600         | 3                                                                                   | Yes                                  | page 422                        |
| NX-8504         | 3                                                                                   | Yes                                  | page 423                        |
| NX-7700         | 2                                                                                   | Yes                                  | page 422                        |
| NX-7600         | 3                                                                                   | Yes                                  | page 422                        |
| NX-7504         | 3                                                                                   | Yes                                  | page 423                        |
| NX-7500         | 3                                                                                   | Yes                                  | page 423                        |
| NX-5700         | 2                                                                                   | Yes                                  | page 422                        |
| NX-5600         | 3                                                                                   | Yes                                  | page 422                        |
| NX-5504         | 2                                                                                   | Yes                                  | page 423                        |
| NX-5500         | 2                                                                                   | Yes                                  | page 423                        |
| NX-3700         | 2                                                                                   | Yes                                  | page 422                        |
| NX-3600         | 2                                                                                   | Yes                                  | page 422                        |
| NX-3500         | 2                                                                                   | Yes                                  | page 424                        |
| NX-2700         | 2                                                                                   | Yes                                  | page 422                        |
| NX-2610         | <b>Not</b> an authorized customer task.<br>Contact Customer Support for assistance. |                                      |                                 |
| NX-2600         | <b>Not</b> an authorized customer task.<br>Contact Customer Support for assistance. |                                      |                                 |
| NX-2500         | <b>Not</b> an authorized customer task.<br>Contact Customer Support for assistance. |                                      |                                 |
| NX-1700         | <b>Not</b> an authorized customer task.<br>Contact Customer Support for assistance. |                                      |                                 |



**WARNING** Do not open the casing of a power supply. Power supplies can only be accessed and serviced by a qualified technician from the manufacturer.

## Replacing a Power Supply in the NX-9700, NX-9610, NX-8700, NX-8600, NX-7700, NX-7600, NX-5700, NX-5600, NX-3700, NX-3600, or NX-2700



**CAUTION** Unplug the power cord before removing the power supply!!!



**Note** The photos are of the NX-x600 series. The power supplies in the NX-x700 appliances look recognizably similar.


The NX-3600 power supplies are oriented 90° counterclockwise from these photos.

### ♦ To access the power supply


Locate the release tab on the right side of the power supply.

1. Turn the screw counter-clockwise to loosen it.


2. To release the power supply from its locking position, squeeze the screw and the release tab together. Then hold it there while you ....



3. ...grip the handle to remove the power supply from the chassis.



4. Once the power supply module is released from its locking position, remove it from the chassis.



To insert a new power supply, repeat the procedure in reverse.



**WARNING** Do not open the casing of a power supply. Power supplies can only be accessed and serviced by a qualified technician from the manufacturer.




## Replacing a Power Supply in the NX-8504, NX-7500, NX-7504, NX-5500, or NX-5504




**CAUTION** Unplug the power cord before removing the power supply!!!

### ♦ To access the power supply

- 1 Locate the release tab on the left side of the power supply.
- 2 Push the release tab to the right to release the power supply from its locking position, as shown below:



- 3 Once the power supply module is released from its locking position, remove it from the chassis.



- 4 To insert a new power supply, repeat the procedure in reverse.



**WARNING** Opening the casing of a power supply voids the warranty. Only a qualified technician from the manufacturer has the authority to access and/or service power supplies.

## Replacing a Power Supply in the NX-3500




**CAUTION** Unplug the power cord before removing the power supply!!!

### ♦ To access the NX-3500 power supply


Locate the release tab on the right side of the power supply.

1. Loosen the screw
2. Push the release tab to the left and hold it there to release the power supply from its locking position while you ....




3. ...grip the handle to remove the power supply from the chassis.

[This picture shows an out-of-service appliance. However, if you were removing the power supply from a "live" unit, you would first unplug the power supply that you're removing and leave the other power supply plugged in and running to prevent an interruption of service.]



4. Once the power supply module is released from its locking position, remove it from the chassis.



To insert a new power supply, repeat the procedure in reverse.



**WARNING** Do not open the casing of a power supply. Power supplies can only be accessed and serviced by a qualified technician from the manufacturer.



## APPENDIX A

# Specifications, Compliance, and Regulatory Statements

This appendix contains specifications, as well as compliance and regulatory statements.

### In This Appendix

- **Model Specifications** See page 426.
- **Warning Statements** See page 433.
- **Compliance Statements** See page 434.
- **Cable Pinouts** See page 436.

## Model Specifications

This section includes general and model-specific specifications for the Silver Peak appliances:

- **Model-specific Specifications** See page 426.
- **Fiber Specifications** See page 431.
- **NX-Series Specifications** See page 431.
- 

### Model-specific Specifications

|                              |                                    | <b>NX-1700</b>                         | <b>NX-2500</b>                         | <b>NX-2600</b>                         |
|------------------------------|------------------------------------|----------------------------------------|----------------------------------------|----------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 4 Mbps                                 | 2 Mbps                                 | 4 Mbps                                 |
|                              | <i>Local Data Store</i>            | 500 GB                                 | 250 GB                                 | 250 GB                                 |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 2 x 10/100/1000 LAN<br>WAN             | 2 x 10/100/1000 LAN<br>WAN             | 2 x 10/100/1000 LAN<br>WAN             |
|                              | <i>Management</i>                  | 2 x 10/100/1000;<br>RS-232 serial port | 2 x 10/100/1000;<br>RS-232 serial port | 2 x 10/100/1000;<br>RS-232 serial port |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 47–63Hz,<br>90 W / 307 BTU  | 100–240VAC 50-60Hz,<br>120 W / 410 BTU | 100–240VAC 50-60Hz,<br>145 W / 496 BTU |
|                              | <i>Power Supplies</i>              | Single                                 | Single                                 | Single                                 |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 1.8 in. (45 mm) 1 RU                   | 1.7 in. (43 mm) 1 RU                   | 1.7 in. (43.5 mm) 1 RU                 |
|                              | <i>Width</i>                       | 17.5 in. (445 mm)                      | 16.8 in. (427 mm)                      | 16.9 in. (430 mm)                      |
|                              | <i>Depth</i>                       | 8.2 in. (209 mm)                       | 14.0 in. (356 mm)                      | 22.4 in. (569 mm)                      |
|                              | <i>Weight</i>                      | 8.5 lbs (3.9 kg)                       | 14 lbs (6.4 kg)                        | 22.0 lbs (10.0 kg)                     |

|                              |                                    | <b>NX-2610</b>                         | <b>NX-2700</b>                         |
|------------------------------|------------------------------------|----------------------------------------|----------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 8 Mbps                                 | 10 Mbps                                |
|                              | <i>Local Data Store</i>            | 500 GB                                 | 1 TB                                   |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 4 x 10/100/1000 LAN<br>WAN             | 4 x 10/100/1000 LAN<br>WAN             |
|                              | <i>Management</i>                  | 2 x 10/100/1000;<br>RS-232 serial port | 2 x 10/100/1000;<br>RS-232 serial port |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 50-60Hz,<br>165 W / 563 BTU | 100–240VAC 47-63Hz,<br>285 W / 973 BTU |
|                              | <i>Power Supplies</i>              | Single                                 | 1+1 redundant                          |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 1.7 in. (43.5 mm) 1 RU                 | 3.5 in. (89 mm) 2 RU                   |
|                              | <i>Width</i>                       | 16.9 in. (430 mm)                      | 16.9 in. (430 mm)                      |
|                              | <i>Depth</i>                       | 22.4 in. (569 mm)                      | 26 in. (660 mm)                        |
|                              | <i>Weight</i>                      | 24.2 lbs (11.0 kg)                     | 40.5 lbs (18.4 kg)                     |

|                              |                                    | <b>NX-3500</b>                         | <b>NX-3600</b>                         | <b>NX-3700</b>                          |
|------------------------------|------------------------------------|----------------------------------------|----------------------------------------|-----------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 10 Mbps                                | 20 Mbps                                | 20 Mbps                                 |
|                              | <i>Local Data Store</i>            | 500 GB                                 | 1 TB                                   | 1 TB                                    |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 2 x 10/100/1000 LAN<br>WAN             | 4 x 10/100/1000 LAN<br>WAN             | 4 x 10/100/1000 LAN<br>WAN              |
|                              | <i>Management</i>                  | 2 x 10/100/1000;<br>RS-232 serial port | 2 x 10/100/1000;<br>RS-232 serial port | 2 x 10/100/1000;<br>RS-232 serial port  |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 50-60Hz,<br>150 W / 512 BTU | 100–240VAC 47-63Hz,<br>250 W / 853 BTU | 100–240VAC 47-63Hz,<br>305 W / 1041 BTU |
|                              | <i>Power Supplies</i>              | 1+1 redundant                          | 1+1 redundant                          | 1+1 redundant                           |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 3.5 in. (89 mm) 2 RU                   | 3.5 in. (89 mm) 2 RU                   | 3.5 in. (89 mm) 2 RU                    |
|                              | <i>Width</i>                       | 17.9 in. (455 mm)                      | 17.0 in. (432 mm)                      | 16.9 in. (430 mm)                       |
|                              | <i>Depth</i>                       | 22.3 in. (566 mm)                      | 26.0 in. (661 mm)                      | 26 in. (660 mm)                         |
|                              | <i>Weight</i>                      | 34 lbs (15.4 kg)                       | 41.0 lbs (18.6 kg)                     | 40.5 lbs (18.4 kg)                      |

|                              |                                    | <b>NX-5500</b>                       | <b>NX-5504</b>                       | <b>NX-5600</b>                       |
|------------------------------|------------------------------------|--------------------------------------|--------------------------------------|--------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 50 Mbps                              | 50 Mbps                              | 50 Mbps                              |
|                              | <i>Local Data Store</i>            | 2 TB                                 | 2 TB                                 | 2 TB                                 |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 2 x 10/100/1000 LAN WAN              | 4 x 10/100/1000 LAN WAN              | 4 x 10/100/1000 LAN WAN              |
|                              | <i>Management</i>                  | 2 x 10/100/1000; RS-232 serial port  | 2 x 10/100/1000; RS-232 serial port  | 2 x 10/100/1000; RS-232 serial port  |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 50-60Hz, 400 W / 1366 BTU | 100–240VAC 50-60Hz, 400 W / 1366 BTU | 100–240VAC 50-60Hz, 440 W / 1501 BTU |
|                              | <i>Power Supplies</i>              | 2+1 redundant                        | 2+1 redundant                        | 2+1 redundant                        |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 5.2 in. (132 mm) 3 RU                | 5.2 in. (132 mm) 3 RU                | 5.2 in. (132 mm) 3 RU                |
|                              | <i>Width</i>                       | 17.7 in. (450 mm)                    | 17.7 in. (450 mm)                    | 17 in. (432 mm)                      |
|                              | <i>Depth</i>                       | 25.5 in. (647 mm)                    | 25.5 in. (647 mm)                    | 26 in. (659 mm)                      |
|                              | <i>Weight</i>                      | 65 lbs (29.5 kg)                     | 65 lbs (29.5 kg)                     | 62 lbs (28.1 kg)                     |

|                              |                                    | <b>NX-5700</b>                       |
|------------------------------|------------------------------------|--------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 50 Mbps                              |
|                              | <i>Local Data Store</i>            | 4 TB                                 |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 4 x 10/100/1000 LAN WAN              |
|                              | <i>Management</i>                  | 2 x 10/100/1000; RS-232 serial port  |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 47-63Hz, 345 W / 1178 BTU |
|                              | <i>Power Supplies</i>              | 1+1 redundant                        |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 3.5 in. (89 mm) 2 RU                 |
|                              | <i>Width</i>                       | 16.9 in. (430 mm)                    |
|                              | <i>Depth</i>                       | 26 in. (660 mm)                      |
|                              | <i>Weight</i>                      | 43 lbs (19.6 kg)                     |

|                              |                                    | <b>NX-7500</b>                          | <b>NX-7504</b>                          | <b>NX-7600</b>                          |
|------------------------------|------------------------------------|-----------------------------------------|-----------------------------------------|-----------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 155 Mbps                                | 155 Mbps                                | 155 Mbps                                |
|                              | <i>Local Data Store</i>            | 2 TB                                    | 2 TB                                    | 3 TB                                    |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 2 x 10/100/1000 LAN<br>WAN              | 4 x 10/100/1000 LAN<br>WAN              | 4 x 10/100/1000 LAN<br>WAN              |
|                              | <i>Management</i>                  | 2 x 10/100/1000;<br>RS-232 serial port  | 2 x 10/100/1000;<br>RS-232 serial port  | 2 x 10/100/1000;<br>RS-232 serial port  |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 50-60Hz,<br>400 W / 1366 BTU | 100–240VAC 50-60Hz,<br>400 W / 1366 BTU | 100–240VAC 50-60Hz,<br>580 W / 1979 BTU |
|                              | <i>Power Supplies</i>              | 2+1 redundant                           | 2+1 redundant                           | 2+1 redundant                           |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 5.2 in. (132 mm) 3 RU                   | 5.2 in. (132 mm) 3 RU                   | 5.2 in. (132 mm) 3 RU                   |
|                              | <i>Width</i>                       | 17.7 in. (450 mm)                       | 17.7 in. (450 mm)                       | 17 in. (432 mm)                         |
|                              | <i>Depth</i>                       | 25.5 in. (647 mm)                       | 25.5 in. (647 mm)                       | 26 in. (659 mm)                         |
|                              | <i>Weight</i>                      | 65 lbs (29.5 kg)                        | 65 lbs (29.5 kg)                        | 68 lbs (30.8 kg)                        |

| <b>NX-7700</b>               |                                    |                                         |
|------------------------------|------------------------------------|-----------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 155 Mbps                                |
|                              | <i>Local Data Store</i>            | 5 TB                                    |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 4 x 10/100/1000 LAN<br>WAN              |
|                              | <i>Management</i>                  | 2 x 10/100/1000;<br>RS-232 serial port  |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 47-63Hz,<br>475 W / 1621 BTU |
|                              | <i>Power Supplies</i>              | 1+1 redundant                           |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 3.5 in. (89 mm) 2 RU                    |
|                              | <i>Width</i>                       | 16.9 in. (430 mm)                       |
|                              | <i>Depth</i>                       | 26 in. (660 mm)                         |
|                              | <i>Weight</i>                      | 44 lbs (20 kg)                          |

|                              |                                    | <b>NX-8504</b>                       | <b>NX-8600</b>                       | <b>NX-8700</b>                                     |
|------------------------------|------------------------------------|--------------------------------------|--------------------------------------|----------------------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 500 Mbps                             | 500 Mbps                             | 622 Mbps                                           |
|                              | <i>Local Data Store</i>            | 8 TB                                 | 8 TB                                 | 5 TB + 4 x 64GB SSD                                |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 4 x 10/100/1000 LAN WAN              | 4 x 10/100/1000 LAN WAN              | 4 x 10/100/1000 LAN WAN; 2 x 10 Gbps fiber LAN WAN |
|                              | <i>Management</i>                  | 2 x 10/100/1000; RS-232 serial port  | 2 x 10/100/1000; RS-232 serial port  | 2 x 10/100/1000; RS-232 serial port                |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 50-60Hz, 530 W / 1809 BTU | 100–240VAC 50-60Hz, 650 W / 2218 BTU | 100-240VAC 47–63Hz, 520 W / 1775 BTU               |
|                              | <i>Power Supplies</i>              | 2+1 redundant                        | 2+1 redundant                        | 1+1 redundant                                      |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 5.2 in. (132 mm) 3 RU                | 5.2 in. (132 mm) 3 RU                | 3.5 in. (89 mm) 2 RU                               |
|                              | <i>Width</i>                       | 17.7 in. (450 mm)                    | 17 in. (432 mm)                      | 16.9 in. (430 mm)                                  |
|                              | <i>Depth</i>                       | 25.5 in. (647 mm)                    | 26 in. (659 mm)                      | 26 in. (660 mm)                                    |
|                              | <i>Weight</i>                      | 71 lbs (32.2 kg)                     | 75 lbs (34.0 kg)                     | 46.5 lbs (21.2 kg)                                 |

|                              |                                    | <b>NX-9610</b>                                                                 | <b>NX-9700</b>                                      |
|------------------------------|------------------------------------|--------------------------------------------------------------------------------|-----------------------------------------------------|
| <b>Capacity</b>              | <i>WAN Capacity (All Features)</i> | 1 Gbps                                                                         | 1 Gbps                                              |
|                              | <i>Local Data Store</i>            | 8 TB                                                                           | 5 TB + 4 x 64GB SSD                                 |
| <b>Connectivity</b>          | <i>LAN/WAN Ethernet</i>            | 4 x 1 Gbps fiber LAN WAN <sup>a</sup> ; 2 x 10 Gbps fiber LAN WAN <sup>b</sup> | 4 x 1 Gbps fiber LAN WAN; 2 x 10 Gbps fiber LAN WAN |
|                              | <i>Management</i>                  | 2 x 10/100/1000; RS-232 serial port                                            | 2 x 10/100/1000; RS-232 serial port                 |
| <b>Power</b>                 | <i>Requirement</i>                 | 100–240VAC 50-60Hz, 682 W / 2327 BTU                                           | 100-240VAC 47–63Hz, 600 W / 2048 BTU                |
|                              | <i>Power Supplies</i>              | 2+1 redundant                                                                  | 1+1 redundant                                       |
| <b>Dimensions and Weight</b> | <i>Height</i>                      | 5.2 in. (132 mm) 3 RU                                                          | 3.5 in. (89 mm) 2 RU                                |
|                              | <i>Width</i>                       | 17 in. (432 mm)                                                                | 16.9 in. (430 mm)                                   |
|                              | <i>Depth</i>                       | 26 in. (659 mm)                                                                | 26 in. (660 mm)                                     |
|                              | <i>Weight</i>                      | 75.5 lbs (34.5 kg)                                                             | 47 lbs (21.2 kg)                                    |

a. The connectors are LC and support both multi-mode 50 micron fiber and multi-mode 62.5 micron fiber.

b. The connectors are LC and support multi-mode 50 micron fiber.



| <b>GX-1000</b>                   |                                        |                                        |
|----------------------------------|----------------------------------------|----------------------------------------|
| <b>Capacity</b>                  | <i>WAN Capacity<br/>(All Features)</i> | N/A                                    |
|                                  | <i>Hard Drive</i>                      | 250 GB                                 |
| <b>Connectivity</b>              | <i>Management</i>                      | 2 x 10/100/1000;<br>RS-232 serial port |
| <b>Power</b>                     | <i>Requirement</i>                     | 100–240VAC 50-60Hz,<br>140 W / 478 BTU |
|                                  | <i>Power Supplies</i>                  | Single                                 |
| <b>Dimensions<br/>and Weight</b> | <i>Height</i>                          | 1.7 in. (43.5 mm) 1 RU                 |
|                                  | <i>Width</i>                           | 16.9 in. (430 mm)                      |
|                                  | <i>Depth</i>                           | 22.4 in. (569 mm)                      |
|                                  | <i>Weight</i>                          | 22.0 lbs (10.0 kg)                     |

## Fiber Specifications

|                | 1 Gbps Fiber Interfaces                                                                                                                   |                                      | 10 Gbps Fiber Interfaces                                                                                                    |                                        |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
|                | <b>lan0 / wan0<br/>Fiber Support</b>                                                                                                      | <b>lan0 / wan0<br/>Fail-to-Close</b> | <b>tlan0 / twan0<br/>Fiber Support</b>                                                                                      | <b>tlan0 / twan0<br/>Fail-to-Close</b> |
| <b>NX-8700</b> | <ul style="list-style-type: none"> <li>4 interfaces</li> <li>LC connectors</li> <li>Support multi-mode 50μ fiber / 62.5μ fiber</li> </ul> | <b>yes</b>                           | <ul style="list-style-type: none"> <li>2 interfaces</li> <li>LC connectors</li> <li>Support multi-mode 50μ fiber</li> </ul> | <b>no</b>                              |
| <b>NX-9610</b> |                                                                                                                                           | <b>no</b>                            |                                                                                                                             | <b>no</b>                              |
| <b>NX-9700</b> |                                                                                                                                           | <b>yes</b>                           |                                                                                                                             | <b>no</b>                              |

## NX-Series Specifications

|                      |                                |                                                                      |
|----------------------|--------------------------------|----------------------------------------------------------------------|
| <b>Environmental</b> | <i>Temperature (Operating)</i> | 10°C to 40°C (50°F to 104°F)                                         |
|                      | <i>Temperature (Storage)</i>   | -40°C to 65°C (-40°F to 149°F)                                       |
|                      | <i>Humidity</i>                | 8% to 90% relative humidity, non-condensing                          |
| <b>Regulatory</b>    | <i>EMC</i>                     | FCC Part 15 Class A, EN 55022 Class A,<br>EN 61000-3-2/3-3, EN 55024 |
|                      | <i>Safety</i>                  | UL/cUL 60950, EN 60950                                               |

## VX Series Specifications and Requirements

Following are the specifications and hardware requirements for the Silver Peak VX virtual appliances:

|                                        | <b>VX-1000</b>                                       | <b>VX-2000</b>                                        | <b>VX-3000</b>                                        | <b>VX-5000</b>                                        |
|----------------------------------------|------------------------------------------------------|-------------------------------------------------------|-------------------------------------------------------|-------------------------------------------------------|
| <b>WAN Capacity</b>                    | 4 Mbps                                               | 10 Mbps                                               | 20 Mbps                                               | 50 Mbps                                               |
| <b>Simultaneous Sessions</b>           | 8,000                                                | 64,000                                                | 64,000                                                | 64,000                                                |
| <b>Encryption</b>                      | Disk / WAN                                           | Disk / WAN                                            | Disk / WAN                                            | Disk / WAN                                            |
| <b>Hypervisors currently supported</b> | VMware ESX or ESXi (4.0 or greater)                  | VMware ESX or ESXi (4.0 or greater)                   | VMware ESX or ESXi (4.0 or greater)                   | VMware ESX or ESXi (4.0 or greater)                   |
| <b>CPUs</b>                            | Two 64-bit x86 CPUs, with a minimum speed of 2.3 GHz | Four 64-bit x86 CPUs, with a minimum speed of 2.3 GHz | Four 64-bit x86 CPUs, with a minimum speed of 2.3 GHz | Four 64-bit x86 CPUs, with a minimum speed of 2.3 GHz |
| <b>Memory (RAM)</b>                    | 2 GB                                                 | 4 GB                                                  | 4 GB                                                  | 8 GB                                                  |
| <b>Disk Space</b>                      | 100 GB of free contiguous disk space                 | 100 GB of free contiguous disk space                  | 100 GB of free contiguous disk space                  | 100 GB of free contiguous disk space                  |
| <b>Network Interfaces</b>              | 2 x 1 Gbps network interfaces                        | 2 x 1 Gbps network interfaces                         | 2 x 1 Gbps network interfaces                         | 2 x 1 Gbps network interfaces                         |

## Warning Statements

### **NX-9600**

- Class 1 Laser Product

## Compliance Statements

This section includes the following required compliance statements:

- **FCC Compliance Statement** See page 434.
- **ICES-003 statement** See page 434.
- **Requirements for Rack-Mount Equipment** See page 434.
- **Requirements for Knurled Thumb Screws** See page 435.

### FCC Compliance Statement

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

### ICES-003 statement

The Class A digital apparatus complies with Canadian ICES-003.

*Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.*

### Requirements for Rack-Mount Equipment

Observe the following requirements for all rack-mount equipment:


- 1 Elevated Operating Ambient Temperature** – If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T<sub>ma</sub>) specified by the manufacturer.
- 2 Reduced Air Flow** – Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.
- 3 Mechanical Loading** – Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.
- 4 Circuit Overloading** – Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring.  
  
Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.
- 5 Reliable Earthing** – Reliable earthing of rack-mounted equipment should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (for example, use of power strips).

## **Requirements for Knurled Thumb Screws**

Thumbscrews should be tightened with a tool after both initial installation and subsequent access to the panel.

## Cable Pinouts

Following is the pinout for the console (RS-232 serial) port, which uses a null modem cable.



**DB-9F**

- |   |       |
|---|-------|
| 2 | (RxD) |
| 3 | (TxD) |
| 5 | (GND) |
| 6 | (DSR) |
| 4 | (DTR) |
| 8 | (CTS) |
| 7 | (RTS) |

# Glossary

**802.1q encapsulation.** Also known as *VLAN tagging*. An IEEE standard (and process) which allows multiple bridged networks to transparently share the same physical network link without leakage of information between networks and, in common usage, the name of the encapsulation protocol used to implement this mechanism over Ethernet networks.

**ACL.** Access Control List.

**ARP.** Address Resolution Protocol. An IP protocol for finding a host's link layer (hardware) address when only its Internet Layer or some other Network Layer address is known.

**asymmetric routing.** When new writes can be made without having to wait for the secondary or remote storage site to also finish its writes.

**asynchronous replication.** A type of disk storage replication, where *write* is considered complete as soon as local storage acknowledges it. Remote storage is updated, but probably with a small lag. Performance is greatly increased, but in case of losing a local storage, the remote storage is not guaranteed to have the current copy of data and most recent data may be lost.

**authentication.** The process of validating the claimed identity of an end user or a device such as a host, server, switch, router, etc.

**authorization.** The act of granting access rights to a user, groups of users, system, or program.

**auto discovery.** Within the NX Series appliances, the ability of an appliance to discover and register with the Global Management System (GMS) server when first deployed.

**auto-negotiation.** The process by which terminating devices automatically negotiate for maximum bandwidth.

**bandwidth.** A rate of data transfer, throughput, or bit rate, measured in bits per second.

**bit.** A binary digit, taking a logical value of either "1" or "0" (also referred to as "true" or "false" respectively). It is also a unit of measurement, the information capacity of one binary digit.

**blan0.** When configuring for gigabit etherchannel bonding, **lan0** plus **lan1** bond to form **blan0**, which uses the **lan0** IP address.

**Bridge mode.** In-line deployment of an appliance, placing it between an Ethernet LAN switch and a WAN edge router.

**bwan0.** When configuring for gigabit etherchannel bonding, **wan0** plus **wan1** bond to form the virtual interface, **bwan0**, which uses the **wan0** IP address.

**bypass.** Refers to *hardware bypass*. If there is a major problem with the appliance hardware, software, or power, all traffic goes through the appliance without any processing. Additionally, you can manually put the appliance into Bypass as an aid to troubleshooting.

**chattiness.** A common problem with naively designed application protocols is that they are too "chatty". That is, they imply too many "round-trip" cycles.

**CIFS.** Common Internet File System. CIFS is the remote file system access protocol used by Windows servers and clients to share files across the network. Some specific capabilities of CIFS include file access, record locking, read/write privileges, change notification, server name resolution, request batching, and server authentication

**CIFS acceleration.** A set of techniques for mitigating the impacts of latency across the WAN. They include read-aheads and write-behinds to pipeline CIFS requests and the respective acknowledgements. This dramatically minimizes roundtrip delays when using CIFS over a WAN.

**CLI.** See *Command Line Interface*.

**client.** An application or system that accesses a remote service on another computer system, known as a *server*, by way of a network.

**Command Line Interface.** A method of configuring the appliance by typing in commands via the local serial interface or remote SSH session. [Peribit]

**CoS.** *Class of Service* (CoS) is a way of managing traffic in a network by grouping similar types of traffic (for example, e-mail, streaming video, voice, large document file transfer) together and treating each type as a class with its own level of service priority. Unlike Quality of Service (QoS) traffic management, Class of Service technologies do not guarantee a level of service in terms of bandwidth and delivery time; they offer a "best-effort." On the other hand, CoS technology is simpler to manage and more scalable as a network grows in structure and traffic volume. One can think of CoS as "coarsely-grained" traffic control and QoS as "finely-grained" traffic control.

**crossflow compression.** A technique that applies compression across various flows of traffic.

**data streaming.** The transfer of data at a steady high-speed rate sufficient to support such applications as high-definition television (HDTV) or the continuous backup copying to a storage medium of the data flow within a computer. Data streaming requires some combination of bandwidth sufficiency and, for real-time human perception of the data, the ability to make sure that enough data is being continuously received without any noticeable time lag.

**datagram.** An independent, self-contained message sent over the network whose arrival, arrival time, and content are not guaranteed.

**default gateway.** A gateway is a router on a computer network, serving as an access point to another network.

**DHCP.** Dynamic Host Configuration Protocol. A TCP/IP protocol that enables PCs and workstations to automatically get temporary or permanent IP addresses (out of a pool) from centrally administered servers.

**DNS.** Domain Naming System or Domain Name Server. It serves as the "phone book" for the Internet by translating human-friendly computer hostnames into IP addresses.



**DSCP.** Differentiated Services Code Point. A 6-bit value that encodes Per-Hop Behavior (PHB) into the 8-bit Differentiated Services (DS) field of the IP packet header. The DS field is the same as the TOS (Type of Service) field.

**domain.** The main purpose of a domain name is to provide a recognizable names to mostly numerically addressed Internet resources. This abstraction allows any resource (for example, website) to be moved to a different physical location in the address topology of the network, globally or locally in an intranet, in effect changing the IP address.

**failover.** The capability to switch over automatically to a redundant or standby computer server, system, or network upon the failure or abnormal termination of the previously active server, system, or network. Failover happens without human intervention and generally without warning, unlike *switchover*.

**FEC.** Forward Error Correction. When Adaptive Forward Error Correction (FEC) is enabled, the appliance introduces a parity packet, which helps detect and correct single-packet loss within a stream of packets, reducing the need for retransmissions. Silver Peak dynamically adjusts how often this parity packet is introduced in response to changing link conditions. This maximizes error correction while minimizing overhead.

**flow.** In a packet switching network, *packet flow* or *traffic flow* is a sequence of packets from a source computer to a destination, which may be another host, a multicast group, or a broadcast domain. As packets traverse successive communication links towards their destination, the packets from one flow (for example, A1, A2, A3) will be intermingled with packets from other flows also traversing the network to form a multiplexed stream (for example, A1, B7, C9, A2, C10, A3). This represents a form of statistical multiplexing because the link is shared as required.

**FTP.** File Transfer Protocol. A network protocol used to exchange and manipulate files over a TCP computer network, such as the Internet. An FTP client may connect to an FTP server to manipulate files on that server.

**full duplex.** Bidirectional, simultaneous two-way communications.

**gateway.** Also called *protocol converters*, can operate at any layer of the OSI model. The job of a gateway is much more complex than that of a router or switch. Typically, a gateway must convert one protocol stack into another.

**GMS.** Global Management System.

**GRE.** Generic Routing Encapsulation. Tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork.

**GUI.** Graphical User Interface.

**half duplex.** A circuit designed for data transmission in both directions, but not at the same time.

**hardware bypass.** If there is a major problem with the appliance hardware, software, or power, all traffic goes through the appliance without any processing. Additionally, you can manually put the appliance into Bypass as an aid to troubleshooting.

**header compression.** This technique can provide additional bandwidth gains by reducing packet header information using specialized compression algorithms.

**high availability.** For maximizing uptime, deploying NX appliances redundantly in 1+1 or N+1 configurations, with failover and load balancing.

**host.** In computer networking, a network host, Internet host or host is a computer connected to the Internet. A network host can host information as well as client and/or server software.

**host address.** The host address, or more properly the host id portion of an IP address is the portion of the address used to identify hosts (which can be any device requiring a Network Interface Card, such as a personal computer or networked printer) on the network.

**HTTP.** HyperText Transfer Protocol. The protocol web browsers use to communicate with web servers.

**HTTPS.** HyperText Transfer Protocol Secure. A combination of the HyperText Transfer Protocol and a cryptographic protocol, for accessing a secure web server.

**ICMP.** Internet Control Message Protocol. An internet protocol used by networked computers' operating systems to manage errors and generate control messages.

**Internet.** A global network of interconnected computers, enabling users to share information along multiple channels.

**IP.** Internet Protocol. Network layer protocol in the TCP/IP stack that enables a connectionless internetwork service.

**IP Address.** An Internet Protocol (IP) address is a numerical identification and logical address that is assigned to devices participating in a computer network utilizing the Internet Protocol for communication between its nodes.

**IPSec.** Internet Protocol Security Protocol.

**IP VPN.** Internet Virtual Private Network.

**LAN.** Local Area Network.

**LAN Rx.** Traffic received from the LAN.

**LAN Tx.** Traffic transmitted to the LAN.

**latency.** A time delay between the moment something is initiated, and the moment one of its effects begins or becomes detectable. Network latency is the time it takes for information to go from a sender to a receiver and back.

**load balancing.** A technique to spread work between two or more computers, network links, CPUs, hard drives, or other resources, in order to get optimal resource utilization, maximize throughput, and minimize response time. Using multiple components with load balancing, instead of a single component, may increase reliability through redundancy. The balancing service is usually provided by a dedicated program or hardware device.

**lossy.** A WAN prone to dropped and out-of-order packets. This is most common on shared networks, like MPLS and Internet VPNs.

**MAPI.** Messaging Application Programming Interface. A Microsoft Windows program interface that enables you to send e-mail from within a Windows application and attach the document you are working on to the e-mail note. Applications that take advantage of MAPI include word processors, spreadsheets, and graphics applications.

**MIB.** Management Information Base. A type of database for managing devices in a communications network.

**Microsoft Exchange.** Messaging and groupware software for Windows from Microsoft. The Exchange server is an Internet-compliant messaging system that runs under Windows systems and can be accessed by web browsers, the Windows In-box, Exchange client or Outlook. The Exchange server also stores files for sharing.

**MPLS.** MultiProtocol Label Switching is an IETF initiative that integrates Layer 2 information into Layer 3 (IP) packets.

**MTU.** Maximum Transmission Unit. The largest size packet that a device can transmit on a network.

**Network Acceleration.** Addresses high WAN latency and TCP chattiness. This is achieved using standard TCP acceleration techniques, such as adjustable windows and selective acknowledgements.

**Network Integrity.** Protects traffic from collateral congestion in a shared service provider network by mitigating the impact of dropped and out-of-order packets.

**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.

**NFS.** Network File System. The file sharing protocol in a UNIX network.

**OOO.** Out-of-Order [packets]

**out-of-path.** Same as **Router mode**. In an out-of-path deployment, policy-based routing (PBR), VRRP, or WCCP redirect the traffic to the Silver Peak appliance for processing.

**packet coalescing.** When packets are small, packet headers consume substantial bandwidth in comparison to the amount of end-user data transferred. Packet coalescing combines multiple user packets traveling between the same two sites into a single coalesced packet. Used in conjunction with header compression, this amortizes a single header over multiple packets thus decreasing overhead, and therefore bandwidth requirements. Packet coalescing is particularly beneficial for web applications, VoIP, and interactive applications, like Citrix.

**pass-through traffic.** Traffic that is sent to the WAN without being optimized.

**payload compression.** Uses algorithms to identify relatively short byte sequences that are repeated frequently over time. These sequences 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

**PBR.** Policy-based routing is a technique used to make routing decisions based on policies set by the network administrator.

**Propagate Link Down.** Forces the WAN interface to go down when the corresponding LAN interface goes down, or vice versa. By default, this option is enabled on the **Configuration - System** page.

**ping.** A programs used to test whether a particular network destination is online, by sending an Internet Control Message Protocol (ICMP) echo request and waiting for a response. [Peribit]

**POC.** Packet Order Correction. To avoid retransmissions that occur when packets arrive out of order, Silver Peak NX appliances use Packet Order Correction (POC) to resequence packets on the far end of a WAN link, as needed.

**QoS.** Quality of Service is the ability to provide different priority to different applications, users, or data flows, or to guarantee a certain level of performance to a data flow. QoS involves several functions: 1) classification of packets into traffic classes based on characteristics such as source, destination addresses,

and/or applications and 2) queuing and service mechanisms that are used to apply service policies based on these classifications, including bandwidth allocation.

**RADIUS.** Remote Authentication Dial In User Service (RADIUS) is a networking protocol that provides centralized Authentication, Authorization and Accounting (AAA) management for computers to connect and use a network service. It is a client/server protocol that uses UDP as transport.

**Router mode.** Out-of-path deployment, where data traffic is redirected by using policy-based routing (PBR), Web Cache Coordination Protocol (WCCP), or Virtual Router Redundancy Protocol (VRRP).

**RTT.** Round-trip time. the time it takes to send a packet to a remote host and receive a response; used to measure delay on a network at a given time. [Peribit]

**SMB.** Server Message Block. An application-level network protocol mainly used to provide shared access to files, printers, serial ports, and miscellaneous communications between nodes on a network.

**SMB2.** Server Message Block, version 2.

**SMTP.** Simple Mail Transfer Protocol. A de facto standard for electronic mail (e-mail) transmissions across the Internet.

**SNMP.** Simple Network Management Protocol. A standard TCP/IP protocol for network management. Network administrators use SNMP to monitor network devices, performance, and security, and to manage configurations and collect statistics.

**SSL.** Secure Socket Layer. These are cryptographic protocols that provide secure communications for such things web browsing, email, and other data transfers over the internet.

**subnet.** A portion of a network that shares a common address component. On TCP/IP networks, subnets are defined as all devices whose IP addresses have the same prefix. For example, all devices with IP addresses that start with 100.100.100. would be part of the same subnet. Dividing a network into subnets is useful for both security and performance reasons. IP networks are divided using a subnet mask.

**switch.** A network device that filters and forwards frames based on the destination address of each frame. The switch operates at Layer-2 (data link layer) of the Open System Interconnection (OSI) model.

**TACACS+.** Terminal Access Controller Access-Control System Plus is a protocol which provides access control for routers, network access servers and other networked computing devices via one or more centralized servers. TACACS+ provides separate authentication, authorization and accounting services. It uses TCP for its transport. Transactions between the TACACS+ client and TACACS+ servers are also authenticated through the use of a shared secret.

**TCP.** Transmission Control Protocol. The error-correcting Transport layer (Layer-4) in the TCP/IP protocol suite. It ensures that all data arrive at the other end accurately and completely intact.

**TCP acceleration.** A set of techniques for mitigating the impacts of latency across the WAN. They include adjustable window sizing and selective acknowledgements.

**TCP/IP.** Transmission Control Protocol/Internet Protocol. A protocol suite for communication between computers, used as a standard for transmitting data over networks and as the basis for standard Internet protocols.

**Telnet.** A terminal emulation protocol used on the Internet and TCP/IP-based networks. A Telnet program allows a user at a terminal or PC to log in to a remote computer and run a program and execute other Unix commands.

**throughput.** The average rate of successful message delivery over a communication channel.

**tunneling.** Encapsulating one type of network protocol (called the payload protocol) within a different delivery protocol. Reasons to use tunneling include carrying a payload over an incompatible delivery network, or to provide a secure path through an untrusted network.

**UDP.** User Datagram Protocol. Part of the TCP/IP protocol suite, it was created to provide a way for applications to access the connectionless features of IP. UDP provides for exchange of datagrams without acknowledgements or guaranteed delivery.

**VLAN.** Virtual Local Area Network. A means by which LAN users on different physical LAN segments are afforded priority access privileges across the LAN backbone so that they appear to be on the same physical segment of an enterprise-level logical LAN.

**VLAN tag.** *See 802.1q encapsulation.*

**VoIP.** Voice-Over-Internet-Protocol. A protocol optimized for the transmission of voice through the Internet or other packet-switched networks.

**VRRP.** Virtual Router Redundancy Protocol is a standard redundancy protocol designed to increase the availability of servicing hosts on the same subnet.

**WAN.** Wide Area Network

**WAN Rx.** Traffic received from the WAN.

**WAN Tx.** Traffic transmitted to the WAN.

**WCCP.** Web Cache Communications Protocol. A Cisco-developed content-routing protocol that provides a mechanism to redirect traffic flows in real-time. It has built-in load balancing, scaling, fault tolerance, and service-assurance (failsafe) mechanisms.

**X.11.** An application redirect protocol; a distributed window system that is based on the client/server model.



# Index

## Numerics

- 10G interfaces 95
- 1G interfaces 95
- 802.1q trunked links
  - See* VLANs

## A

- Access Control Lists 135–144
  - application groups in 156
  - as reusable MATCH criteria 132
  - characteristics of 135
  - creating 136–138
  - how they filter traffic 145–146
  - modifying an ACL rule 139
  - removing an ACL 141–144
  - removing an ACL rule 140

- ACLs
  - See* Access Control Lists

- Alarm Log Viewer 316

- alarms
  - clearing 402
  - current 402
    - viewing 406–407
  - handling of 408
  - list of types and text 403–404
  - severity levels 402

- appliance configuration file
  - See* configuration file

- Appliance Manager
  - accessing 74
  - Home - Summary page 76–80
  - log in 75
  - menu structure
    - Administration menus 86–87
    - Alarm menus 88
    - Configuration menus 80–82
    - Maintenance menus 87
    - Monitoring menus 84–85
    - Reporting menus 83
  - system requirements 74
  - tables, managing 88
- application groups
  - creating 156–157

- modifying 158
  - properties of 156
  - removing from the database 159
  - using in MATCH criteria 156
- applications
  - built-in
    - list of 147–151
    - viewing 151
  - custom, defining and creating 152–155
  - historical statistics 244, 248–250
  - realtime statistics 264, 270–272
- asymmetric networks or flows
  - See* flow redirection
- Audit Log 317
- authentication and authorization 9, 337, 338
- authorization 337
  - for SNMP v3 user 331
- Auto Optimization 162
  - SET actions diagram 166
- auto optimization
  - Bridge mode handshaking 163
  - Router mode handshaking 164

## B

- bandwidth
  - historical statistics 244, 253–254
  - realtime statistics 264, 295–296
- bandwidth management
  - auto bandwidth 189
  - best practices 184–187
  - Configuration pages for 187
  - configuring minimum and maximum values 185–186
  - maximum system bandwidth, determining 188
  - for multiple tunnels 184
  - multiple traffic classes 184
  - for pass-through traffic 190
- bandwidth shaping 116
  - See also* **Configuration - Tunnels**
  - page
- banners

- login message 354
  - Message of the Day 354
- bitmasks 88
- blan0**
  - See* etherchannel bonding
- bridge mode
  - 2-port bridge 11, 96
  - 4-port 12–14
  - 4-port bonding 96
  - 4-port, flat 96
  - 4-port, standard 96
  - realtime statistics 265, 307

- bwan0**
  - See* etherchannel bonding
- bypass
  - See* System Bypass

## C

- cabling 68
- CIFS Acceleration 213, 214
  - disabled in Tunnel Compatibility Mode 126
- cluster interface 230, 233
- configuration
  - Initial Configuration Wizard 356–357
  - saving the running configuration 91
- configuration file
  - downloading 380, 381
    - from a URL 383
    - from an FTP server 385
    - from an SCP server 384
    - from a local disk 382
  - saving 375, 376
    - to a local disk 377
    - to an FTP server 379
    - to an SCP server 378
  - viewing
    - last saved version 373
    - the running config file 374
- configuration, cabling 68
- configuration
  - Initial Configuration Wizard 69–71

current flows

- customizing which columns display 275, 276
- details of 280, 281–284
- filter for specific results 277–280
- orange background 278
- pink background 278
- realtime statistics 264, 273–285
- reduced tunnel functionality 278
- resetting for improved performance 285
- resetting the default display 276
- selecting filters 274
- stale connection 278
- unaccelerated TCP 285
- See also* flows

## D

datapath connectivity 112, 308

date and time

- conventions 90

Debug Dump 318, 319

debug files

- Debug Dump 318, 319
- deleting 324
- Log 318, 319
- saving to
  - a local PC 321
  - a remote server 320
  - an FTP server 323
  - an SCP server 322
- Show Tech 318, 319
- Snapshot 318, 319
- TCP Dump Result 318, 319
- types of 318
- See also* logs

debugging 214

deployment

- Initial Configuration Wizard 356–357

deployment options 10, 96–98

- in-line
  - Bridge mode - four ports 12–14
  - Bridge mode - two ports 11
- out-of-path 15
  - with Policy-Based Routing (PBR) Redirection 16
  - with Policy-Based-Routing (PBR) and VRRP redundant Silver

- Peak appliances 19
  - with VRRP peering to WAN router 18
- with Web Cache Coordination Protocol (WCCP) 17
- with Web Cache Coordination Protocol (WCCP) redundant Silver Peak appliances 20

DNS server, adding 113

DSCP markings 180, 181

- apply to
  - pass-through traffic 197–199
- applying to
  - optimized traffic 194–196
- definitions list 199–200

## E

encapsulation 116

encryption

- for SNMP v3 user 331

hard disk 360

etherchannel bonding 94

- configuring 99
- gigabit, for 4-port devices 95, 99

Ethernet ports

- See* etherchannel bonding

Events Log 315

## F

FEC

- See* Forward Error Correction

field definitions 91

flow counts 292

- historical statistics 244, 255–256
- realtime statistics 264, 297–298
- TCP and non-TCP 291
- See also* current flows

flow exports, for NetFlow 114

flow redirection

- asymmetric networks and flows 226
- configuration example 229–233
- for LAN-initiated traffic 228
- realtime statistics 265, 302–303
- for removing asymmetry 226
- reporting 234
- for WAN-initiated traffic 227

flows, processing of 128

Forward Error Correction

- historical statistics 258–260

- realtime statistics 264, 291, 292, 300–301
- in Tunnel Compatibility Mode 126

FTP server capability in appliance 325

## H

hard disk replacement 409, 410

- NX-2610 420
- NX-3500 419
- NX-36000 416
- NX-8504, NX-7500, NX-7504, NX-5500, NX-5504 417, 418
- NX-9610, NX-8600, NX-7600, NX-5600 415
- NX-9610, NX-8600, NX-7600, NX-5600, NX-3600 414
- NX-9700, NX-8700, NX-7700, NX-5700, NX-3700, NX-2700 412

hardware bypass mode

- See* System Bypass

hardware components

- NX-1700 3
- NX-2500 3
- NX-2600 3
- NX-2610 3
- NX-2700 4
- NX-3500 4
- NX-3600 4
- NX-3700 4
- NX-5500 4
- NX-5600 5
- NX-5700 5
- NX-7500 5
- NX-7600 5
- NX-7700 5
- NX-8500 6
- NX-8600 6
- NX-8700 6
- NX-9610 6
- NX-9700 6

hardware installation tasks 22

help icons 91

Hostname

- modifying 95
- See also* Initial Config Wizard

## I

inbound traffic 245, 266, 290

Initial Configuration Wizard 69–71, 356–357



interfaces  
   1G and 10G 95  
   configuring from table 106  
   manually configuring for DHCP 104  
   realtime statistics 265, 305–306

IP routes  
   datapath connectivity 112  
   realtime statistics 265, 308  
   removing a LAN route from the routing table 112  
   static routes for management traffic 108  
   unreachable 112

IP routes, configuring  
   default gateways for management interfaces 107  
   default LAN-side gateway 112  
   next-hop address(es) for LAN-side networks 107, 111  
   WAN next-hop address(es) 107

## L

latency 292  
   historical statistics 257  
   realtime statistics 264, 299  
   statistics 244, 291

LEDs  
   *See also* power, connecting and verifying  
   verifying connectivity 66

log settings  
   configuring parameters for event and alarm logs 310  
   local logging 312  
   minimum severity level 310, 314  
   remote syslog server 313

login message banner 354

logs  
   Alarm Log Viewer 316  
   Audit Log 317  
   deleting files 324  
   Event Log Viewer 315

lost packets  
   *See* FEC

## M

management interface, configuring  
   next-hop IP address 109  
   static routes 108

management methods 9

management route  
   removing from the routing table 110

MATCH criteria 162, 166  
   5-tuple 128  
   application groups in 156  
   basis of 128  
   configuring 131  
   how they filter traffic 145–146  
   specifying applications and protocols in 133–134  
   understanding 130–134  
   using ACLs to summarize 132

Max System Bandwidth  
   configuring 94  
   *See also* bandwidth management, best practices

Message of the Day 354

MIBs, list of standard and proprietary 327

minimum severity level 310  
   modifying the remote receiver 314

## N

naming objects 90

NetFlow  
   configuring flow exports for 114  
   realtime statistics 265, 304

netmask notation 88  
   table of 89

Network Acceleration  
   description of 8

network connectivity, testing 386–388  
   *See also* ping, traceroute, and tcpdump

Network Integrity 292  
   description of 8  
   historical statistics 244, 258–260  
   realtime statistics 264, 300–301  
   *See* Forward Error Correction (FEC)  
   *See* Packet Order Correction (POC)

Network Memory 214  
   benefit scenarios 212  
   definition 212  
   description of 8  
   disabled in Tunnel Compatibility Mode 126  
   erasing 399  
   hard disk encryption 360  
   pre-positioning data into 325, 326  
   *See also* Optimization policies

networking parameters  
   related menus 103

next-hops, list of 265, 308

NTP server, configuring 101–102

NX-8700  
   configuring 1G or 10G 95

NX-9610  
   configuring 1G or 10G 95

NX-9700  
   configuring 1G or 10G 95

## O

CIFS Acceleration  
   *See also* Optimization policies 213

Optimization policies 128, 129, 166, 167, 212–224  
   activating a new policy 224  
   adding a new Optimization map 220–222  
   adding an entry to a map 217–218  
   Configuration page organization 216  
   default behaviors 129  
   deleting an entry 220  
   deleting an inactive map 223  
   editing an active entry 219  
   when the appliance can apply them 216

Payload Compression  
   *See also* Optimization policies 213

TCP Acceleration  
   *See also* Optimization policies 213

optimized traffic flows  
   how to begin 129  
   processing of 128

outbound traffic 245, 266, 290

out-of-order packets  
   *See* POC

Out-of-Order Packets (OOP)  
   *See* Packet Order Correction

## P

packet coalescing  
   in Tunnel Compatibility Mode 126

Packet Order Correction  
   historical statistics 258–260  
   realtime statistics 264, 291, 292, 300–301  
   in Tunnel Compatibility Mode 126

pass-through traffic  
   applying DSCP markings 197–199  
   configuring bandwidth limits 190

password

- for SNMP v3 user 331
- guidelines for creating 90
- guidelines for user accounts 332
- Payload Compression 213, 214
  - disabled 126
- pie charts, viewing options for 246, 268
- ping 387, 389–390
- POC
  - See* Packet Order Correction
- power supply replacement 409
  - NX-3500 424
  - NX-8504, NX-7500, NX-7504, NX-5500, NX-5504 423
  - NX-9700, NX-9610, NX-8700, NX-8600, NX-7700, NX-7600, NX-5700, NX-5600, NX-3700, NX-3600, NX-2700 422
- power, connecting and verifying
  - NX-1700 57
  - NX-2500 57
  - NX-2600, NX-2610 58
  - NX-2700 63
  - NX-3500 59
  - NX-3700 63
  - NX-5500, NX-5504, NX-7500, NX-7504, NX-8504 60–61
  - NX-5600, NX-7600, NX-8600, NX-9610 62
  - NX-5700 63
  - NX-7700 63
  - NX-8700 63
  - NX-9700 63
- pre-positioning data 325, 326
- Propagate Link Down 94, 95

## Q

- QoS policies 128, 179–209
  - activating a new policy 209
  - adding a new QoS map 206–207
  - adding an entry to a map 202–203
  - Configuration page organization 201
  - default behavior 180
  - default behaviors 129
  - deleting an entry 205
  - deleting an inactive map 208
  - editing an entry 204
  - effective use of 180
- QoS shaping and marking 126
- QoS statistics
  - tunnel, realtime 264, 286–287

## R

- rack mount instructions
  - NX-1700 25–26
  - NX-2500 27
  - NX-2600, NX-2610 28–33
  - NX-2700, NX-3700, NX-5700, NX-7700, NX-8700, NX-9700 51–56
  - NX-3500 33
  - NX-3600 34–40
  - NX-5500, NX-5504, NX-7500, NX-7504, NX-8504 41–44
  - NX-5600, NX-7600, NX-8600, NX-9610 45–50
- RADIUS 337
  - configuration 339–340
  - example, integrating NX appliance into network 341–345
- reboot clean 400
- rebooting the appliance 400
- reduction
  - historical statistics 244, 251–252
  - realtime statistics 264, 293–294
- Reorder Wait Time 292
- restarting the appliance 400
- Route policies 128, 161–178
  - activating a new policy 178
  - adding a new Route map 175–177
  - adding an entry to a map 172–173
  - auto optimization 162, 163–164
  - auto-optimizing TCP flows 129
  - Configuration page organization 171
  - default behaviors 129
  - deleting an entry 175
  - deleting an inactive map 178
  - editing an active entry 174
  - optimizing non-TCP traffic 129
  - where to direct flows 165–170
- router mode
  - dual-homed 97
  - dual-homed bonding 97
  - Out-of-Path with Policy-Based Routing (PBR) redirection 16
  - Out-of-Path with Policy-Based-Routing (PBR) and VRRP redundant Silver Peak appliances 19
  - Out-of-Path with VRRP peering to WAN router 18
  - Out-of-Path with Web Cache Coordination Protocol (WCCP) 17

- Out-of-Path with Web Cache Coordination Protocol (WCCP)
  - redundant Silver Peak appliances 20
- standard 97
- standard bonding 97
- running configuration 91

## S

- Save 160
- server virtualization 7
- session idle time-out 338
- SET actions 128, 129, 130, 131, 135, 137, 145
  - in Optimization policies 212, 214, 216
  - in QoS policies 180–183, 201
  - in Route policies 162, 165–170, 171
- Show Tech 318, 319
- shutdown 400
- site preparation 22–24
- slash notation 88
- Snapshot 318, 319
- SNMP
  - adding an SNMP v3 user 331
  - configuring SNMP settings 328
  - loading SNMP MIBs 327
  - trap receiver
    - adding 329
    - modifying 330
    - removing 330
- software management
  - installing a software image
    - from a URL 366
    - from an FTP server 369, 370
    - from an SCP server 367, 368
    - from the local disk 365
    - into a partition 364
  - options 361–364
  - switching partitions 371
- software version
  - listed 360
- statistics
  - about viewing 244–247, 266–269
  - counters
    - clearing non-destructively 267
    - view since reboot 266
  - Delta Stats 267
  - historical
    - applications 244, 248–250

- bandwidth 244, 253–254
  - flow counts 244, 255–256
  - Forward Error Correction
    - 244, 258–260
  - latency 244, 257
  - Network Integrity 244, 258–260
  - Packet Order Correction 244, 258–260
  - reduction 244, 251–252
  - summary report 244, 261–262
  - pre-defined time periods 247
  - realtime
    - applications 264, 270–272
    - bandwidth 264, 295–296
    - bridge mode 265, 307
    - current flows 264, 273–285
    - flow counts 264, 297–298
    - flow redirection 265, 302–303
    - Forward Error Correction
      - 264, 300–301
    - interfaces 265, 305–306
    - IP routes 265, 308
    - latency 264, 299
    - NetFlow 265, 304
    - Network Integrity 264, 300–301
    - Packet Order Correction 264, 300–301
    - reduction 264, 293–294
    - tunnel 264, 288–292
    - tunnel QoS 264, 286–287
  - refreshing 267
  - Support, contacting 358
  - syslog server
    - See* log settings
  - System Bypass 94, 95, 117, 360
  - system configuration settings
    - date and time 100
    - NTP server, configuring 101–102
  - system deployment
    - See* deployment
  - system information, displayed 360
  - system requirements for Appliance Manager 74
- T**
- table data, exporting 247, 269
  - TACACS+ 337
    - configuration 346–347
  - example, configuring server to
    - authenticate NX appliance users 347–353
  - TCP Acceleration 213, 214
    - disabled 126
  - TCP flows
    - See* flow counts
    - asymmetric
      - See* flow redirection
  - tcpdump 318, 319, 388, 393–396
    - retrieving results 397–398
  - Technical Support, contacting 358
  - time periods, how measured 247
  - timestamp 90
  - tlan0** 95
  - traceroute 388, 391–392
  - traffic
    - direction of flows 245, 266
    - inbound 245, 266
    - outbound 245, 266
    - processing flows with policies 128
  - traffic classes
    - configuring for tunnels and
      - pass-through traffic 191–193
    - defaults 180
    - for optimized flows 181
    - shaped pass-through traffic 182
    - tunnels 181
  - tunnel
    - characteristics 116
    - creating a 117–120
    - deleting a 125
    - diagram of directing flows to 165–166
    - editing a 123, 124
    - encapsulation 126
    - parameters
      - FEC 123, 124, 126
      - POC 126
      - Pre Shared Key 124
      - Reorder Wait 124
    - QoS shaping and marking 126
    - QoS statistics, realtime 264, 286–287
    - realtime statistics 264, 288–292
    - reduced functionality, indicating 126
    - states, list of 122
    - status, verifying 121
    - traffic
      - how Route Policies affect 116
    - Tunnel Compatibility Mode 126
      - disabled optimizations 126
  - preserved functionalities 126
- U**
- uptime, of appliance 360
  - user access management 9
  - user accounts
    - accessing 333
    - admin* and *monitor* privileges 334
    - authentication and authorization 337, 338
    - built-in user database 333, 337
    - creating 334
    - deleting 336
    - modifying 335
    - password guidelines 332
    - session idle time-out 338
- V**
- virtual appliances 3, 7
  - VLANs 236–242
    - configuring a VLAN IP interface 237
    - in Current Flows Details 283
    - setting tags in outgoing WAN-side packets 238
      - 2-port Bridge 239
      - 4-port Bridge 240
      - bonded 4-port Bridge 242
      - flat 4-port Bridge 241
    - tags 128
    - why create a VLAN IP interface 236
  - VX appliances 3, 7
- W**
- WAN interface
    - modifying total bandwidth 94, 95
  - web
    - protocol settings 355
    - user settings 355







Silver Peak

Silver Peak Systems, Inc.  
4500 Great America Parkway, Suite 100  
Santa Clara, CA 95054

1.877.210.7325  
+1.408.935.1850

[www.silver-peak.com](http://www.silver-peak.com)