

Central Management System (CMS) 5.1 Administrator's Guide

*Central Management and Configuration Software for
Sequence Reducer and Sequence Mirror Devices*



Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
USA
408-745-9500
www.juniper.net

Part number 100872 Rev. 001

Central Management System (CMS) Administrator's Guide

Copyright © 2005 Juniper Networks, Inc. All Rights Reserved. Printed in USA.

Copyright © 2003-2005 Peribit Networks, Inc. All Rights Reserved. Printed in USA.

Juniper Networks, the Juniper Networks logo, NetScreen, NetScreen Technologies, the NetScreen logo, NetScreen-Global Pro, ScreenOS, and GigaScreen are registered trademarks of Juniper Networks, Inc. in the United States and other countries.

The following are trademarks of Juniper Networks, Inc.: Molecular Sequence Reduction, MSR, Network Sequence Mirroring, NSM, Packet Flow Acceleration, PFA, Application Flow Acceleration, AppFlow, Fast Connection Setup, Active Flow Pipelining, AFP, Policy-Based Multipath, PBM, Sequence Mirror, SR, Sequence Mirror, SM, Sequence Reduction System, SRS, Central Management System, CMS, and My Peribit. All other products and services are trademarks, registered trademarks, service marks or registered service marks of their respective owners. All specifications are subject to change without notice.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer or otherwise revise this publication without notice.

U.S. GOVERNMENT RIGHTS

Use, duplication, or disclosure by the U.S. Government of any of the programs included in this product shipment is subject to restrictions set forth in the Juniper Networks, Inc. SOFTWARE LICENSE AGREEMENT AND LIMITED WARRANTY and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DRAS 252.227-7013(c)(ii) (OCT 1988), FAR 12.212(a)(1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

This software product is the property of Juniper Networks, Inc. and its licensors, and is subject to the Clickwrap License Agreement accompanying this software product. By installing or using the software product, you agree to be bound to the terms of the Clickwrap License Agreement. You may not modify, translate, reverse engineer, decompile, disassemble or otherwise attempt to reconstruct or discover the source code of the software product. The Clickwrap License Agreement contains additional restrictions and disclaimers.

Any use, duplication, or disclosure by the U.S. government of any of the code included in this software product is subject to restrictions as set forth in the Clickwrap License Agreement accompanying this software product, and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (OCT 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable. Juniper Networks, Inc.

This product includes code licensed from RSA Security, Inc. Some portions licensed from IBM are available at <http://oss.software.ibm.com/icu4j/>.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org>). Copyright © 2000 The Apache Software Foundation. All rights reserved. A copy of the Apache Software License terms, restrictions and disclaimers is available at <http://www.apache.org/licenses/LICENSE>.

This product includes software developed by the ExoLab Project (<http://www.exolab.org>). Copyright © 2000-2002 Intalio Inc. All rights reserved. A copy of the license terms, restrictions and disclaimers for this software is available at <http://castor.exolab.org/license.html>.

Contents

Preface	11
Audience	11
Document Contents	11
Document Conventions	13
Typographical Conventions	13
Technical Support	13
Obtaining Additional Product Information	13
 Chapter 1 Introduction	 15
About CMS	15
What's New in Version 5.1	16
How CMS Works	17
Understanding CMS	18
CMS Support of Device Software Versions	18
Logging In to CMS	18
CMS Web Console Interface	19
Where to Go Next	20
 Chapter 2 Installing CMS	 21
System Requirements	21
Pre-Installation Tasks	22
Installing CMS	23
Uninstalling CMS	27
Logging In for the First Time	27
Recommended Configuration Tasks	30
Where to Go Next	31

Chapter 3 Managing Devices	33
Viewing Devices	33
Managing Devices	36
Viewing Device Events	37
Loading Device Boot Images	38
Rolling Back Device Boot Images	40
Rebooting Devices	42
Viewing Device Configuration Summaries	43
Analyzing Device Configurations	45
Loading Device Configurations	48
Rolling Back Device Configurations	51
Backing Up Device Configurations	52
Restoring Device Configurations	54
Retrieving Device Files	56
Applying a Registration Server Password	58
Putting Devices in Safe Mode	60
Accessing the SRS Web Console from CMS	61
Exporting Community and Device Information	61
Managing CMS Schedules	62
Managing Scheduled Tasks	62
Exporting a Schedule Log	66
Chapter 4 Managing Device Configurations	67
Overview of Device Configurations	67
Configuration Settings for SRS 5.1 and 5.0 Devices	68
Downloading Global and Partial Configurations	70
Consistency Checking	72
Tracking Configuration Versions	72
Tips for Managing Configurations	73
Viewing Configurations	73
Managing Configurations	75
Extracting Configurations	75
Duplicating Configurations	77
Creating New Configurations with Factory Defaults	78
Comparing Configurations	80
Displaying Configurations	81

Viewing Configuration History	82
Deleting Configurations	82
Defining Configuration Settings	83
Configuring Device Settings	87
Configuring Device Addresses	88
Defining Communities	90
Configuring Time Zone Settings	92
Configuring the ARP Table	93
Advertising Reduction Subnets	94
Defining Outbound QoS Exclusions	96
Adding Static Routes	97
Configuring Router Polling	99
Configuring Multi-Path Addresses	101
Configuring the RADIUS Source Address	103
Configuring Basic Setup Parameters	103
Configuring the Interface Settings	104
Configuring NTP	107
Enabling SNMP	108
Enabling Syslog Reporting	109
Configuring Dynamic Local Routes	110
Enabling Route-Based Router Balancing	112
Designating a Registration Server	114
Generating NetFlow Records	115
Configuring AAA Settings	116
Selecting Authentication Methods	117
Enabling Authorization Checking	120
Defining RADIUS Servers	121
Defining Local Users	123
Securing Operator Access	125
Securing Front Panel Access	126
Configuring Application Settings	127
Default Application Definitions for SRS 5.1 Devices	128
Viewing the Application Overview	130
Configuring Application Definitions	131
Testing New Application Definitions	136
Assigning Applications to Traffic Classes	136

Monitoring Applications	138
Configuring Reduction Settings	139
Configuring Endpoints for Reduction Tunnels	139
Configuring Network Sequence Mirroring	141
Reducing Applications	144
Configuring Remote Routes	146
Configuring Tunnel Load Balancing Policies	147
Configuring Default Assemblers	149
Defining Preferred Assemblers	151
Configuring Tunnel Mode Settings	153
Configuring QoS Settings	154
Using Outbound QoS to Enhance Performance	155
Understanding Outbound QoS	156
Traffic Classes and Bandwidths	157
QoS Templates and Endpoints	158
WAN Circuit Speeds and Router Overhead	159
Dedicated, Oversubscribed, and Variable Rate WANs	160
Direct Setup Versus Wizard Configuration Results	162
Class Priorities and Excess Bandwidth Allocation	164
ToS/DSCP Values	166
Unadvertised Subnets	166
Procedure for Configuring Outbound QoS Policies	166
Using the Outbound QoS Setup Wizard	168
Defining Outbound QoS Settings by Endpoint	176
Defining Outbound QoS Templates	179
Defining Outbound QoS Endpoints	180
Changing Outbound ToS/DSCP Values	184
Starting and Stopping Outbound QoS	187
Configuring Inbound QoS Policies	188
Configuring Traffic Acceleration	190
Overview of Packet Flow Acceleration	190
Active Flow Pipelining	191
Forward Error Correction	192
Fast Connection Setup	192
Overview of Application Flow Acceleration	193
Microsoft CIFS and Microsoft Exchange Acceleration	194

HTTP Acceleration	196
Enabling Acceleration by Endpoint	198
Enabling Acceleration by Application	203
Enabling Active Flow Pipelining by Application	203
Enabling Fast Connection Setup by Application	204
Enabling Microsoft CIFS Acceleration	205
Enabling Microsoft Exchange Acceleration	208
Enabling HTTP Acceleration	210
Configuring Advanced Setup Parameters	212
Configuring the Feature/Topology Settings	212
Configuring Source/Destination Filters	216
Defining the Prime Time	218
Configuring Packet Interception	220
Configuring Packet Interception for Off-Path Devices	221
RIP Router/Switch Configuration Commands	224
WCCP Router Configuration Commands	228
External Policy-Based Router Commands	229
Alternatives to Packet Interception	230
Configuring WAN Performance Monitoring	231
Adding CLI Commands to Configurations	234
Configuring Policy-Based Multi-Path	235
Procedure for Configuring Multi-Path	236
Enabling Policy-Based Multi-Path	237
Defining Multi-Path Templates	238
Defining Multi-Path Endpoints	240
Configuring Routers to Support Multi-Path	243
Configuring IPSec	245
Default IPSec Policy	245
IPSec Implementation Details	246
Procedure for Configuring IPSec Policies	247
Defining IPSec Settings by Endpoint	247
Defining IPSec Templates	249
Defining the Default IPSec Policy	251

Chapter 5 Automatic Deployment of Devices	253
About Automatic Deployment	253
Configuring Auto-Deployment	254
Auto-Deployment Procedure	254
Defining Deployment Groups	255
Defining Deployment Records	257
Viewing the Auto-Deployment Status	259
Configuring License Management	261
Licensing Procedure	261
Importing and Validating Authorization Codes	262
Generating and Applying Licenses	264
Viewing the License Status	266
 Chapter 6 Monitoring Performance	 269
Viewing and Printing Reports	269
Configuring the My Peribit Page	270
Viewing Reports on the Monitor Page	274
WAN Statistics	274
WAN Performance Statistics	275
WAN Throughput Statistics	280
WAN Application Summary	282
Reduction Statistics	283
Data Reduction Statistics	283
Application Summary Statistics	291
Passthrough Statistics	293
Packet Size Distribution Statistics	295
Monitoring Tunnel Status	296
QoS Statistics	300
Outbound QoS Statistics	300
Inbound QoS Statistics	304
Acceleration Statistics	308
Active Flow Pipelining Statistics	308
Fast Connection Setup Statistics	311
CIFS and Exchange Acceleration Statistics	313
HTTP Acceleration Statistics	315

Top Traffic Statistics	317
Executive Summary	321
Chapter 7 CMS Setup and Administration	325
Changing User Passwords	325
Viewing Logged In Users	326
Uploading a Boot Image	327
Administering Devices	328
Managing Communities	328
Managing Device Groups	331
Generating a Diagnostic File	333
Administering CMS	334
Defining CMS User Accounts	335
Defining User Groups	337
Controlling Client Device Access to CMS	339
Defining the Session Timeout	340
Configuring FTP Server Parameters	341
Enabling Syslog Reporting	342
Entering a Permanent License Key	343
Stopping and Starting the Scheduler	344
Changing the Web Server Port	345
Configuring Data Collection and Retention	346
Setting WAN Reporting Thresholds	348
Viewing the Polling Catch-Up and Failure Logs	349
Viewing System Logs	351
Backing Up and Restoring the Database	352
Manual Database Backups	352
Automatic Database Backups	353
Moving CMS to Another Disk Drive	354
Purging Temporary Java Files	355

Appendix A CMS Licenses	357
Appendix B Device Events	359
Appendix C Understanding Exported Data Results	363
NetFlow Version 5 Export.	363
Performance Statistics Export	365
General Device Information	366
Data Section Information	366
System Session Statistics	367
Reduction Session Statistics	370
Application Session Statistics	371
WAN Statistics	372
Application Flow Acceleration Statistics	372
Bandwidth Management Statistics	373
WAN Performance Statistics	374
Inbound Traffic By Port Statistics	375
Top Traffic Export.	376
Appendix D Common Application Port Numbers	377
Glossary	379
Index	383

Preface

Welcome to the Central Management System (CMS) — a powerful management and configuration tool for Sequence Reducer™ and Sequence Mirror™ devices. This section describes the audience, organization, and typographical conventions used in this manual.

Audience

This manual is intended for administrators who install and use CMS, and for network managers who monitor device performance. Readers are assumed to be familiar with their network architecture and devices, and can perform basic network configuration procedures.

Document Contents

- **Chapter 1, “Introduction”**

This chapter provides an overview of CMS, and describes the new features in this release.

- **Chapter 2, “Installing CMS”**

This chapter describes how to install the CMS software.

- **Chapter 3, “Managing Devices”**

This chapter describes how to centrally manage devices in a community by performing such tasks as loading new configurations and SRS™ boot images on selected devices. It also describes how to use the scheduler to manage scheduled tasks.

- **Chapter 4, “Managing Device Configurations”**

This chapter describes how to create and maintain global and partial configurations in CMS.

- **Chapter 5, “Automatic Deployment of Devices”**

This chapter describes how to configure new devices automatically, and how to distribute permanent licenses to devices that have evaluation licenses.

- **Chapter 6, “Monitoring Performance”**

This chapter describes how to monitor the percentage of data reduction, outbound bandwidth management by traffic class, and reduction tunnel status for the devices in each community.

- **Chapter 7, “CMS Setup and Administration”**

This chapter describes CMS administration tasks, such as importing communities and defining user accounts.

- **Appendix A, “CMS Licenses”**

This appendix describes the evaluation and permanent licenses for CMS.

- **Appendix B, “Device Events”**

This appendix describes the critical- and error-level Syslog messages generated by the devices and displayed in the Devices page as “events.” It also describes the appropriate action to take if a device encounters one of these events.

- **Appendix C, “Understanding Exported Data Results”**

This appendix describes the contents of the statistics file that CMS can retrieve from a device.

- **Appendix D, “Common Application Port Numbers”**

This appendix lists common application port numbers, as listed by the Internet Assigned Numbers Authority (IANA).

- **Glossary**

The glossary includes definitions of networking terms as well as terms specific to devices and CMS.

Document Conventions

This section describes conventions used throughout this manual.

Typographical Conventions

Table 1 lists the typographical conventions used throughout this manual.

Table 1 Typographical Conventions

Convention	Meaning	Example
boldface	Names of buttons or keys you should press.	Click Submit .
<code>courier font</code>	Text that you enter from the keyboard.	Enter the following command: <code>a:\setup</code>
Angle brackets	Variables that you must substitute another value for.	set ip <device's IP address>
<i>italics</i>	Names of manuals, directories, files, or Uniform Resource Locators (URLs).	The address of Juniper's web site is <i>http://www.juniper.net</i> .

Technical Support

Our commitment to create products and services that enable our customer's success is reflected in our Technical Assistance Center (TAC), and our comprehensive support programs.

For technical support, use the following methods:

- Go to *http://www.juniper.net/support*
- Send email to *support@juniper.net*.
- Call +1-888-314-JTAC (U.S, Canada, and Mexico) or +1-408-745-9500

Obtaining Additional Product Information

In addition to this *Central Management System Administrator's Guide*, refer to the *CMS 5.1 Release Notes* document enclosed with the product. Also refer to the *Sequence Reducer/Sequence Mirror Operator's Guide* and the *Quick Start* cards enclosed with each device.

For additional product information, visit our web site at *http://www.juniper.net/products/appaccel/wan*.

Chapter 1 Introduction

This chapter introduces the Central Management System (CMS) and covers the following topics:

- “About CMS” in the next section
- “What’s New in Version 5.1” on page 16
- “How CMS Works” on page 17
- “Understanding CMS” on page 18

About CMS

CMS provides easy and extensive central configuration and management for Sequence Reducer and Sequence Mirror devices in geographically dispersed locations. CMS can manage up to 2000 devices in multiple communities. CMS offers the following benefits:

- **Cost effective** — CMS reduces the cost of ownership for Sequence Reducer and Sequence Mirror devices by creating a single location from which to manage all devices and leverage configurations on devices already deployed in the network.
- **Eases configuration** — Using CMS, you can quickly and easily configure tens or hundreds of newly deployed devices, modify the configuration of already deployed devices, and view and manage the newly created WAN capacity generated by our Molecular Sequence Reduction (MSR)[™] and Network Sequence Mirroring (NSM)[™] technology.
- **Simplifies software deployment** — CMS dramatically simplifies the configuration and management of software upgrades. From a single location, and in a single operation, you can upgrade all devices in the same community to a new software version.
- **Creates global policies** — CMS allows network managers to centrally manage and modify global and device-specific configuration settings on all devices. Global settings include basic and advanced setup options, such as for NTP and SNMP, authentication settings, application definitions, outbound QoS settings, and the applications being reduced, monitored, and accelerated.
- **Schedules all tasks** — Using CMS, you can schedule all device management tasks to be performed at the optimal time for the individual location.

- **Centrally views all data reduction results** — CMS provides a single, clear window into the performance of devices around the globe. It presents historical per-tunnel and per-application data reduction statistics for each device.
- **Centrally views global device and tunnel status** — Using CMS, you can immediately view the status of each deployed device and all reduction tunnels.

All features are available through the CMS Web console, which is a Web-based graphical user interface. Up to 50 users can access the Web console simultaneously. You can control access to CMS with user accounts and passwords, as well as access control lists.

What's New in Version 5.1

CMS 5.1 has the following new features:

- Supports devices running SRS 5.x.
- All device configuration settings can now be controlled through CMS.
- Device groups let you create arbitrary groups of devices that are independent of communities. You can then grant access to Device groups let you limit access to specific devices in any community.
- User groups let you grant read or read/write access to selected communities and device groups for a specific set of users.
- The My Peribit page includes new charts for WAN loss, latency, and availability, measured by WAN performance monitoring or Policy-Based Multipath.
- New monitoring reports include:
 - Executive report summarizing reduction results, WAN application traffic, and WAN performance
 - CIFS, HTTP, and Microsoft Exchange application acceleration
 - Color-coded performance matrices for overall WAN performance, loss, latency, and availability
 - Throughput and Application Summary for traffic to and from the WAN
 - Application Summary and Passthrough reports for reduction results
- New catch-up polling mechanism can collect up to two weeks of past data when polling resumes for previously unreachable devices.

How CMS Works

CMS is deployed on a single Microsoft Windows Server 2000 or Windows Server 2003 server in your network (Figure 1-1). CMS includes a Web server that can be accessed by multiple remote Web consoles using secure Web access.

A CMS Web console is a workstation in your network that supports the Microsoft Internet Explorer 6.0 (or later) Web browser. You can access the Web by directing the browser to the IP address or host name of the CMS server (to use the host name, the host name must have a DNS entry.)

Figure 1-1 shows a logical flow of the communication between the Sequence Reducer and Sequence Mirror devices, a CMS server, and the CMS Web consoles. Configuration data between the devices and the CMS server is securely transmitted via a proprietary protocol. Monitoring data is collected from the devices in clear text (compressed). Data between the CMS server and the Web consoles is securely transmitted via HTTPS.

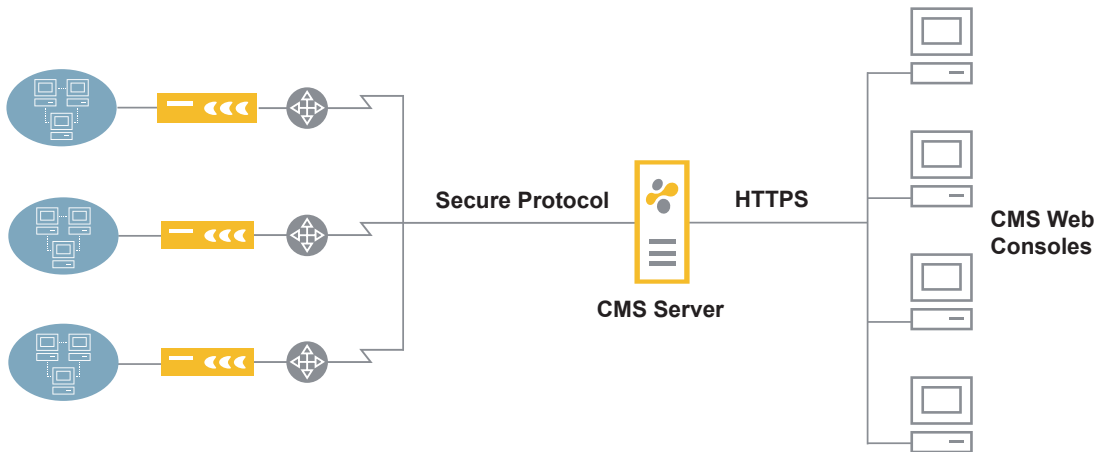


Figure 1-1 CMS Communication

Understanding CMS

The following sections provide general information about CMS.

- “CMS Support of Device Software Versions” in the next section.
- “Logging In to CMS” on page 18.
- “CMS Web Console Interface” on page 19.

CMS Support of Device Software Versions

CMS 5.1 manages devices running SRS version 5.0 and greater. Devices running SRS versions prior to 5.0 are displayed in some Web console pages (such as the Devices page), but they cannot be managed through CMS.

Logging In to CMS

When you log in to the CMS Web console for the first time, you must specify the user name “root” and a default password. The root user account provides access to all CMS functions. This level of access is known as Admin access. With Admin access, you can create up to 49 other user accounts and specify the level of access for each user, as described in “Defining CMS User Accounts” on page 335.


Up to 50 users can access the Web console at the same time. If two or more users modify the same settings concurrently, the last set of saved changes is used.

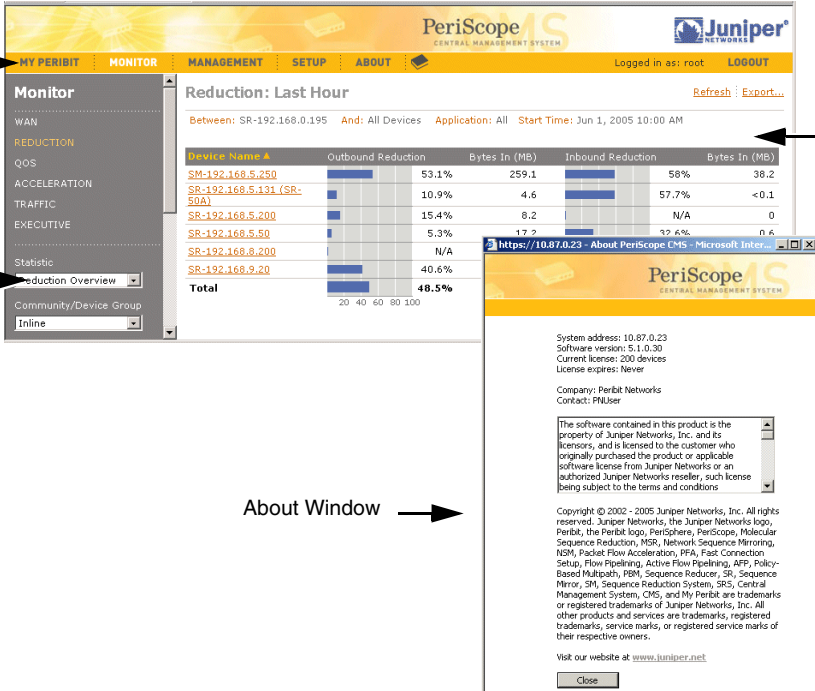
To log out of the Web console, click **LOGOUT** in the menu frame of any page. Users are logged out automatically if their sessions are inactive for the session timeout time (default is 30 minutes).

NOTE:

CMS Web Console Interface

The menu frame of the CMS Web console (Figure 1-2) identifies the user account used to log in and provides the following links:

- **MY PERIBIT** — Select and view a personalized set of performance charts specific to the user account.
- **MONITOR** — Monitor tunnel status and performance statistics.
- **MANAGEMENT** — Manage devices, configurations, automatic deployment, and scheduled tasks.
- **SETUP** — Administer CMS, such as add and delete user accounts, and import communities.
- **ABOUT** — View CMS server address, software version, and license information.
-  — Open a PDF version of this manual.
- **LOGOUT** — Log out of the CMS Web console.



The screenshot displays the CMS Web Console interface. The top navigation bar includes links for MY PERIBIT, MONITOR, MANAGEMENT, SETUP, and ABOUT. The user is logged in as 'root'. The main content area shows the 'Monitor' page with a table of performance statistics for the last hour. A left-hand navigation frame contains links for WAN, REDUCTION, QOS, ACCELERATION, TRAFFIC, EXECUTIVE, and a dropdown for Reduction Overview. An 'About' window is open in the foreground, displaying system information and a license agreement.

Menu Frame →

Left-hand navigation frame →

About Window →

Data Frame →

Device Name	Outbound Reduction	Bytes In (MB)	Inbound Reduction	Bytes In (MB)
SM-192.168.5.250	53.1%	259.1	58%	38.2
SR-192.168.5.131 (SR-50A)	10.9%	4.6	57.7%	<0.1
SR-192.168.5.200	15.4%	8.2	N/A	0
SR-192.168.5.50	5.3%	17.2	32.6%	0.6
SR-192.168.8.200	N/A			
SR-192.168.9.20	40.6%			
Total	48.5%			

System address: 10.87.0.23
 Software version: 5.1.0.30
 Current license: 200 devices
 License expires: Never
 Company: Peribit Networks
 Contact: PNUser

Copyright © 2002 - 2005 Juniper Networks, Inc. All rights reserved. Juniper Networks, Inc. and its licensors, and is licensed to the customer who originally purchased the product or applicable software license from Juniper Networks or an authorized Juniper Networks reseller, such license being subject to the terms and conditions

Visit our website at www.juniper.net

Close

Figure 1-2 CMS Web Console Interface

Where to Go Next

The left-hand navigation frame provides various sub-menu items, and the data frame displays the device monitoring and configuration data.

Where to Go Next

You are ready to install the CMS software. Proceed to “Installing CMS” on page 21.

Chapter 2 Installing CMS

This chapter describes the installation procedure for the Central Management System (CMS) and covers the following topics:

- “System Requirements” in the next section
- “Pre-Installation Tasks” on page 22
- “Installing CMS” on page 23
- “Logging In for the First Time” on page 27
- “Recommended Configuration Tasks” on page 30

System Requirements

Verify that the designated CMS server meets or exceeds the following hardware and software requirements:

NOTE: CMS should be installed on a dedicated server that is not used for any other applications. The installation optimizes some network parameters for CMS, which should not noticeably affect the system.

- Microsoft Windows Server 2003, or a Windows Server 2000 with Service Pack 3 or 4
- Table 2-1 shows the recommended and minimum CPU, memory, and disk space requirements for each range of devices being managed. These estimates assume a dedicated server with high speed drives, and a 30-minute polling interval.

Table 2-1 CPU, Memory, and Disk Space Requirements

Devices	Pentium 4 CPU (GHz)	RAM (GB)	Estimated Disk Space (GB)
Under 100	2.0+ (min. 1.8)	1.0 (min. 768 MB)	40+ (min. 40)
100 to 500	2.8+ (min. 2.0)	1.5 (min. 1.0)	60+ (min. 40)
500 to 1000	3.0+ (min. 2.8)	2.0 (min. 1.5)	80+ (min. 60)
1000 to 1500	3.2+ (min. 3.0)	3.0 (min. 2.0)	100+ (min. 80)
1500 to 2000	3.2+ dual CPU (min. 3.2)	4.0 (min. 3.0)	120+ (min. 100)

- CD-ROM drive
- Video display with 1024 x 768 resolution
- 10/100 Ethernet Network Interface Controller (NIC)
- A user account with administrator privileges (to perform the installation)
- Microsoft FTP Server installed and running, with an “anonymous” or password-protected user account that has read/write access to the FTP home directory

Pre-Installation Tasks

Complete all of the following pre-installation tasks:

- Verify that the TEMP environment variable for the system account is set to an NTFS drive with 100 MB of free disk space for the temporary files.

NOTE: An error occurs if the disk specified by TEMP has insufficient space, even if you install CMS on a separate disk with sufficient free space.

- Verify that the system date, time, and time zone are accurate for your location. In addition to the time zone setting in the Windows Date/Time properties dialog box, check the time zone environment variable. Refer to your Microsoft Windows documentation for more information.
- Determine if port 443 on the server is already used by IIS (the Windows Web server), or any other server. Port 443 is the default port used by the CMS Web server. If another server uses port 443, disable the server or specify port 8443 for the CMS Web server during installation. Port 443 or 8443 is required to support auto-deployment of devices.
- Verify that TCP port 443 (or 8443), and TCP and UDP ports 3577 and 3578 are not blocked by firewalls or other devices. CMS uses ports 3577 and 3578 to communicate with the devices.
- Determine if the Sun™ Microsystems™ Java™ Runtime Environment (JRE™), which is a component of the Java 2 Platform, Standard Edition (J2SE™), is on your system. If JRE version 1.5.0 is not present, the CMS installation wizard will install it.
- Reserve a static IP address for the CMS server.

- If the Microsoft FTP Server must be installed and running on the CMS server.

To install the FTP Server on Windows Server 2000:

- a. Click **Start > Settings > Control Panel**, and double-click **Add/Remove Programs**.
- b. Double-click **Add/Remove Windows Components**.
- c. Select **Internet Information Services (IIS)**, and click **Details**.
- d. In the IIS window, select the check box for **File Transfer Protocol Server** and click **OK**.
- e. Click **Next** to install the service. When prompted, insert the Microsoft Windows 2000 Server CD into the CD drive.

To install the FTP Server on Windows Server 2003:

- a. Install the FTP service as described at <http://support.microsoft.com/?kbid=323384>.
- b. Enable write permission for the FTP service, as described at <http://support.microsoft.com/default.aspx?scid=kb;en-us;309007&sd=tech>.

Installing CMS

To install the CMS software on your system:

1. Log in to the Microsoft Windows 2000 or 2003 server as a user with administrator privileges. Next, close all windows and exit all programs, including any anti-virus programs running on the desktop.
2. If you are upgrading from CMS 5.0 to CMS 5.1 (the upgrade will take several minutes and cannot be reversed):

NOTE: All SRS 4.0 configurations are deleted. Devices running SRS versions prior to 5.0 are listed on some CMS pages, such as the Devices page, but they cannot be managed through CMS.

- a. Stop the CMS service on the server:
 - a. Click **Start > Run**, enter “services.msc” and click **OK**.
 - b. In the Services window, right-click on **PeriScopeCMS** and click **Stop**.
- b. Back up the 5.0 database, as described in “Backing Up and Restoring the Database” on page 352.
3. Insert the CMS CD into the server’s CD drive.

After installation files are extracted, a welcome window for the installation wizard is displayed. If the welcome window does not appear, you can access the installation program on the CD.

NOTE: Click **OK** for all security or AntiSpyware prompts encountered during the installation and initial setup of CMS.

4. Click **Next**. The CMS license agreement appears. Read the agreement carefully. To accept the terms of the agreement, click **Yes**. The Customer Information window opens (Figure 2-1).

Figure 2-1 Entering Customer Information

5. Enter customer information:
 - a. Enter a user and company name if the fields are not already filled in.

- b. If you have a permanent license key, enter it in the **License Key** field. If you do not have a permanent license key, leave “Evaluation” in the **License Key** field. For more information about licenses, refer to “CMS Licenses” on page 357.
- c. Click **Next**. The Choose Install Type window opens (Figure 2-2).

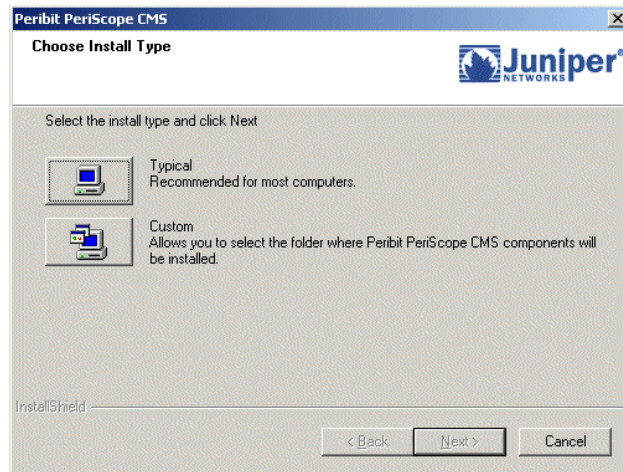


Figure 2-2 Selecting the Installation Type

6. Select a Typical or Custom installation, as follows:
 - Click **Typical** to do the following:
 - Install the CMS files in *C:\Program Files\Peribit\CMS*.
 - Install JRE version 1.5.0 in *C:\Program Files\Java\j2rel.5.0_01* if it is not already installed on your system.
 - Set the Web server port to 443 (the default HTTPS port). If port 443 is currently used by IIS or some other Web server, you are prompted to enter another port number (enter port 8443).
 - To change any of the default settings, click **Custom**: to open the Custom Settings window (Figure 2-3).

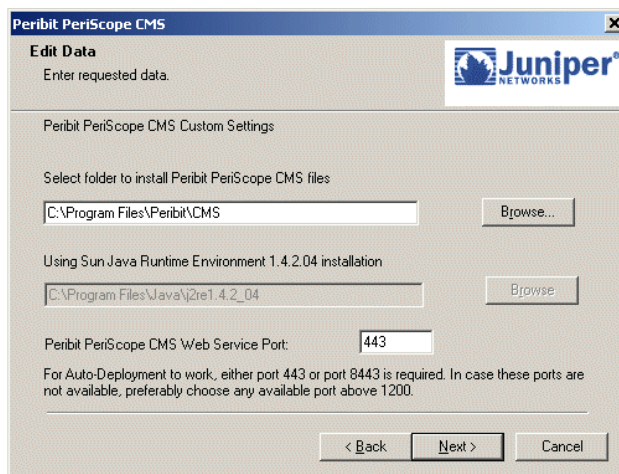


Figure 2-3 Customizing CMS Installation

- a. To change the locations of the CMS files, click **Browse** and use the Windows Explorer to navigate to the desired locations.
 - b. If the default Web server port number (443) is already in use, enter port number 8443. If 8443 is also in use, specify a port number above 1200. Note that you cannot auto-deploy devices unless the port number is 443 or 8443.
 - c. Click **Next**.
7. To change any of the previous settings, click **Back**. If you are satisfied with the settings, click **Next** to start the installation.

If TCP or UDP port number 3577 is in use, you are prompted to enter another port number (CMS listens on this port to collect performance data from the devices).

8. When the installation is complete, a window displays the URL to use to access CMS. Click **Finish**. The restart window is displayed.

You must restart the system to activate CMS. Before restarting the system, remove any disks or CDs from the drives.

9. To restart the system, select **Yes** and then click **Finish**.

Uninstalling CMS

To uninstall CMS, use the Microsoft Windows Add/Remove Programs function in the Control Panel. The uninstall wizard allows you to delete the CMS data and configuration folders, which include all files related to CMS, including the license, communities, users, and passwords. If you are removing CMS from your system, you can safely delete these files.

If the JRE was installed by the installation wizard, the uninstall wizard also lets you delete it from the system.

Logging In for the First Time

After installing CMS, you must log into the Web console and perform some basic administration.

You can log into the CMS Web console from any workstation in your network. The Web console supports Microsoft Internet Explorer version 5.5 and greater. Data is securely transmitted from the CMS server to the Web browser via HTTPS.

To log in to the CMS Web console:

1. From a workstation in your network, start your Web browser and enter the following URL:

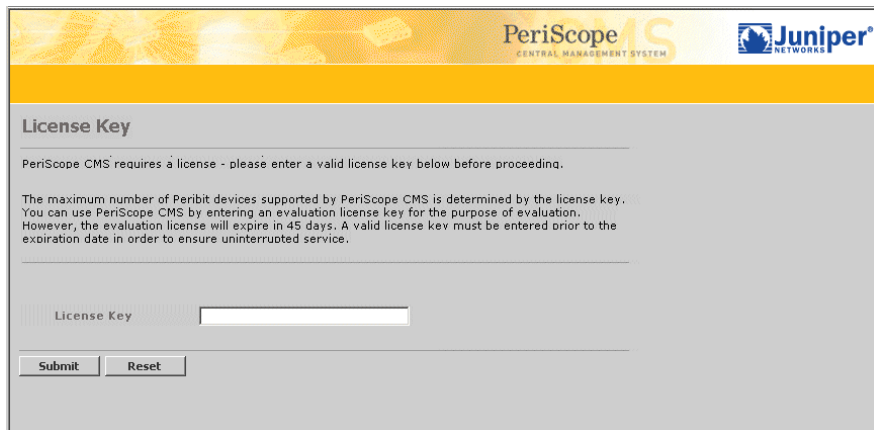
`https://<IP address of the CMS server>:<port number>`

Be sure to use “https” instead of “http”. Also, if you have changed the CMS Web server port number from 443 to 8443, you must include “:8443” after the IP address. For example:

`https://10.10.0.1:8443`

If you have not changed the Web server port number from the default of 443, you can omit the colon and port number after the IP address.

2. If you did not enter a license key during installation, the License Key page opens (Figure 2-4).

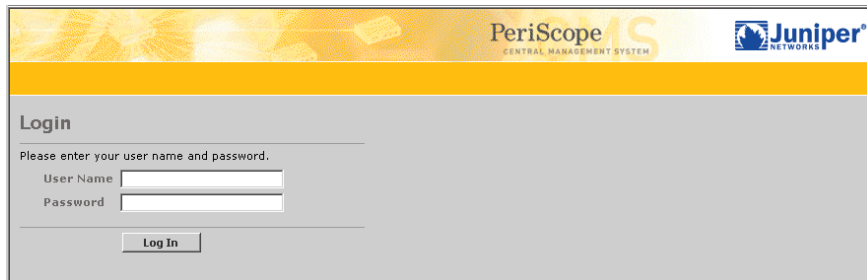


The screenshot shows the 'License Key' page of the PeriScope CMS. The page has a yellow header with the 'PeriScope CMS' logo and the Juniper Networks logo. Below the header, the title 'License Key' is displayed. The main content area contains a message: 'PeriScope CMS requires a license - please enter a valid license key below before proceeding.' followed by a detailed explanation of the evaluation license. At the bottom, there is a text input field labeled 'License Key' and two buttons: 'Submit' and 'Reset'.

Figure 2-4 Entering the License Key

To obtain a license, go to https://www.juniper.net/generate_license. For an evaluation license (good for 10 devices and 45 days), click **CMS 5.1 Evaluation**, enter the IP address of the CMS server and other information, and click **Continue**. You can then copy the generated key into the **License Key** field here, and click **Submit**.

3. Depending on your Web browser settings, the Security Alert dialog box may appear. Click **Yes** to open the Login page (Figure 2-5).



The screenshot shows the 'Login' page of the PeriScope CMS. The page has a yellow header with the 'PeriScope CMS' logo and the Juniper Networks logo. Below the header, the title 'Login' is displayed. The main content area contains a message: 'Please enter your user name and password.' followed by two text input fields: 'User Name' and 'Password'. At the bottom, there is a 'Log In' button.

Figure 2-5 Logging In For The First Time

4. Enter the following user name and password, and click **Log In**.
 - User name: **root**
 - Password: **peribit**

The Change CMS Administrator Password page opens (Figure 2-6).

Figure 2-6 Changing the Default Password

5. Enter a new password for the root user in the **New Password** and **Verify New Password** fields, and click **Submit**. The password is case-sensitive.
6. A blank Communities page opens. To manage the devices in each community, you must import the communities defined on each device that acts as a registration server. For more information about registration servers, refer to “Designating a Registration Server” on page 114. If you do not yet have a registration server, you can import communities at a later time (refer to “Managing Communities” on page 328).

To import communities into CMS:

- a. Click **Import** to open the Communities > Import page (Figure 2-7).

Figure 2-7 Importing Communities to CMS

- b. In the Communities > Import page, enter the IP address and password of a device that acts as a primary registration server, and click **Submit**.

- c. Select the check box next to each community, click **Import**, and click **OK**.

Note that the Default community on each registration server becomes “Default - <IP address>” in CMS.

The CMS quick setup is complete. You are now ready to perform additional administrative tasks. For more information, see “Recommended Configuration Tasks” in the next section.

Recommended Configuration Tasks

Now that CMS is initially configured, you should perform the following tasks (the first item is required):

- Create at least one user group, as described in “Defining User Groups” on page 337.
- The clocks on all devices, including the CMS server, should be synchronized to the same Simple Network Time Protocol (SNTP) server or server hierarchy. To use CMS to configure an NTP server for your devices, refer to “Configuring NTP” on page 107.

If you do not have an SNTP server in your network, you can use the address of the CMS server. During installation, CMS enables the Windows SNTP server. Be sure to verify that port 123 (UDP) is not blocked by firewalls or other devices.

If you use some other SNTP server to synchronize your devices, the Windows SNTP agent on the CMS server should be pointed to the same SNTP server.

- If you have multiple registration servers, import the communities from each server, as described in “Managing Communities” on page 328.
- Upload SRS boot images to the CMS server, as described in “Uploading a Boot Image” on page 327. An uploaded image can then be downloaded to selected devices.
- Use CMS to retrieve and analyze the differences between the configurations of selected devices. Extracted configurations can be used as a starting point in managing your device configurations. For more information about analyzing a configuration, see “Analyzing Device Configurations” on page 45. For more information about extracting a configuration, see “Extracting Configurations” on page 75.

Analyzing configurations also helps you select a global configuration that can be modified and loaded on other devices. For more information about modifying and loading a configuration, see “Defining Configuration Settings” on page 83 and “Loading Device Configurations” on page 48.

Where to Go Next

To view the devices that CMS discovers for the Community, proceed to Chapter 3, “Managing Devices”. To create user accounts or perform additional administrative functions, proceed to Chapter 7, “CMS Setup and Administration”.

Where to Go Next

Chapter 3 Managing Devices

This chapter describes how to use CMS to centrally manage communities of Sequence Reducer and Sequence Mirror devices. It covers the following topics:

- “Viewing Devices” in the next section
- “Managing Devices” on page 36
- “Accessing the SRS Web Console from CMS” on page 61
- “Exporting Community and Device Information” on page 61
- “Managing CMS Schedules” on page 62

Viewing Devices

The Devices page lets you view the devices in each community or device group, execute tasks for selected devices, and open the SRS Web console for a specific device.

To view the devices in each community or device group:

1. Click **MANAGEMENT** in the menu frame.

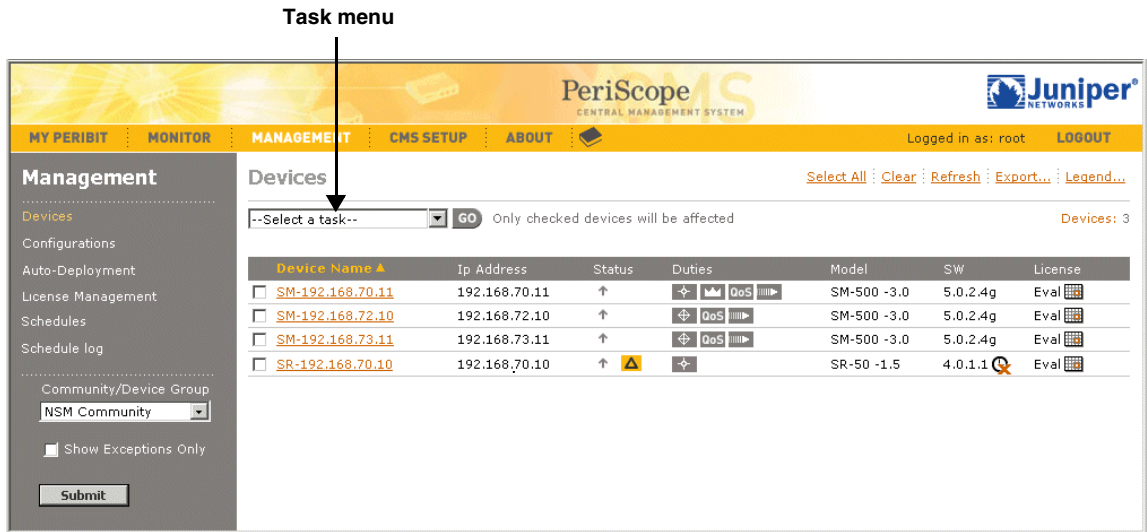




Figure 3-1 Devices Page

2. Select a community or device group from the **Community/Device Group** menu, and click **Submit**. If the device polling takes too long, you can click **Stop** in the upper-right corner of the page. Polling continues, but you can then execute tasks for one or more devices, such as loading a new boot image, as described in “Managing Devices” on page 36.

From the Devices page, you can also:

- View the status, hardware model, software version, license, and functional properties of each device. An asterisk (*) next to the software version indicates that the device did not respond to the last query.
- Click **Legend** to view a brief description of the icons used on the Devices page (refer to Table 3-1 for more information).
- Click **Refresh** to view the latest device status information. Click the column headers to change the sort.
- Click **Show Exceptions Only** to view only devices that are not responding or that have:
 - Events that occurred in the last 24 hours
 - Tasks that failed
 - An expired license or a registration server password different from CMS
- Click  in the **Status** column to view device events, as described in “Viewing Device Events” on page 37.
- Open the SRS Web console for a device by clicking the device name, as described in “Accessing the SRS Web Console from CMS” on page 61.
- Click **Export** to export the device information to a CSV file, as described in “Exporting Community and Device Information” on page 61.
- Click  in the **SW** column to view the details of failed tasks, as described in “Managing CMS Schedules” on page 62.

NOTE: Devices with SRS versions prior to 5.0 are displayed without a check box to indicate that they cannot be managed with CMS.

Table 3-1 Device Icons




























Icon	Table Column	Description
	Status	The device is up and running.
	Status	The device is operating in safe mode.
	Status	Registration server password on the device does not match the password in CMS.
	Status	One or more critical- or error-level events occurred in the past 24 hours. Move the cursor over the icon to view the number of events. Click the icon to view the event details.
	Status	Storage space on the device is low. Move the cursor over the icon to view the number of bytes remaining.
	Status	The device is not reachable or has never responded. An asterisk (*) after the software version indicates no response to the last query.
	Status	The device is NOT using an NTP server to maintain an accurate device time. An NTP server is recommended to ensure the accuracy of hourly reports (refer to “Configuring NTP” on page 107).
	Duties	The device is a hub in a Hub and Spoke topology.
	Duties	The device is a spoke in a Hub and Spoke topology. By default, a spoke reduces and assembles data only for the hub devices.
	Duties	The device is part of a mesh topology.
	Duties	The device is the primary registration server.
	Duties	The device is the secondary registration server.
	Duties	The device is a backup for one or more devices. The icon flashes when the backup device is active.
	Duties	The device is part of a multi-node configuration.
	Duties	Outbound Quality of Service (bandwidth management) is enabled.
	Duties	Packet Flow Acceleration is enabled.
	Duties	Network Sequence Mirroring (NSM) is enabled.
	Duties	Policy-Based Multi-Path is enabled.
	Duties	IPSec encryption is enabled.

Table 3-1 Device Icons (Continued)

Icon	Table Column	Description
	Duties	Packet interception using RIP is enabled.
	Duties	Packet interception using WCCP is enabled.
	Duties	Packet interception using external routing is enabled.
	Device	Indicates that you cannot change the device's boot image or configuration
	SW	One or more scheduled tasks is pending. Click the icon to view more information about the scheduled tasks
	SW	One or more scheduled tasks has failed. Click the icon to view more information about the failed tasks.
	License	License key has an expiration date. Move the cursor over the icon to view the number of days remaining.
	License	Licensed throughput is exceeded.

Managing Devices


The following sections describe the device management tasks:

- “Viewing Device Events” in the next section
- “Loading Device Boot Images” on page 38.
- “Rolling Back Device Boot Images” on page 40.
- “Rebooting Devices” on page 42.
- “Viewing Device Configuration Summaries” on page 43.
- “Analyzing Device Configurations” on page 45.
- “Loading Device Configurations” on page 48.
- “Rolling Back Device Configurations” on page 51.
- “Backing Up Device Configurations” on page 52
- “Restoring Device Configurations” on page 54
- “Retrieving Device Files” on page 56.

- “Applying a Registration Server Password” on page 58.
- “Putting Devices in Safe Mode” on page 60.

Viewing Device Events

To view the details of device events:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Click  in the Status column to open the Events window (Figure 3-2). The event icon is displayed only if one or more critical- or error-level events occurred on the device in the past 24 hours.

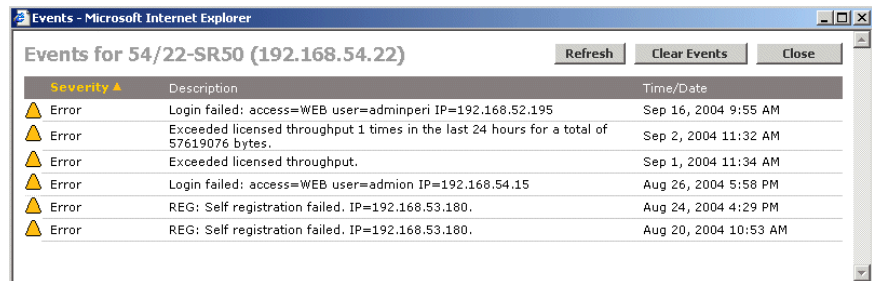


Figure 3-2 Viewing Device Events

The Events window displays the severity level, description, and date and time of the last 20 events for the device. To update the list, click **Refresh**.

3. To delete all events and remove the event icon from the Devices page, click **Clear Events**. This also clears the events on the device.
4. To close the Events window, click **Close**.

For a list of the possible events and recommended actions, refer to Appendix B, “Device Events.”

Loading Device Boot Images

After you load an SRS boot image on the CMS server, you can globally distribute the image to selected devices in a community or device group. To upload a boot image, refer to “Uploading a Boot Image” on page 327.

Loading a boot image on a device does not affect the device configurations. All configuration information is preserved. It is strongly recommended that all devices in the same community have the same boot image.

Loading a boot image involves two tasks:

- Load the boot image from CMS to selected devices.
- Reboot the devices to activate the new boot image. The reboot can be done automatically or scheduled as a separate task.

When loading a boot image to multiple devices, you should schedule the reboot separately after verifying that the boot image was loaded successfully on each device. This lets you activate the new boot image on all devices simultaneously. To verify a task was successful, refer to “Managing Scheduled Tasks” on page 62.

If you have any problems after upgrading to a new boot image, you can roll the boot image back to the previous version, as described in “Rolling Back Device Boot Images” in the next section.

CAUTION: You can downgrade devices to a previous version of SRS. However, avoid downgrading whenever possible. Downgrading may cause unpredictable behavior because the configuration and other data files were created with the later release.

Monitor downgraded devices carefully. If problems occur, restore or roll back the configuration to the one used with the older boot image, if possible (refer to “Rolling Back Device Configurations” on page 51 and “Restoring Device Configurations” on page 54).


To load a boot image on selected devices:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices where you want to load the boot image or click **Select All**.
3. From the Task menu, select **Image > Load** and click **Go**.

The screenshot shows the 'Periscope CMS' interface with a navigation bar at the top containing 'MY PERIBIT', 'MONITOR', 'MANAGEMENT', 'CMS SETUP', and 'ABOUT'. The 'MANAGEMENT' tab is active. On the left, a 'Management' sidebar lists options: 'Devices' (highlighted), 'Configurations', 'Auto-Deployment', 'License Management', 'Schedules', and 'Schedule log'. Below these is a 'Community/Device Group' dropdown set to 'NSM Community' and a 'Show Exceptions Only' checkbox. A 'Submit' button is at the bottom of the sidebar. The main content area is titled 'Devices > Load image'. It contains a 'Load the following boot image to the selected device(s). [Show selected devices](#)' instruction. A 'Boot image' dropdown menu is set to '.....'. The 'Schedule' section has two radio buttons: 'Load now' (selected) and 'Delay loading until:'. The 'Delay loading until' section includes 'Time' (HH:MM) and 'Date' (MM/DD/YYYY) fields, with 'AM' and 'PM' radio buttons. The 'Reboot' section has a checkbox 'Reboot device(s) after loading boot image'. The 'Downgrade' section has a checkbox 'Allow image downgrade' and a detailed warning text: 'If you enable image downgrade, the selected image will be loaded on the selected device(s) even if its version is older than the image currently running on the device. You should enable this option only after careful consideration of the consequences, especially if downgrading to an older major version of the image. The configuration and other data files created by the current image running on the device may cause unpredictable behavior when the device is restarted with the older image. This may include the device not being able to become fully operational. If you enable this option, please monitor the device (s) carefully to ensure they are operating normally.' At the bottom are 'Submit' and 'Cancel' buttons.

Figure 3-3 Loading a Boot Image on Devices

4. Specify the following information:

- | | |
|------------|---|
| Boot image | <p>Select the SRS boot image you want to load. The default naming convention of boot images is:</p> <p style="text-align: center;">srs<rdm><bb>.<zip or bin></p> <p>where:</p> <p><r> is the major release number.</p> <p><d> is the minor release number.</p> <p><m> is the maintenance release number.</p> <p><bb> is the build number.</p> <p>The file extension must be “zip” or “bin”.</p> |
| Schedule | <p>Select Load now or select Delay loading until and enter a future time and date:</p> <ul style="list-style-type: none"> • Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. • Enter a date in MM/DD/YYYY format or click  and select the month and date. |

Reboot	<p>Click the check box to reboot the device after the image is loaded. The loaded image is not activated until the device is rebooted.</p> <p>To schedule the reboot as a separate task, which is recommended when updating multiple devices, refer to “Rebooting Devices” on page 42.</p> <p>NOTE: When you reboot a device, all unsaved configuration data is lost.</p>
Downgrade	<p>Click the check box if the selected boot image is older than the current version.</p> <p>CAUTION: Downgrading to a previous boot image may cause unpredictable behavior and should be avoided whenever possible.</p>

5. To review the devices you selected, click **Show selected devices**.
6. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

If problems occur after upgrading to a new boot image, you can roll the boot image back to the previous version, as described in “Rolling Back Device Boot Images” in the next section.

Rolling Back Device Boot Images

When you load a SRS boot image from CMS, each device retains the previous boot image. If problems occur with the new image, you can roll back to the previous version. During a rollback, each device reverts to the previous image and deletes the current image.

NOTE: You can roll back the boot image on a device only if you loaded the boot image from CMS.

Rolling back the boot image does not affect the device configurations. All configuration information is preserved. It is strongly recommended that all devices in the same community have the same boot image.

Rolling back a boot image involves two tasks:

- CMS directs the specified devices to roll back to the previous image.
- Reboot the devices to activate the rolled back boot image. The reboot can be done automatically or scheduled as a separate task.

When rolling back the boot image on multiple devices, you should schedule the reboot separately after verifying that the rollback was successful on each device. This lets you activate the boot image on all devices simultaneously. To verify a task was successful, refer to “Managing Scheduled Tasks” on page 62.

To roll back the boot image on selected devices:


1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices where you want to roll back the boot image or click **Select All**.
3. From the Task menu, select **Image > Rollback** and click **Go**.

Figure 3-4 Rolling Back the Boot Image

4. Specify the following information:

Schedule

Select **Roll back now** or select **Delay roll back until** and enter a future time and date:

- Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click **AM** or **PM**. Note that midnight can be 0:0 AM or 12:00 AM.
- Enter a date in MM/DD/YYYY format or click  and select the month and date.

Reboot Click the check box to reboot the device after the rolled back image is loaded. The loaded image is not activated until the device is rebooted.

To schedule the reboot as a separate task, which is recommended when updating multiple devices, refer to “Rebooting Devices” on page 42.

NOTE: When you reboot a device, all unsaved configuration data is lost.

5. To review the devices you selected, click **Show selected devices**.
6. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

Rebooting Devices

You must reboot a device to activate a loaded or rolled back boot image or to reactivate data reduction on a device that is in safe mode.

NOTE: When you reboot a device, all unsaved configuration data is lost.

To reboot selected devices:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices that you want to reboot or click **Select All**.
3. From the Task menu, select **Reboot** and click **Go**.

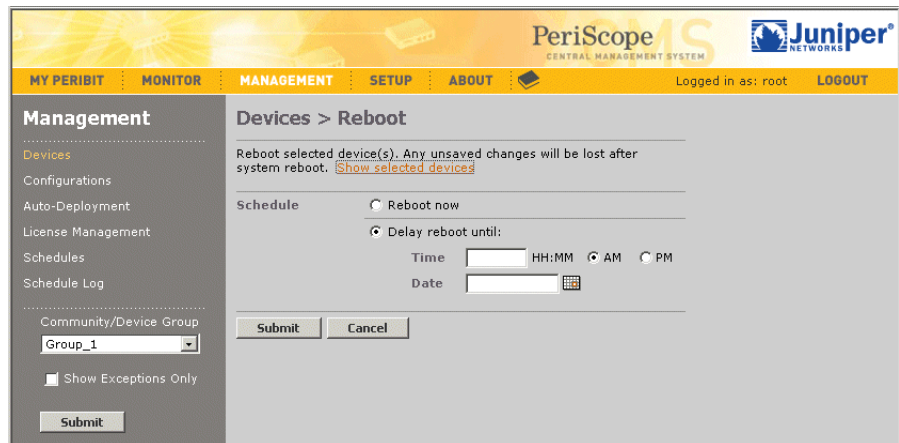



Figure 3-5 Rebooting a Device

4. Select **Reboot now** or select **Delay reboot until** and enter a time and date:

- Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click **AM** or **PM**. Note that midnight can be 0:0 AM or 12:00 AM.
- Enter a date in MM/DD/YYYY format or click  and select the month and date

5. To review the devices you selected, click **Show selected devices**.

6. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

Viewing Device Configuration Summaries

If you have loaded configurations to one or more devices from CMS, you can view a summary of the last set of configurations loaded on each device. You can also verify whether a device has any unsaved settings that can be defined in a global configuration. Unsaved settings are lost when you load a global configuration.

To view configuration summaries:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices for which you want to view a summary or click **Select All**.

3. From the Task menu, select **Configuration > Summary** and click **Go**.

Device Name	IP Address	Config Name	Config Type	Created	Applied	
denver	10.10.190.25	config03	Global	10/24/03 11:57 AM	10/24/03 6:23 PM	VERIFY
		app-28	Applications	10/24/03 3:47 PM	10/24/03 6:23 PM	
		red_west_07	Reduction	10/24/03 3:14 PM	10/24/03 6:23 PM	
long_beach	10.10.191.25	config03	Global	09/01/02 11:57 AM	09/04/02 2:34 AM	VERIFY
los_angeles	10.10.192.25	N/A	N/A	N/A	N/A	
monterey	10.10.193.25	config03	Global	09/01/02 11:57 AM	09/04/02 2:34 AM	VERIFY
oakland	222.222.222.222	config03	Global	09/01/02 11:57 AM	09/04/02 2:34 AM	VERIFY

Figure 3-6 Configuration Summary

The following information is displayed for each selected device:

- Name and type of the last global and partial configurations downloaded from CMS.
- Date and time the “load configuration” task was created and submitted to the scheduler.
- Date and time the configuration was applied to the device. The created and applied times are different if the load task was scheduled for a future time.

If no configurations have been loaded from CMS for a device, **N/A** is displayed for the above fields.

4. Click **Verify** for a device to check for differences between the saved and running configurations. All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.

The Verify windows shows device-specific settings in bold italics. Color-coded lines indicate the following:

- **Blue**. Settings unique to the saved configuration in the left column.
- **Yellow**. Settings unique to the running configuration in the right column.
- **Pink**. Settings that are different between the two configurations.

When you are done viewing the configuration, click **Close**.

To capture unsaved settings, you can extract the running configuration, as described in “Extracting Configurations” on page 75. Alternatively, you can incorporate the unsaved changes in an existing configuration (refer to “Defining Configuration Settings” on page 83).

Analyzing Device Configurations

CMS lets you analyze the differences between the configurations on two or more devices in a community. This is particularly useful if the devices were installed and configured without CMS. Based on the analysis, you can eliminate unnecessary differences between devices, and extract global or partial configurations that you can load on other devices.

To analyze configurations:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices you want to analyze or click **Select All**.
3. From the Task menu, select **Configuration > Analyze** and click **Go**.

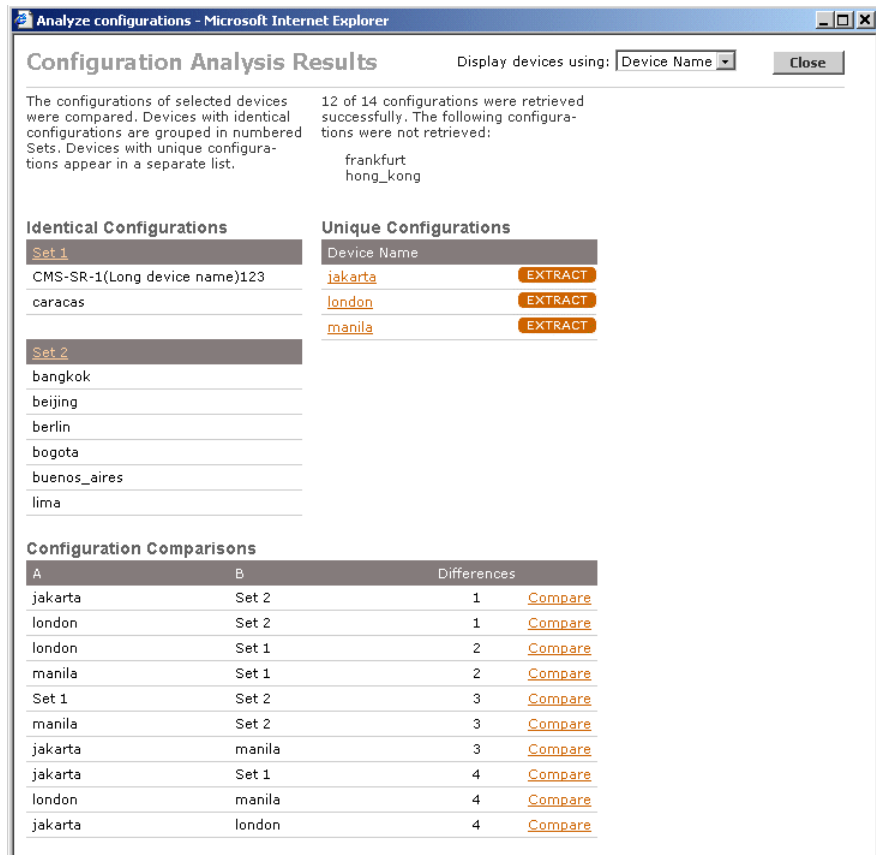


Figure 3-7 Analyzing Configurations

The Configuration Analysis Results window includes the following:

- Devices from which a configuration could not be retrieved
- Sets of devices that have identical configurations
- Devices that have unique configurations
- Comparisons of each unique pair of configurations and the number of differences between them.

NOTE: The number of differences indicates the number of different blocks of settings, not the number of different lines.

4. To view devices by IP address, rather than by name, select **IP Address** from the menu at the top of the window.

5. To view or compare the settings that can be defined in a global configuration in CMS, do one of the following. All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.
 - a. Click **Set** <number> to view of one of the identical configuration sets.
 - b. Click the device name or IP address to view a unique configuration.
 - c. Click **Compare** next to the two configurations you want to compare. A line-by-line comparison of the settings that can be defined in a global configuration is displayed. Color-coded lines indicate the following:
 - **Blue**. Settings unique to the configuration in the left column.
 - **Yellow**. Settings unique to the configuration in the right column.
 - **Pink**. Settings that are different between the two configurations.

When you are done viewing the configurations, click **Close**.

6. To create a global configuration from a device's running configuration, click **EXTRACT** next to the device, enter a configuration name and description, and click **Submit**. Only the settings that can be defined in a global configuration in CMS are extracted.

The extracted configuration is added to the Configurations page. You can then edit the configuration and load it on selected devices, as described in “Defining Configuration Settings” on page 83, and “Loading Device Configurations” on page 48.

7. When you are done viewing the Configuration Analysis Results window, click **Close**.

Loading Device Configurations

CMS lets you load a configuration on selected devices in a community or device group. A loaded configuration can consist of a global configuration and/or one or more partial configurations. The configuration changes take effect immediately.

Downloaded configurations are processed as follows:

- **Global configuration OR partial configurations.** The selected settings in the global or partial configurations override the corresponding settings on each device.
- **Global AND partial configurations.** The selected settings in the partial configurations override the settings in the global configuration, and the result overrides the corresponding settings on each device.

Any settings not defined in the downloaded configurations are left unchanged on each device, provided the settings were saved in the startup configuration file. Device settings are replaced, not supplemented. For example, if a device has four traffic classes, and you load a configuration that has two classes, the resulting device configuration will have two traffic classes, not six.

Settings that can be specified only by CLI commands are retained on each device unless they are overridden by commands in the CLI section of a global or partial configuration.

You can preview the results for each device before submitting the task. For more information about global and partial configurations, refer to “Overview of Device Configurations” on page 67.

To load a configuration:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices where you want to load a configuration or click **Select All**.
3. If you have previously loaded configurations from CMS, check each device for unsaved configuration settings (refer to “Viewing Device Configuration Summaries” on page 43). Unsaved settings are lost when a new configuration is loaded.


NOTE: Do not select devices running different versions of SRS. The SRS 5.0 and 5.1 configuration files are not compatible.

4. From the Task menu, select **Configuration > Load** and click **Go**.

Figure 3-8 Loading a Configuration

5. Specify the following information:

- | | |
|----------------------|--|
| Global Configuration | To load a global configuration, select a global configuration from the menu. Click History to view the selected configuration and its history of changes. To create global configurations, refer to “Managing Configurations” on page 75. |
| Partials | To load only partial configurations, select Do not load global configuration . To load partial configurations, click Yes and select up to one of each type of partial configuration. The settings in each partial configuration replace the corresponding settings in the selected global configuration (if any). Click History to view each selected configuration and its history of changes. |

Schedule	<p>Select Load now or select Delay loading until and enter a future time and date:</p> <ul style="list-style-type: none"> • Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. • Enter a date in MM/DD/YYYY format or click  and select the month and date. • To load the configuration periodically, select a recurring retrieval interval (Daily, Weekly, or Monthly). Monthly schedules that start on the 31st are permanently reset to the last day of the shortest month encountered.
Reboot	<p>Click the check box to reboot the device after the configuration is loaded. The new configuration takes effect with or without a reboot. However, a reboot is recommended when you make substantial changes to a configuration or when you change the topology parameter.</p> <p>NOTE: When you reboot a device, all unsaved configuration data is lost.</p>

6. To review the devices you selected, click **Show selected devices**.
7. Click **Preview** to view the resulting configuration for the first device. The blue text indicates the configuration settings that must be specified by CLI commands, a Device Settings partial configuration, or the SRS Web console. Click **Next** to preview the configuration for each selected device.

To open the SRS Web console, refer to “Accessing the SRS Web Console from CMS” on page 61. To change a global or partial configuration, refer to “Defining Configuration Settings” on page 83.

All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator’s Guide*.

8. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

Rolling Back Device Configurations

When you load a configuration on one or more devices from CMS, each device's previous configuration is retained in CMS. If problems occur with the new configuration, you can roll back to the previous version.

NOTE: You can roll back the configuration on a device only once, and only if you have loaded the current configuration from CMS.

To roll back the configuration on selected devices:


1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices where you want to roll back a configuration or click **Select All**.
3. From the Task menu, select **Configuration > Rollback** and click **Go**.

The screenshot displays the PeriScope CMS interface. The top navigation bar includes links for MY PERIBIT, MONITOR, MANAGEMENT, CMS SETUP, and ABOUT. The left sidebar under 'Management' lists various options, with 'Devices' highlighted. The main content area is titled 'Devices > Roll back configuration'. It contains a message stating that the configuration of selected devices will be rolled back to the state that existed prior to the most recent configuration download, with a link to 'Show selected devices'. Below this, there are two radio buttons for 'Schedule': 'Roll back now' and 'Delay roll back until'. The 'Delay roll back until' option is selected, and it includes fields for 'Time' (HH:MM) and 'Date'. There is also a checkbox for 'Reboot' labeled 'Reboot device(s) after rolling back configuration'. At the bottom, there are three buttons: 'Submit', 'Preview...', and 'Cancel'. The 'Community/Device Group' dropdown menu is set to 'community 1', and the 'Show Exceptions Only' checkbox is unchecked.

Figure 3-9 Rolling Back the Configuration

4. To view the rollback configuration (if any), click **Preview**.

5. Specify the following information:

- | | |
|----------|--|
| Schedule | <p>Select Roll back now or select Delay loading until and enter a future time and date:</p> <ul style="list-style-type: none"> • Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. • Enter a date in MM/DD/YYYY format or click  and select the month and date. |
| Reboot | <p>Click the check box to reboot the device after the configuration is rolled back. The configuration takes effect with or without a reboot. However, a reboot is recommended if the rolled back configuration has substantial changes or has a different topology setting.</p> <p>NOTE: When you reboot a device, all unsaved configuration data is lost.</p> |

6. To review the devices you selected, click **Show selected devices**.

7. Click **Preview** to view the rollback configuration of the first device. Click **Next** and **Back** to page through all the configurations. When you're done, click **Close**.

8. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

Backing Up Device Configurations

You can schedule the configuration file (*startup.cfg*) to be backed up periodically for each device running SRS 4.x or higher. Each configuration file is archived in the following directory on the CMS server:

```
<CMS file location>\data\configuration\config\device\RCS
```

The default <CMS file location> is C:\Program Files\Peribit\CMS.


The backup configuration files are archived as “Version 1.1”, “Version 1.2”, and so on. A new version is archived only if changes have occurred since the last backup. To restore an archived configuration file, refer to “Restoring Device Configurations” on page 54.

To back up device configurations:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices you want to back up or click **Select All**.
3. From the Task menu, select **Configuration > Backup** and click **Go**.

Figure 3-10 Backing up Configuration Files

4. Specify the following information:

- | | |
|----------|--|
| Backup | Select whether you want to save the running configuration before doing the backup (the default). Alternatively, you can back up the current saved configuration or the running configuration. |
| Schedule | <p>Select Backup now or select Delay backup until and enter a future time and date:</p> <ul style="list-style-type: none"> • Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. • Enter a date in MM/DD/YYYY format or click  and select the month and date. • To back up the configuration periodically, select a recurring retrieval interval (Daily, Weekly, or Monthly). Monthly schedules that start on the 31st are permanently reset to the last day of the shortest month encountered. |

5. To review the devices you selected, click **Show selected devices**.
6. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task and access the backed up configuration file, refer to “Managing Scheduled Tasks” on page 62.

Restoring Device Configurations

If you periodically back up the configuration file (*startup.cfg*) for each device, you can restore a backup configuration at any time. Note that each backup contains device-specific settings, so it can be restored only to its original device. To back up configuration files, refer to “Backing Up Device Configurations” on page 52.

To restore an archived configuration file:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select one device where you want to restore the configuration.
3. From the Task menu, select **Configuration > Restore** and click **Go**.

The screenshot shows the PeriScope IS Central Management System interface. The top navigation bar includes 'MY PERIBIT', 'MONITOR', 'MANAGEMENT', 'SETUP', and 'ABOUT'. The 'MANAGEMENT' tab is active, and the breadcrumb trail shows 'Devices > Restore Startup File'. The left sidebar contains a 'Management' menu with options like 'Devices', 'Configurations', 'Auto-Deployment', 'License Management', 'Schedules', and 'Schedule Log'. The 'Community/Device Group' dropdown is set to 'Group_1'. The main content area is titled 'Devices > Restore Startup File' and contains the following text and form elements:

Restore the selected version of the startup file to the selected device. [Show selected device](#)

Note that the parameters in the selected startup file will **replace** those on the devices. Please verify that you have set all the parameters correctly before loading the configuration on the devices.

Click on the "View" button to view the selected version.

Version:

Schedule:


- ☐ Restore now
- ☒ Delay restoring until:
 - Time: ☒ AM ☐ PM
 - Date:
 - Recurrence:

Reboot: ☐ Reboot device after restoring startup configuration

At the bottom are 'Submit' and 'Cancel' buttons.

Figure 3-11 Restoring a Configuration File

4. Specify the following information:

- | | |
|----------|--|
| Version | <p>Select the configuration to be restored. The most recent backup has the highest “1.n” version number. The list is empty if you have no backups for the selected device.</p> <p>Click VIEW to verify the settings in the selected configuration. All configuration settings are saved as CLI commands (the commands are described in the <i>Sequence Reducer/Sequence Mirror Operator’s Guide</i>).</p> |
| Schedule | <p>Select Restore now or select Delay restore until and enter a future time and date:</p> <ul style="list-style-type: none"> • Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. • Enter a date in MM/DD/YYYY format or click  and select the month and date. • To restore the configuration periodically, select a recurring retrieval interval (Daily, Weekly, or Monthly). Monthly schedules that start on the 31st are permanently reset to the last day of the shortest month encountered. |
| Reboot | <p>Click the check box to reboot the device after the configuration is restored. The configuration takes effect with or without a reboot. However, a reboot is recommended if the restored configuration has substantial changes or has a different topology setting.</p> <p>NOTE: When you reboot a device, all unsaved configuration data is lost.</p> |

5. To verify the device you selected, click **Show selected device**.

6. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

Retrieving Device Files

You can schedule the following files to be retrieved periodically from selected devices:

- **Diagnostic file.** Current configuration and the most recent system log and access control log.
- **System Log.** Critical, error, and information messages related to the operation of the device.
- **Access Control Log.** Log of each user who accessed the device in the last five days. Includes the workstation IP address for HTTPS and SSH access, and any configuration changes made by the user.
- **Monitor statistics.** All performance data from last month through the hour of the retrieval. The statistics are described in Appendix C, “Understanding Exported Data Results.”
- **Flow statistics.** Top traffic flows collected on the device. This file is empty if there are no top traffic statistics on the device.

The monitor and flow statistics are in CSV format, and can be imported into a spreadsheet or other data analysis program.

For each device, the retrieved files are compressed in ZIP or TAR file (TAR files are used only for diagnostic files). The files for each device are stored in the following directory on the CMS server:

<CMS file location>\data\download\<device ip address>

The *<CMS file location>* is *D:\Program Files\Peribit\CMS*, unless it was changed during installation.

NOTE: To retrieve device files, the Microsoft FTP server must be installed and running on the CMS server. Also, if you retrieve device files on a recurring basis, be sure that adequate disk space is available.

You can access the retrieved files from the Schedules page and download the files to the hard disk of your CMS Web console. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

To retrieve files from selected devices:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices that you want to retrieve files from or click **Select All**.
3. From the Task menu, select **Diagnostics/Statistics** and click **Go**.

PeriScope CMS
CENTRAL MANAGEMENT SYSTEM

Juniper
NETWORKS

MY PERIBIT MONITOR MANAGEMENT CMS SETUP ABOUT Logged in as: root LOGOUT

Management

Devices

Configurations

Auto-Deployment

License Management

Schedules

Schedule log

Community/Device Group

NSM Community

☐ Show Exceptions Only

Submit

Devices > Retrieve Statistics/Diagnostic Files

Retrieve statistics/diagnostic files from the selected device(s). [Show selected devices](#)

If only Monitor statistics and/or Flow Statistics is selected then you can retrieve them on an hourly basis.

Retrieve Files:

☐ Diagnostic file

☐ System Log

☐ Access Control Log

☐ Monitor statistics

☐ Flow statistics (Top talker data)

Schedule

☐ Retrieve now

☒ Delay retrieval until:

Time HH:MM ☒ AM ☐ PM

Date


Recurrence No Recurrence

The file downloaded for each device will be stored in the sub-directory identified by the device's ip address under the directory "C:\Program Files\Peribit\CMS\data\download". All files for each device will be bundled into one ZIP file.

Submit **Cancel**

Figure 3-12 Retrieving Statistics/Diagnostic Files

4. Select the checkbox next to the files you want to retrieve. If you select the diagnostic file, the system and access control logs are included, so you do not need to select them separately.

5. Select **Retrieve now** or select **Delay retrieval until** and enter a time and date, and a recurring retrieval interval (optional):
 - Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click **AM** or **PM**. Note that midnight can be 0:0 AM or 12:00 AM.
 - Enter a date in MM/DD/YYYY format or click  and select the month and date.
 - To retrieve the selected files periodically, select a recurring retrieval interval (**Hourly**, **Daily**, **Weekly**, or **Monthly**). The hourly interval is available if only monitor and/or flow statistics are selected. Monthly schedules that start on the 31st are permanently reset to the last day of the shortest month encountered.
6. To review the devices you selected, click **Show selected devices**.
7. Click **Submit** to submit the task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task and access the retrieved files, refer to “Managing Scheduled Tasks” on page 62.

Applying a Registration Server Password

Each device accesses the registration server periodically to identify the other devices in the same community. If you change the registration server’s password, you must apply the new password to all devices in each community managed by the registration server.

To change a registration server’s password, do the following:


- Use the SRS Web console to change the password on the registration server.
- Enter the new password in CMS (refer to “Managing Communities” on page 328).
- Apply the new password to the devices in each community, as described below.

NOTE: All devices reporting to the same registration server must use the same registration server password.

To apply a new registration server password to all devices in a community:

1. On the Devices page, select a community from the **Community/Device Group** menu.
2. Click **Select All** to select all devices in the community.
3. From the Task menu, select **Apply password** and click **Go**.

Figure 3-13 Applying a Registration Server Password

4. Select **Apply password now** or select **Delay until** and enter a time and date:
 - Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click **AM** or **PM**. Note that midnight can be 0:0 AM or 12:00 AM.
 - Enter a date in MM/DD/YYYY format or click  and select the month and date
5. To review the devices you selected, click **Show selected devices**.
6. Click **Submit** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

Putting Devices in Safe Mode

If you have problems with your network or a specific device, you can put the device in safe mode. In safe mode, the device is powered on and can be managed over the network, but all traffic is passed through without data reduction or assembly.

NOTE: Putting a device in safe mode reboots the device, so any unsaved configuration data is lost.

To put devices in safe mode:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Select the devices that you want to put in safe mode.
3. From the Task menu, select **Put in 'SAFE' mode** and click **Go**.

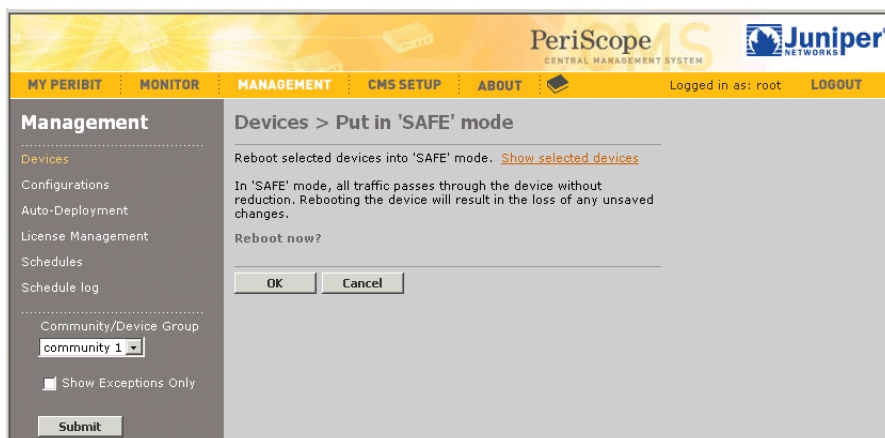


Figure 3-14 Putting Devices in Safe Mode

4. To review the devices you selected, click **Show selected devices**.
5. Click **OK** to submit this task, or click **Cancel**.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 62.

6. Perform the appropriate troubleshooting and diagnostics. When you are done, reboot the devices to enable data reduction (refer to “Rebooting Devices” on page 42).

Accessing the SRS Web Console from CMS

You can configure individual devices by accessing the SRS Web console from CMS. To configure a device:

1. On the Devices page, click the name of the device that you want to configure.
2. Enter the user name and password. The SRS Web console opens.
3. For complete information about configuring a device from the SRS Web console, see the *Sequence Reducer/Sequence Mirror Operator's Guide*.

Exporting Community and Device Information

Information about the devices in a selected community or device group can be exported to a file in comma-separated variable (CSV) format. The CSV file can then be imported into a spreadsheet program (such as Microsoft Excel) or other data evaluation program.

The exported file contains the following information:

- Community or device group name
- Registration server IP address
- Date and time of the export
- The following information for each device:
 - Device name.
 - IP address.
 - Model number.
 - Serial number.
 - Local MAC address.
 - Remote MAC address.
 - License speed.
 - Software version.

To export device information:

1. On the Devices page, select a community or device group from the **Community/Device Group** menu.
2. Click **Export** in the upper-right corner of the page.
3. To save the file to a local hard disk, click **Save** and specify a file name and location.

Managing CMS Schedules

The following sections describe how to use the CMS scheduler:

- “Managing Scheduled Tasks” on page 62
- “Exporting a Schedule Log” on page 66

Managing Scheduled Tasks

The Schedules page lets you view all device tasks scheduled in the last 15 days, including tasks that are pending, in-process, successful, failed, or cancelled. To view tasks older than 15 days, refer to “Exporting a Schedule Log” on page 66.


The following scheduling actions are available:

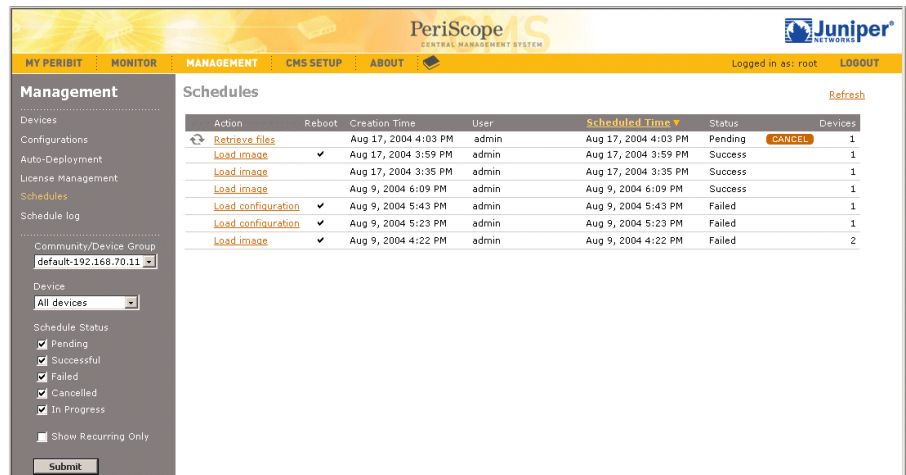
- **Acknowledge.** Failed tasks can be acknowledged, which removes the failed task icon from the Devices page.
- **Add Devices.** Pending and recurring tasks can be updated to include more devices.
- **Cancel.** Pending tasks can be cancelled for all or selected devices.
- **Reschedule.** Failed and pending tasks can be rescheduled.

Also, for tasks that retrieve files from a device or back up a configuration, you can open the retrieved files or download them to a local disk.

If your network is having problems, you can stop and restart the scheduler, as described in “Stopping and Starting the Scheduler” on page 344.

To manage the scheduled tasks:

1. Click **MANAGEMENT** in the menu frame, and then click **Schedules** in the left-hand navigation frame. Alternatively, click  in the Devices page to view the failed tasks for a device.
2. Change one or more of the following parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu, or select **All Communities** to view the tasks are scheduled for all devices.
 - Select a name from the **Device** menu to view tasks only for the selected device. The default is **All devices**.
 - Click the check box next to the status of the tasks you want to view, such as **Failed** or **Pending**. Click **Show Recurring Only** to view tasks that have the selected status AND are run periodically.





Action	Reboot	Creation Time	User	Scheduled Time	Status	Devices
 Retrieve files		Aug 17, 2004 4:03 PM	admin	Aug 17, 2004 4:03 PM	Pending	1
Load image	✓	Aug 17, 2004 3:59 PM	admin	Aug 17, 2004 3:59 PM	Success	1
Load image		Aug 17, 2004 3:35 PM	admin	Aug 17, 2004 3:35 PM	Success	1
Load image		Aug 9, 2004 6:09 PM	admin	Aug 9, 2004 6:09 PM	Success	1
Load configuration	✓	Aug 9, 2004 5:43 PM	admin	Aug 9, 2004 5:43 PM	Failed	1
Load configuration	✓	Aug 9, 2004 5:23 PM	admin	Aug 9, 2004 5:23 PM	Failed	1
Load image	✓	Aug 9, 2004 4:22 PM	admin	Aug 9, 2004 4:22 PM	Failed	2

Figure 3-15 Schedules Page

The Schedules page provides the following information for each task:

- **Action**—Task name. The  icon indicates a recurring task. Move the cursor over the icon to view the frequency and the next run time.
- **Reboot**—A check mark indicates that the task includes a reboot after the task is performed.
- **Creation Time**—Date and time the task was submitted to the scheduler.

- **User**—ID of the user who submitted the task.
- **Scheduled Time**—Date and time that the task is scheduled. For a recurring schedule that has run once, this is the scheduled time of the last run.
- **Status**—The status of the task. For a recurring schedule that has run once, this is the status of the last run.
- **Devices**—Number of devices for which the task is scheduled.

NOTE: A “Failed” status indicates the task has failed for at least one device.

3. To cancel a pending task for all devices, click the **CANCEL** button. The status of the task is changed from “Pending” to “Cancelled”.

To cancel a pending task for specific devices, view the results of a task, or perform other scheduling functions, click the name of the appropriate task.

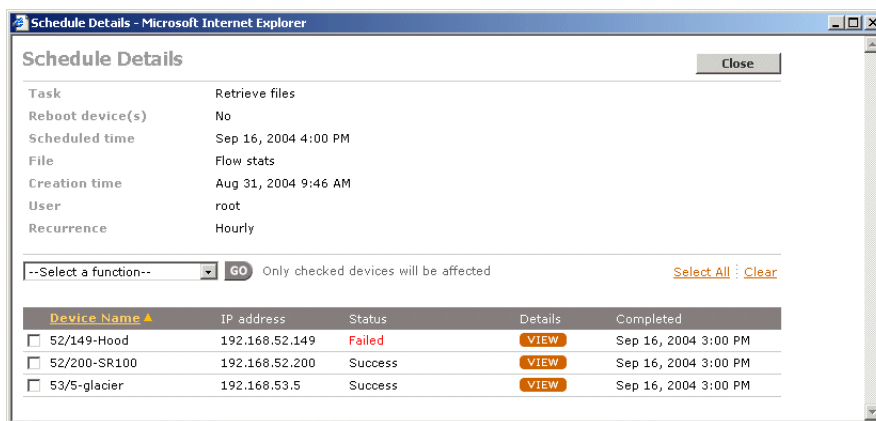


Figure 3-16 Viewing Task Details

The Schedule Details page displays general task information, plus the status, details and completion time for each device the task is scheduled for. The **Details** column contains one of the following:

- Details about a failed task
- A link to a retrieved file (for one-time retrievals). Click the link to open or save the file to a local disk.

- A **VIEW** button for a recurring schedule. Click the button to view the details of each run for a specific device. A recurring backup or file retrieval includes a link to each retrieved file.

4. To change, acknowledge, reschedule, or cancel a scheduled task:

Add devices to pending or recurring tasks

1. Select **Add Devices** on the task menu, and click **Go**.

2. Select the devices to be added and click **Submit**.

Reschedule failed or pending tasks

1. Click the check box next to the devices for which you want to reschedule the task.

2. Select **Reschedule** on the task menu, and click **Go**.

3. Reschedule the task and click **Submit**.


For each selected device, the status is changed to “Failed: Rescheduled” for failed tasks or “Cancelled: Rescheduled” for pending tasks, and a new “Pending” task is added to the Schedules page.

If you reschedule a pending task for all its devices, the original task is changed to “Cancelled” on the Schedules page.

Acknowledge failed tasks

1. Click the check box next to the devices for which you want to acknowledge the failed task.

2. Select **Acknowledge ‘Failed’ Status** on the task menu, and click **Go**.

The  icon is removed from the selected devices on the Devices page, and the status is changed to “Failed: Acknowledged” on the Schedule Details page.

Cancel pending tasks

1. Click the check box next to the devices for which you want to acknowledge the failed task.

2. Select **Cancel** on the task menu, and click **Go**.

For each selected device, the status is changed to “Cancelled”. If the task is still pending for some devices, the task remains on the Schedules page as a “Pending” task.

Exporting a Schedule Log

You can export a schedule log containing information about tasks submitted to the scheduler for a particular community or device group. The file contains the following information:

- Community name or device group
- Date and time that the schedule log is exported from CMS
- Task identification number and the task itself
- If a reboot was scheduled after the task
- Device name and IP address
- Scheduled date and time for the task
- Status of the task
- Files associated with the task
- Additional task details (if any)
- Date and time that the task was completed
- Creation time
- User who scheduled the task

You can save the file in CSV format on a local disk, and then import its contents into a spreadsheet program (such as Microsoft Excel).

To export the schedule log:

1. Click **MANAGEMENT** in the menu frame, and then click **Schedule log** in the left-hand navigation frame.
2. Select a community or device group from the **Community/Device Group** menu, and click **Submit**.
3. To save the file to a local disk, click **Save** and specify the file location.

Chapter 4 Managing Device Configurations

This chapter describes how to use the CMS to generate and manage device configurations. It covers the following topics:

- “Overview of Device Configurations” on page 67.
- “Viewing Configurations” on page 73.
- “Managing Configurations” on page 75.
- “Defining Configuration Settings” on page 83.

Overview of Device Configurations

You can use CMS to define multiple sets of configuration settings that can be selectively combined and downloaded to one or more devices in a community. You can define two types of configurations in CMS:

- **Global configurations.** Includes all configuration settings that can be defined in the device Web console for SRS 5.1 devices, except for the packet capture password and the IPSec Setup Wizard. You can also append CLI commands to enable features that are available only through the CLI.
- **Partial configurations.** Includes one type of configuration settings defined in a global configuration. Partial configurations let you change specific configuration settings, such as for QoS or data reduction, without having to create an entire global configuration for each minor change to the common settings shared by most devices.

The Device Settings partial configuration can be used to configure settings for one device, such as the IP address. Alternatively, you can define device-specific settings through the device Web console (refer to “Accessing the SRS Web Console from CMS” on page 61).

All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator’s Guide*.

Configuration Settings for SRS 5.1 and 5.0 Devices

Table 4-1 lists the configuration settings that can be defined in CMS for each type of partial configuration. Except for Device Settings, a global configuration includes all of the partial configuration settings. Most of the partial configurations correspond to parameter groupings in the device Web console.

Note the differences between SRS 5.1 and 5.0 configurations. For example, for SRS 5.1, CLI commands can be specified under Advanced Settings; for SRS 5.0, CLI commands can be added at the end of a global configuration.

Table 4-1 Partial Configurations for SRS 5.x Devices

Type	CMS Settings	Notes
Device Settings	<ul style="list-style-type: none"> • Addresses • Communities (5.1 only) • Time zone • ARP • Reduction subnets • Outbound QoS exclusions • Static local routes • Dynamic local routes (router polling) • Multi-Path (secondary IP address) • RADIUS source IP address (5.1 only) 	
Basic Setup	<ul style="list-style-type: none"> • Interfaces • Time (NTP servers) • SNMP • Syslog server • Dynamic local routes (OSPF/RIP) • Router balancing • Registration server • NetFlow 	<ul style="list-style-type: none"> • To apply license keys from CMS, refer to “Automatic Deployment of Devices” on page 253. • To apply a new registration server password to multiple devices, refer to “Applying a Registration Server Password” on page 58.
AAA	<ul style="list-style-type: none"> • Authentication • Authorization • RADIUS • Local users • Operator access • Front panel access 	<ul style="list-style-type: none"> • Packet capture passwords cannot be defined in CMS.
Applications	<ul style="list-style-type: none"> • Overview (5.1 only) • Application definitions • Traffic classes • Monitoring (5.1 only) 	

Table 4-1 Partial Configurations for SRS 5.x Devices (Continued)

Type	CMS Settings	Notes
Reduction	<ul style="list-style-type: none"> Endpoints Network Sequencing Mirroring (5.1 only) Application filter Remote routes Load balancing Default assemblers Preferred assemblers Tunnel mode 	<ul style="list-style-type: none"> For SRS 5.0 devices, Network Sequence Mirroring must be configured on the device. If you upgrade a 5.0 Reduction partial configuration to 5.1, application monitoring is removed from the application filter (monitoring is defined under Applications in 5.1)
QoS	<ul style="list-style-type: none"> Setup Wizard Overview Traffic classes (in 5.0 global configurations) Templates Endpoints ToS/DSCP Start/stop Inbound Qos 	
Acceleration	<ul style="list-style-type: none"> Overview Flow Pipelining (5.0 only) Fast Connection Setup Active Flow Pipelining CIFS applications (5.1 only) HTTP applications (5.1 only) Exchange applications (5.1 only) 	
Advanced Setup	<ul style="list-style-type: none"> Topology Source/destination filter Prime time Packet interception WAN performance monitor (5.1 only) CLI commands (5.1 only) 	
Multi-Path	<ul style="list-style-type: none"> Start/stop Templates Endpoints 	
IPSec	<ul style="list-style-type: none"> Overview Templates Default policy 	<ul style="list-style-type: none"> The IPSec Setup Wizard can be run only from device Web console.

Downloading Global and Partial Configurations

When you download configuration settings to a device, you can select one global and zero or more partial configurations, or just a combination of partial configurations. Downloaded configurations are processed as follows:

- **Global configuration OR partial configurations.** The selected settings in the global or partial configurations override the corresponding settings on each device.
- **Global AND partial configurations.** The selected settings in the partial configurations override the settings in the global configuration, and the result overrides the corresponding settings on each device.

Configuration settings are selected by setting check boxes. Figure 4-1 shows a global configuration with the **Features/Topology** settings selected. Selecting a check box enables the default settings. You can then click the link next to the check box to change the default. If you clear the check box, the corresponding settings are removed from the configuration.

Figure 4-1 Selecting Configuration Settings

When you load a configuration on a device, only the checked settings affect the device. Device settings that can be specified only by CLI commands can be overridden by commands in the CLI section of a global or partial configuration.

Note that device settings are replaced, not supplemented. For example, if a device has four traffic classes, and you load a configuration that has two classes, the resulting device configuration will have two traffic classes, not six.

Figure 4-2 shows how downloaded global and partial configurations are added to a device’s configuration in the *startup.cfg* file (any unsaved device settings will be lost). The gaps indicate unchecked settings in the global and partial configurations. Any CLI commands specified in the CLI section of the downloaded configurations are appended to the *startup.cfg* file on the device.

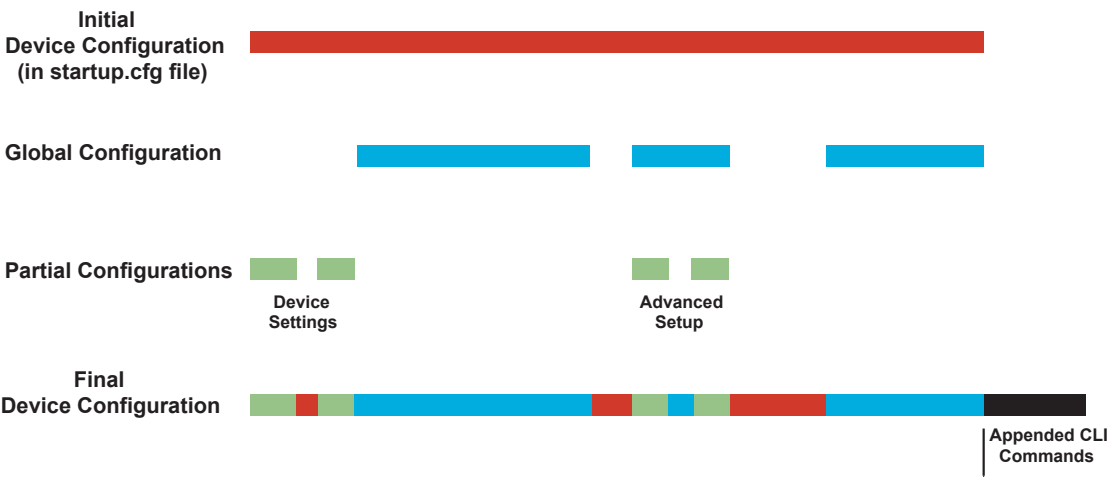


Figure 4-2 Downloading Global and Partial Configurations on a Device

For example, new devices have a default topology setting of “mesh,” with the range of devices set to zero (the lowest range). If you use a hub and spoke topology, you can create a global configuration for the spoke devices and an Advanced Setup partial configuration that specifies the hub setting and the appropriate range of devices in the community.

To configure a new device as a hub, simply download the two configurations. Table 4-2 shows the topology CLI commands in each configuration.

Table 4-2 Combining Global and Partial Configurations

Global Configuration	Advanced Setup Configuration	Resulting Device Configuration
config reduction set topology-type spoke	config reduction set topology-type hub config reduction set topology-size 1	config reduction set topology-type hub config reduction set topology-size 1
.		Plus all other downloaded settings.
.		
.		

NOTE: Some default settings have no explicit CLI commands. For example, an Advanced Setup configuration with the default topology setting would have no topology commands, but loading the configuration on a device would erase the current topology commands (if any).

For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*. To download global configuration settings from CMS, refer to “Loading Device Configurations” on page 48.

Consistency Checking

When you save a configuration in CMS, you are prompted to correct any incomplete settings. For example, if CIFS acceleration is enabled, Application Flow Pipelining (AFP) must also be enabled. You can save a configuration that has errors, and fix the errors at a later time.

When you load a configuration on a device, an error occurs if the resulting configuration would have incomplete or inconsistent settings. For example, if you download a configuration that has the default topology settings (Application Flow Acceleration is disabled), the task will fail for any device that is already using CIFS, Exchange, or HTTP acceleration.

Tracking Configuration Versions

A version history is maintained for each global and partial configuration defined in CMS. The first version is 1.1, and subsequent versions are numbered 1.2, 1.3, and so on. You can enter a description of the changes when you save a new version, and you can view or compare any of the previous versions.

In addition, when you display the CLI commands in a configuration, the first line specifies the format version:

- **Version 2.1.** Compatible with SRS 5.1 devices only
- **Version 2.0.** Compatible with SRS 5.0 devices only

The format versions prevent configurations from being loaded on incompatible devices.

Tips for Managing Configurations

Review the following tips for managing global configuration settings:

- Analyze your existing device configurations to determine which configurations you want to extract and maintain on CMS. For more information, refer to “Analyzing Device Configurations” on page 45 and “Extracting Configurations” on page 75.
- Maintain a small number of unique global configurations, and define partial configurations to customize specific settings, such as topology settings, for selected devices and communities.
- Assign configuration names that reflect the contents of the configuration. Examples of configuration names include “hub-config,” “bandwidth-policy,” and “community-east.”
- Use the version history to track configuration changes and compare previous versions with the current version (refer to “Viewing Configuration History” on page 82).

Viewing Configurations

The Configurations page lets you view, generate, and manage global configuration settings in CMS. Note that all configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.

To view the configurations defined in CMS:

1. Click **MANAGEMENT** in the menu frame, and then click **Configurations** in the left-hand navigation frame.
2. Select the type of configurations you want to view from the **Show Configurations** menu (all, global, or partial), and click **Submit**. If you select **Partial Configurations**, you can select a specific type.

Task menu

Configurations

--Select a task-- GO Only checked configurations will be affected [Select All](#) [Clear](#)

Config Name ▲	Config Type	Description	Compatibility	Requires	Modified Date
<input type="checkbox"/> App_defaults	Applications	Initial app definitions (Extracted from Test)	SRS 5.0		12/11/03 10:13 AM
<input type="checkbox"/> Hub_topology	Advanced Setup	Topology setting for hubs	SRS 5.0		12/2/03 5:11 PM
<input type="checkbox"/> Main_config	Global	Main base config (Extracted from 192.168.8.200)	SRS 5.0		12/2/03 4:51 PM
<input type="checkbox"/> PFA_apps	Acceleration	Default PFA app settings	SRS 5.0	App_defaults	12/11/03 3:34 PM
<input type="checkbox"/> Reduction_defaults	Reduction	Default reduction settings	SRS 5.0	App_defaults	12/11/03 10:12 AM

Figure 4-3 Configurations Page

From the Configurations page, you can:

- View the type, description, compatibility, and last-modified date for each global and partial configuration. Some partial configurations (QoS, Reduction, Acceleration, and Multi-Path) must reference the application definitions in a global or Applications configuration. The **Requires** column indicates the referenced configuration.
- Click the column headers to change the sort.
- Execute a configuration task, such as generating new configurations, as described in “Managing Configurations” in the next section.
- Click the configuration name to change the settings, as described in “Defining Configuration Settings” on page 83.

Managing Configurations

The following sections describe how to define and manage global configurations in CMS:

- “Extracting Configurations” in the next section
- “Duplicating Configurations” on page 77
- “Creating New Configurations with Factory Defaults” on page 78
- “Comparing Configurations” on page 80.
- “Displaying Configurations” on page 81
- “Viewing Configuration History” on page 82
- “Deleting Configurations” on page 82

Extracting Configurations

You can define new global configurations by extracting the saved configuration (the *startup.cfg* file) from a selected device. To determine which device configuration to extract, you can analyze the configurations of your current devices, as described in “Analyzing Device Configurations” on page 45.

You can also define new partial configurations by extracting the related configuration settings, such as application definitions, from a device or a global configuration. After you extract and modify configurations, you can load them on selected devices, as described in “Loading Device Configurations” on page 48.

You cannot extract Acceleration, Reduction, QoS, or Multi-Path partial configurations from a device because they must reference the applications and traffic classes in another configuration.

NOTE: An extracted configuration has all settings enabled (all check boxes are selected). Any settings that cannot be defined in the CMS Web console are shown in the CLI section of the Advanced Setup configuration settings. You should review and edit (or disable) these settings before loading the extracted configuration on another device.

To extract a configuration from a device or a global configuration:

1. On the Configurations page, select **Extract** on the Task menu, and click **Go**.

Management

Devices
Configurations
Auto-Deployment
License Management
Schedules
Schedule log

Show Configurations
All Configurations

Submit

Configurations > Extract

This page allows you to extract and save a configuration. You can extract a global configuration from a Peribit device or you can extract a partial configuration from a global configuration. Only configuration parameters that can be edited using Periscope CMS are extracted.

☒ Extract global configuration from Peribit device
☐ Extract partial configuration from Peribit device
☐ Extract partial configuration from global configuration

Community:
Device:

Enter a name and description for the extracted configuration

Configuration name:
Description:

Submit Cancel

Figure 4-4 Extracting Configurations

2. Do one of the following:
 - a. To extract a global configuration from a device, select a community from the **Community** menu, and select a device name from the **Device** menu.
 - b. To extract a partial configuration from a device, click **Extract partial configuration from Peribit device**, select a community from the **Community** menu, select a device name from the **Device** menu, and select the configuration type from the **Partial Config Type** menu.
 - c. To extract a partial configuration from a global configuration, click **Extract partial configuration from global configuration**, select a global configuration from the **Global Configuration** menu, and select the configuration type from the **Partial Config Type** menu.
3. Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_).
4. Enter a description of the configuration. The text “(Extracted from <source>)” is appended to the description, where <source> is the device IP address or the global configuration name.

5. Click **Submit** to add the new configuration to the Configurations page.
Extracting a global configuration from a device may take some time.
6. Click the configuration name to change its settings, as described in “Defining Configuration Settings” on page 83.

Duplicating Configurations

You can define new configurations by copying and modifying an existing global or partial configuration. After you copy and modify configurations, you can load them on selected devices, as described in “Loading Device Configurations” on page 48.

To copy an existing global or partial configuration:

1. On the Configurations page, select the checkbox next to the configuration that you want to copy.
2. From the Task menu, select **Duplicate** on the task menu, and click **Go**.

Figure 4-5 Duplicating a Configuration

3. Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_).
4. Enter a description of the configuration. The text “(Duplicated from <source>)” is appended to the description, where <source> is the name of the global or partial configuration.
5. Click **Submit** to add the new configuration to the Configurations page.
6. Click the configuration name to change its settings, as described in “Defining Configuration Settings” on page 83.

Creating New Configurations with Factory Defaults

You can create new global or partial configurations without extracting or copying the configuration from another source. In this case, all parameters have the same default values as a new SR-50 before Quick Setup is run.

NOTE: A new global configuration cannot be loaded on a device unless you change the default administrator password. Also, if you specify an incorrect registration server address, the device will lose access to the other devices in the community, and CMS will lose access to the device within 24 hours.

To ensure that the password and registration server are correct, create new global configurations by extracting them from a working device (refer to “Extracting Configurations” on page 75).

After you create new configurations, you can load them on selected devices, as described in “Loading Device Configurations” on page 48.

To create a new configuration with factory defaults:

1. On the Configurations page, select **New** on the Task menu, and click **Go**.

The screenshot shows the PeriScope CMS web interface. The top navigation bar includes 'MY PERIBIT', 'MONITOR', 'MANAGEMENT', 'SETUP', and 'ABOUT'. The 'MANAGEMENT' tab is active. On the left, the 'Management' sidebar lists 'Devices', 'Configurations' (highlighted), 'Auto-Deployment', 'License Management', 'Schedules', and 'Schedule log'. Below this is a 'Show Configurations' section with a dropdown set to 'All Configurations' and a 'Submit' button. The main content area is titled 'Configuration > New' and contains the following text: 'Create new configuration. A new configuration has the same default values for all the parameters as a new device. Please be aware that you need to completely configure this configuration before loading it on any device. In particular, if important parameters like the administrator password and registration server information are not configured, you will not be able to access the device via the Web or SSH and the device will not be able to communicate with the other devices in the community. If you already have fully configured Peribit devices and are now using PeriScope CMS to manage them, you should consider creating the configuration using the Extract task.' Below the text are form fields: 'Configuration name' (text input), 'Description' (text input), 'Compatibility' (radio buttons for 5.1 and 5.0, with 5.1 selected), 'Config Type' (radio buttons for Global and Partial, with Global selected), and a dropdown menu for 'AAA' (currently showing 'AAA'). At the bottom are 'Submit' and 'Cancel' buttons.

Figure 4-6 Creating a New Configuration with Factory Defaults

Specify the following information:

Configuration name	Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_).
Description	Enter a configuration description (up to 100 characters).
Compatibility	Select 5.1 or 5.0 to indicate the version of SRS on the devices where the new configuration will be loaded.
Config Type	Select the new configuration type: <ul style="list-style-type: none"> • Global. Contains all settings that can be defined in CMS (except device-specific settings). • Partial. Contains one group of settings. For Reduction, QoS, Acceleration, or Multi-Path configurations, you must also select a global configuration or an Applications partial configuration that contains the application definitions. For a Multi-Path partial configuration, the selected configuration must also specify QoS traffic classes.

2. Click **Submit** to add the new configuration to the Configurations page.

3. Click the configuration name to change its settings, as described in “Defining Configuration Settings” on page 83. For a new global or AAA partial configuration, you must change the default password for the **admin** account (refer to “Defining Local Users” on page 123).

You should also review all the default settings (all items that have the check box selected in the left frame), as described in the following sections:

- “Configuring Application Definitions” on page 131
- “Monitoring Applications” on page 138
- “Reducing Applications” on page 144
- “Configuring Tunnel Mode Settings” on page 153
- “Enabling Active Flow Pipelining by Application” on page 203
- “Enabling Fast Connection Setup by Application” on page 204

NOTE: If you do not specify the Features/Topology settings, they default to all features except Application Flow Acceleration, and a mesh topology with the lowest range of devices (refer to “Configuring the Feature/Topology Settings” on page 212).

Comparing Configurations

To view a line-by-line comparison of the CLI commands in two configurations:

1. On the Configurations page, select the check box next to two configurations that you want to compare.
2. From the Task menu, select **Compare** and click **Go**.

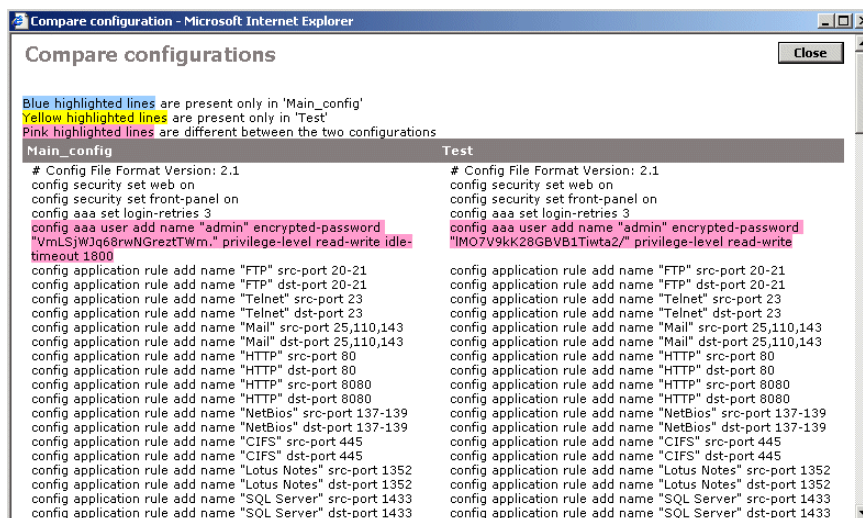


Figure 4-7 Comparing Configurations

The Compare configurations window displays a line-by-line comparison of the settings that can be defined in a global configuration. Color-coded lines indicate the following:

- **Blue**. Settings unique to the configuration in the left column.
- **Yellow**. Settings unique to the configuration in the right column.
- **Pink**. Settings that are different between the two configurations.

3. When you are done viewing the configurations, click **Close**.

For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.

Displaying Configurations

To view the CLI commands in the latest version of a configuration:

1. On the Configurations page, select the check box next to the configuration you want to view.
2. From the Task menu, select **Display** and click **Go**.

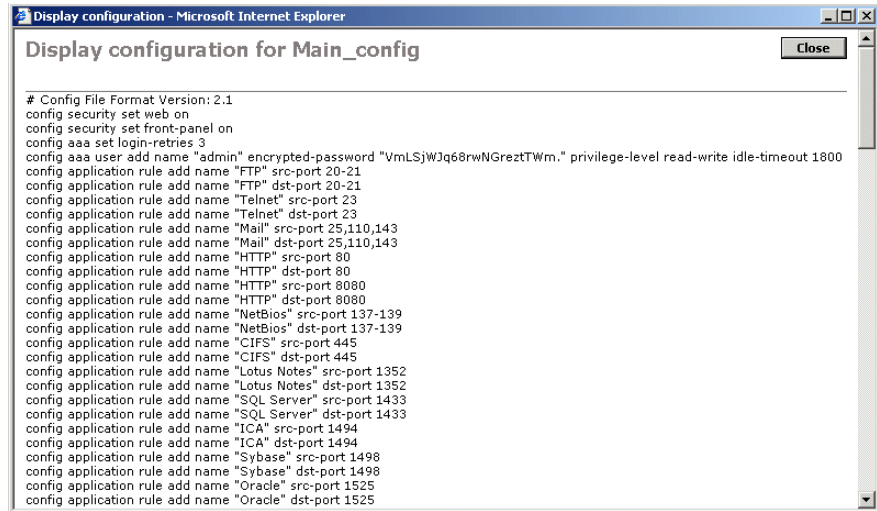


Figure 4-8 Displaying a Configuration

3. When you are finished viewing the configuration, click **Close**.

For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.

Viewing Configuration History

You can view a history of the changes to each global and partial configuration defined in CMS. Each previous version is retained, along with a description of the changes to each version, the time of the change, and the user responsible. You can view or compare any two versions of a configuration.

To view a configuration's history:

1. On the Configurations page, select a configuration.
2. From the Task menu, select **History** and click **Go**.

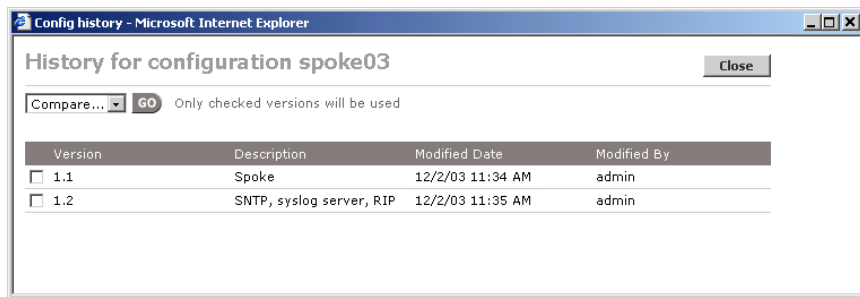


Figure 4-9 Viewing Configuration History

For each version, the History window displays a version number, description of the change, the date and time of the change, and the user responsible. CMS assigns version number 1.1 to a new configuration, and increments the number each time the configuration is changed (1.2, 1.3, and so on).

3. To view a line-by-line comparison of two versions, select the two versions that you want to compare, select **Compare**, and click **Go**.
4. To view the contents of a version, select the version that you want to view, select **Display**, and click **Go**.

Deleting Configurations

To delete configurations on CMS:

1. On the Configurations page, select the check box next to the configurations you want to delete or click **Select All**.
2. From the Task menu, select **Delete** and click **Go**.
3. At the confirmation prompt, click **OK** to delete the configurations.

NOTE: You cannot delete a configuration that is referenced by an auto-deployment group or by a QoS, Reduction, Acceleration, or Multi-Path partial configuration.

Defining Configuration Settings

After you generate a new configuration, you can define or change its settings. Remember that if you create a configuration as described in “Creating New Configurations with Factory Defaults” on page 78, all parameters have the same default settings as a new SR-50 device before Quick Setup is run.

All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*. To load a configuration on selected devices, refer to “Loading Device Configurations” on page 48.

NOTE: If you load a configuration that specifies an incorrect registration server IP address, CMS will lose access to the device in 24 hours, and the device will lose access to the other devices in the community.

To define configuration parameters:

1. On the Configurations page, click the name of a configuration.

The Configuration window opens for the selected global or partial configuration. Figure 4-10 shows a global configuration for SRS 5.1.

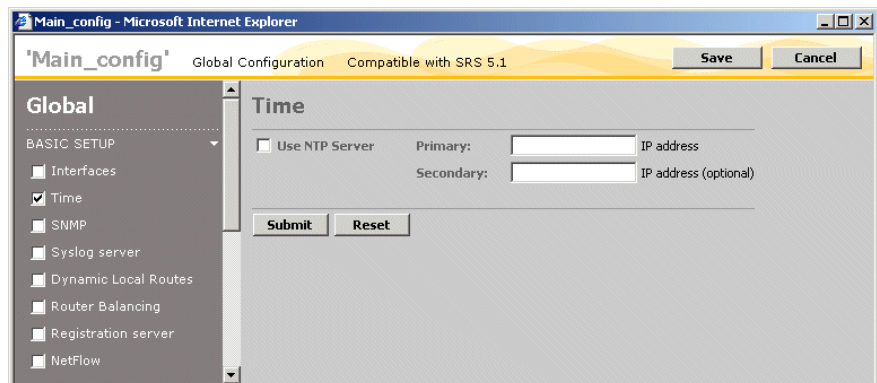


Figure 4-10 Editing a Global Configuration

Figure 4-11 shows a Basic Setup partial configuration.

Figure 4-11 Editing a Partial Configuration

2. To change a setting, select the check box next to the setting in the left-hand navigation frame, and then select the related page link. You can then change the setting and click **Submit**.

When you extract a configuration from a device, all the check boxes are selected, so you should review each setting before loading the configuration on another device. When you create a new configuration, only the check boxes for the default settings are selected.

NOTE: Clearing a check box deletes the associated settings. When you load a configuration, only the checked settings affect the device.

Refer to the sections listed in Table 4-3 for instructions on configuring each parameter.

3. When you are done changing the configuration, click **Save**, enter a description of the changes, and click **OK**. If a System Error page is displayed listing missing or incorrect settings, click **Back** to correct the problems, and then click **Save** again.

If there are no errors, CMS creates an updated configuration with a new version number. The version number and change description can be viewed in the configuration history (refer to “Viewing Configuration History” on page 82).

If you close the Configuration window without clicking **Save**, all the submitted changes are discarded.

Table 4-3 lists the sections that describe each group of configuration parameters for SRS 5.1 devices. The Device Settings must be defined in a partial configuration. Each of the other parameter groups can be defined in a global or partial configuration. To define configurations for SRS 5.0 devices, refer to the *CMS 5.0 Administrator's Guide*.

Table 4-3 Directory of Configuration Parameters

Parameter Group	Sections
Device Settings	<ul style="list-style-type: none"> "Configuring Device Addresses" on page 88 "Defining Communities" on page 90 "Configuring Time Zone Settings" on page 92 "Configuring the ARP Table" on page 93 "Advertising Reduction Subnets" on page 94 "Defining Outbound QoS Exclusions" on page 96 "Adding Static Routes" on page 97 "Configuring Router Polling" on page 99 "Configuring Multi-Path Addresses" on page 101 "Configuring the RADIUS Source Address" on page 103
Basic Setup	<ul style="list-style-type: none"> "Configuring the Interface Settings" on page 104 "Configuring NTP" on page 107 "Enabling SNMP" on page 108 "Enabling Syslog Reporting" on page 109 "Configuring Dynamic Local Routes" on page 110 "Enabling Route-Based Router Balancing" on page 112 "Designating a Registration Server" on page 114 "Generating NetFlow Records" on page 115
AAA	<ul style="list-style-type: none"> "Selecting Authentication Methods" on page 117 "Enabling Authorization Checking" on page 120 "Defining RADIUS Servers" on page 121 "Defining Local Users" on page 123 "Securing Operator Access" on page 125 "Securing Front Panel Access" on page 126

Table 4-3 Directory of Configuration Parameters (Continued)

Parameter Group	Sections
Applications	<ul style="list-style-type: none"> “Configuring Application Settings” on page 127 “Viewing the Application Overview” on page 130 “Configuring Application Definitions” on page 131 “Assigning Applications to Traffic Classes” on page 136 “Monitoring Applications” on page 138
Reduction	<ul style="list-style-type: none"> “Configuring Endpoints for Reduction Tunnels” on page 139 “Configuring Network Sequence Mirroring” on page 141 “Reducing Applications” on page 144 “Configuring Remote Routes” on page 146 “Configuring Tunnel Load Balancing Policies” on page 147 “Configuring Default Assemblers” on page 149 “Defining Preferred Assemblers” on page 151 “Configuring Tunnel Mode Settings” on page 153
QoS	<ul style="list-style-type: none"> “Understanding Outbound QoS” on page 156 “Using the Outbound QoS Setup Wizard” on page 168 “Defining Outbound QoS Settings by Endpoint” on page 176 “Defining Outbound QoS Templates” on page 179 “Defining Outbound QoS Endpoints” on page 180 “Changing Outbound ToS/DSCP Values” on page 184 “Starting and Stopping Outbound QoS” on page 187 “Configuring Inbound QoS Policies” on page 188
Acceleration	<ul style="list-style-type: none"> “Configuring Traffic Acceleration” on page 190 “Enabling Acceleration by Endpoint” on page 198 “Enabling Acceleration by Application” on page 203
Advanced Setup	<ul style="list-style-type: none"> “Configuring the Feature/Topology Settings” on page 212 “Configuring Source/Destination Filters” on page 216 “Defining the Prime Time” on page 218 “Configuring Packet Interception” on page 220 “Configuring WAN Performance Monitoring” on page 231 “Adding CLI Commands to Configurations” on page 234

Table 4-3 Directory of Configuration Parameters (Continued)

Parameter Group	Sections
Multi-Path	“Enabling Policy-Based Multi-Path” on page 237 “Defining Multi-Path Templates” on page 238 “Defining Multi-Path Endpoints” on page 240 “Configuring Routers to Support Multi-Path” on page 243
IPSec	“Defining IPSec Settings by Endpoint” on page 247 “Defining IPSec Templates” on page 249 “Defining the Default IPSec Policy” on page 251

Configuring Device Settings

The Device Settings partial configuration lets you define device-specific configuration settings for a single SRS 5.x device. Alternatively, you can define these settings in the device Web console (refer to “Accessing the SRS Web Console from CMS” on page 61).

If you use automatic deployment, a Device Settings partial configuration is generated automatically for each auto-deployed device (refer to “Automatic Deployment of Devices” on page 253).

The following sections describe the configuration settings that can be defined in a Device Settings partial configuration:

- “Configuring Device Addresses” on page 88
- “Defining Communities” on page 90
- “Configuring Time Zone Settings” on page 92
- “Configuring the ARP Table” on page 93
- “Advertising Reduction Subnets” on page 94
- “Defining Outbound QoS Exclusions” on page 96
- “Adding Static Routes” on page 97
- “Configuring Router Polling” on page 99
- “Configuring Multi-Path Addresses” on page 101
- “Configuring the RADIUS Source Address” on page 103

Configuring Device Addresses

The Addresses page of the Device Settings partial configuration lets you specify the device's IP address, subnet mask, and default gateway, as well as add device and administrator contact information, and the DNS servers used to resolve IP addresses on the Traffic report.

To specify the network address and contact information:

1. In the Device Settings partial configuration window, click **Addresses** in the left-hand navigation frame and select the check box.

The screenshot shows a web browser window titled "https://192.168.5.28 - Device_1 - Microsoft Internet Explorer". The page is titled "Device_1" and "Device Settings Configuration". The left sidebar shows "Device Settings" with a list of options: ☒ Addresses, ☐ Time Zone, ☐ ARP, ☐ Reduction Subnets, ☐ Outbound QoS Exclusions, ☐ Static Local Routes, ☐ Dynamic Local Routes, and ☐ Multi-Path. Below this list, a note states: "If an item is not checked, the settings on the corresponding page will be determined by the Global Configuration. If an item is checked, the settings on the corresponding page will take precedence over the Global Configuration." The main content area is titled "Addresses" and contains a note: "Fields marked with an asterisk (*) are required. It is necessary to configure Domain Name and DNS servers only if you want to use DNS to resolve IP addresses for Traffic reports." The form includes the following fields: "IP Address*", "Subnet Mask*", "Default Gateway*", "Device Name", "Device Location", "Contact Information" (with a sub-note: "Enter name, phone number or email address of the person who will be supporting the device."), "Domain Name", and "DNS Servers" (with a sub-note: "Enter up to 3 IP addresses, one per line."). At the bottom are "Submit" and "Reset" buttons.

Figure 4-12 Configuring Network Address and Contact Information

2. Specify the following information:

IP address	Enter the IP address of the device. NOTE: If you change the IP address or subnet mask, you must reboot the device. To change the address of a registration server, you must first transfer the registration server to another device (refer to the <i>Sequence Reducer/Sequence Mirror Operator's Guide</i>).
Subnet mask	Specify the network portion of the IP address. For example, "255.255.255.0" indicates that the first 24 bits of the IP address are used for the network portion of the address.

Default gateway	Enter the IP address of the default router (must be on the same subnet as the device).
Device name	<p>Enter the device name (up to 30 characters) displayed in the banner of the Web console and in CLI prompts (default is the IP address). Do not use colons (:), asterisks (*) question marks (?) or angle brackets (< >) in device names.</p> <p>Device name changes are propagated to the other devices in the community the next time the device checks in with the registration server.</p>

3. Optionally, specify the following:

Device location	Enter a description of the device's physical location.
Contact information	Enter the contact information for the device administrator.
Domain name	<p>Enter the local DNS domain name of the device (up to 256 characters). The domain name must include at least one period, but not as the first or last character.</p> <p>When an IP address in the local domain is resolved by one of the specified DNS servers, the local domain name is prepended to the host name shown on the Traffic report.</p> <p>If this field is left blank, only the host names are shown for resolved IP addresses in the local domain. Resolved addresses outside the local domain include the domain name returned by the DNS server.</p>
DNS servers	Enter the IP addresses of up to three DNS servers (one per line) that can be used to resolve IP addresses on the Traffic report in the device Web console.

4. Click **Submit to enter the changes, or click **Reset** to discard them.**

Defining Communities

At least one device must be designated as a registration server. The registration server stores the network information for all devices that report to it, and identifies a community for each device. Each device contacts the registration server periodically to identify the other devices in the same community, and then attempts to form a reduction tunnel to each of those devices.

Since data reduction occurs only between devices in the same community, you can optimize performance in large deployments by limiting the number of devices in each community.

To configure the communities for a registration server:

1. In the Device Settings partial configuration window, click **Communities** in the left-hand navigation frame and select the check box.

NOTE: The Addresses section is selected automatically because the device address is needed to form the full name of the default community (“default-<IP address>”). Verify that the IP address, subnet mask, and gateway are defined in the Addresses section (refer to “Configuring Device Addresses” on page 88).

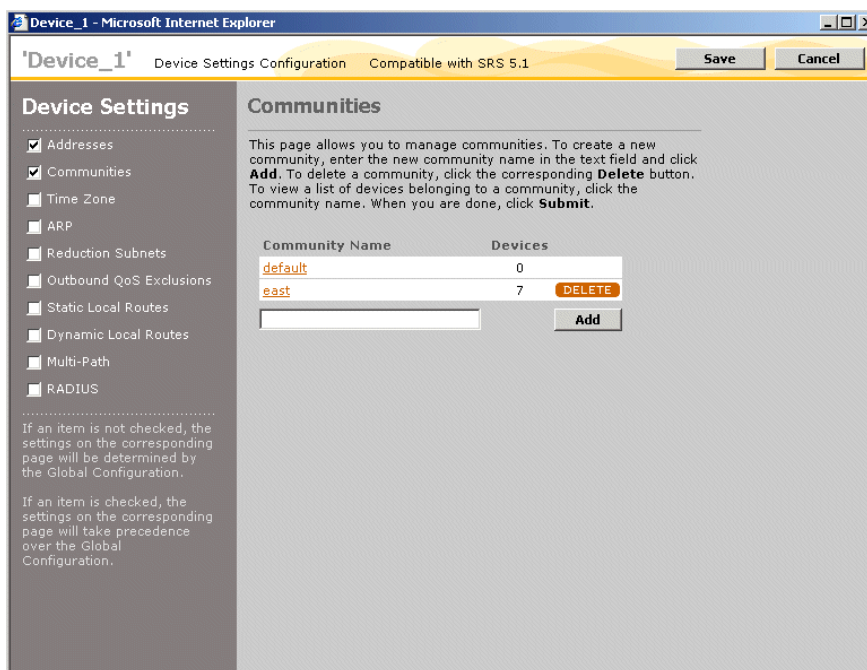


Figure 4-13 Defining Communities for a Registration Server

On the Communities page you can:

- Add a community. Enter a community name (up to 31 characters), and click **Add**.
 - Delete a community. Click **Delete** next to the appropriate community names. The devices in a deleted community are moved to the Default community if they do not belong to any other user-defined communities.
2. To define the devices in a community, click the community name, and click **Add/Remove Endpoints**.
- a. Select a community from the **Community** menu. The device name and IP address are shown for each device in the selected community. The IP address is enclosed in parentheses.
 - b. Select the devices you want to add to the current community, and click **Add**. To remove devices from the “Members of community” list, select the devices and click **Remove**.
 - c. Repeat Steps **a** and **b** for each community (a device can belong to multiple communities). When you download the configuration to a registration server, any devices that report to a different registration server are ignored.
 - d. If one or more devices you want to add are not listed for a community, you can add the devices manually. Click **Manual Entry**, enter the device IP addresses (one per line), and click **Submit**.
 - e. Click **Submit** to enter the changes, or click **Cancel** to discard them, and then click **Done** to return to the Communities page.

NOTE: As new communities are added to a registration server, they must be imported into CMS (refer to “Managing Communities” on page 328). CMS queries the registration server(s) each day and automatically incorporates any changes to the imported communities.

Configuring Time Zone Settings

The Device Settings partial configuration lets you specify a device's time zone and whether the device uses Daylight Savings Time. To specify a Network Time Protocol (NTP) server, refer to “Configuring NTP” on page 107.

NOTE: When you view reports in the device's time, the reported device times will be correct only if the time zone is set correctly.

To configure the time zone settings:

1. In the Device Settings partial configuration window, click **Time Zone** in the left-hand navigation frame and select the check box.

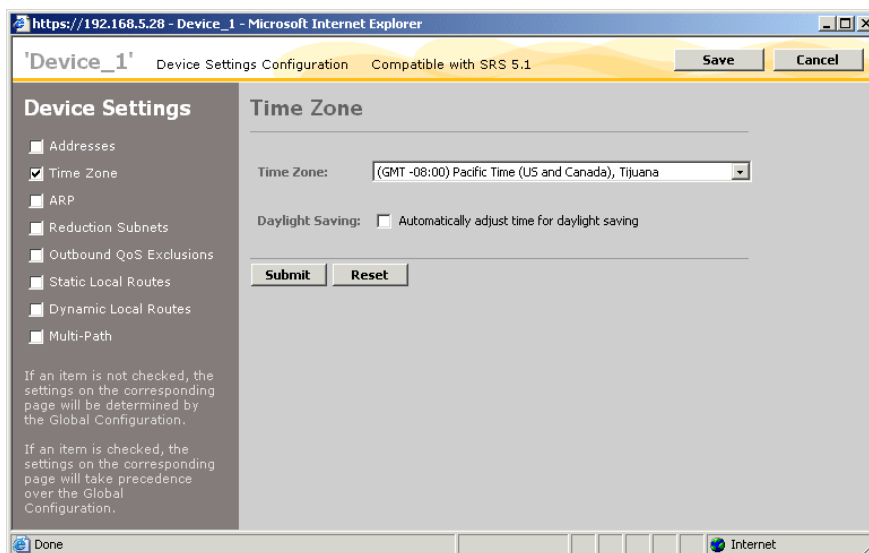


Figure 4-14 Configuring the Time Settings for a Device

2. Select the time zone of the device.
3. Select **Automatically adjust time for daylight savings**, if applicable.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring the ARP Table

The Address Resolution Protocol (ARP) is used to:

- Determine whether the gateway for a route is on the Local or Remote interface
- Discover the hardware (MAC) addresses of devices that are directly addressable on the Local and Remote interfaces

For devices that do not respond to ARP requests, you can add static ARP entries that map their IP addresses to their MAC addresses.

To add static entries to the ARP table:

1. In the Device Settings partial configuration window, click **ARP** in the left-hand navigation frame and select the check box.

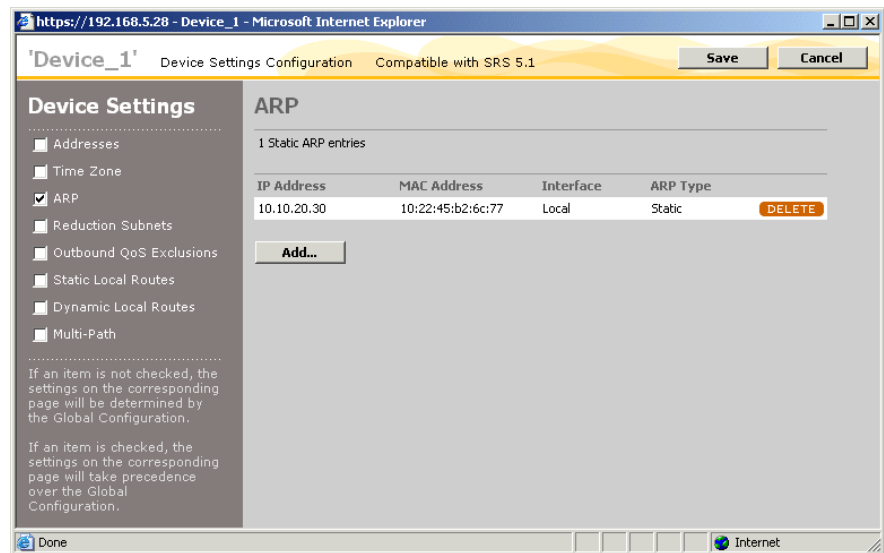


Figure 4-15 Viewing the ARP Table

2. To add one or more static ARP entries, click **Add**, enter the IP address and its associated MAC address, and select the **Local** or **Remote** interface. You can add up to five entries at one time.
3. Click **Submit** to enter the new entries, or click **Cancel** to discard them.
4. To delete an ARP entry, click **DELETE** next to the entry.

Advertising Reduction Subnets

Reduction subnets are the subnets on the LAN side of the device that you can selectively advertise to the other devices in the community. The other devices can then reduce and accelerate traffic sent to the advertised subnets. Initially, the only reduction subnet is the subnet where the device is installed. To enable dynamic discovery of LAN-side subnets, refer to “Configuring Dynamic Local Routes” on page 110.

The set of subnets advertised by each device is called a “netmap.” By default, only the subnets you select are advertised. You can enable the advertisement of all subnets or just selected subnets. In Figure 4-16, each device has two subnets on its LAN side.

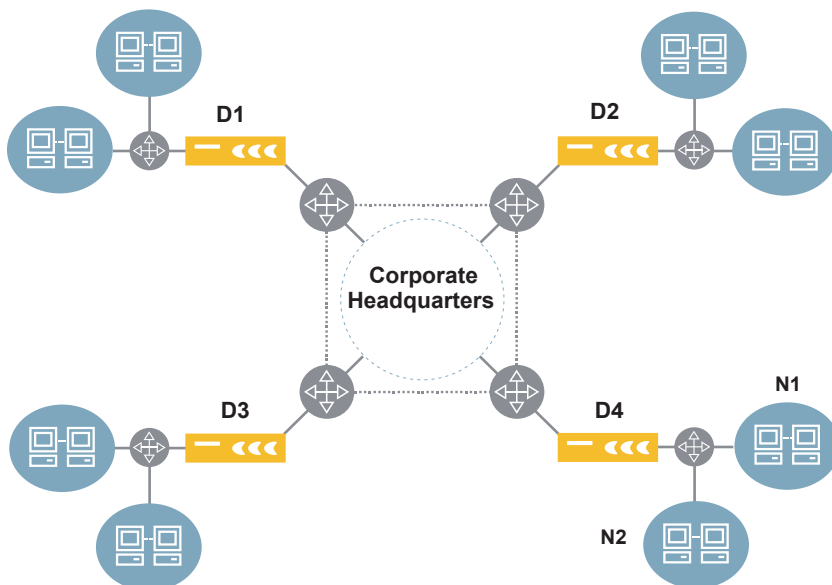


Figure 4-16 Selecting Specific Subnets for Data Reduction

If D4 advertises subnet N1, but not subnet N2, traffic destined for subnet N1 is reduced by the other devices and assembled by D4. However, traffic destined for subnet N2 passes through all devices without reduction.

You can also control reduction by application, as described in “Reducing Applications” on page 144 and by source/destination address, as described in “Configuring Source/Destination Filters” on page 216.

NOTE: If a host or gateway in an advertised subnet becomes unreachable, the device dynamically adjusts the advertised subnets to exclude (“carve out”) the unreachable address.

To advertise reduction subnets:

1. In the Device Settings partial configuration window, click **Reduction Subnets** in the left-hand navigation frame and select the check box.

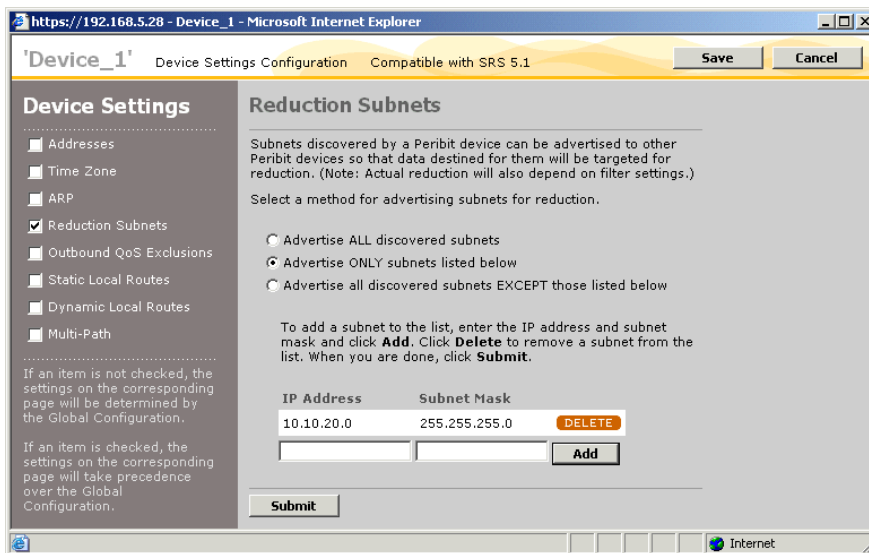


Figure 4-17 Configuring Reduction Subnets

2. Select one of the following parameters for the reduction subnet list:
 - **Advertise ALL discovered subnets.** All subnets discovered by the device are advertised.
 - **Advertise ONLY subnets listed below.** Only the specified subnets are advertised. For each subnet you want to advertise, enter the IP address and subnet mask, and click **Add**. To delete a subnet, click **DELETE**.
 - **Advertise all discovered subnets EXCEPT those listed below.** All discovered subnets are advertised, except the ones you specify.

NOTE: Be careful to advertise only the LAN-side subnets that the device can access. Do not use the ALL option if the device is installed off-path (refer to “Configuring Packet Interception” on page 220) or if the WAN reduction subnet option is enabled manually, such as in some VLAN environments. In these cases, all discovered LAN- and WAN-side subnets are eligible for advertisement.

3. Click **Submit** to enter the changes.

Defining Outbound QoS Exclusions

Each device can manage the outbound bandwidth for one or more remote devices (endpoints). If necessary, specific LAN/WAN address or subnet pairs can be excluded from bandwidth management.

NOTE: Traffic bursts between excluded addresses are unrestrained by QoS priority or bandwidth considerations, and may cause other traffic to be dropped by the router.

To exclude one or more LAN/WAN pairs of addresses or subnets from bandwidth management:

1. In the Device Settings partial configuration window, click **Outbound QoS Exclusions** in the left-hand navigation frame and select the check box.

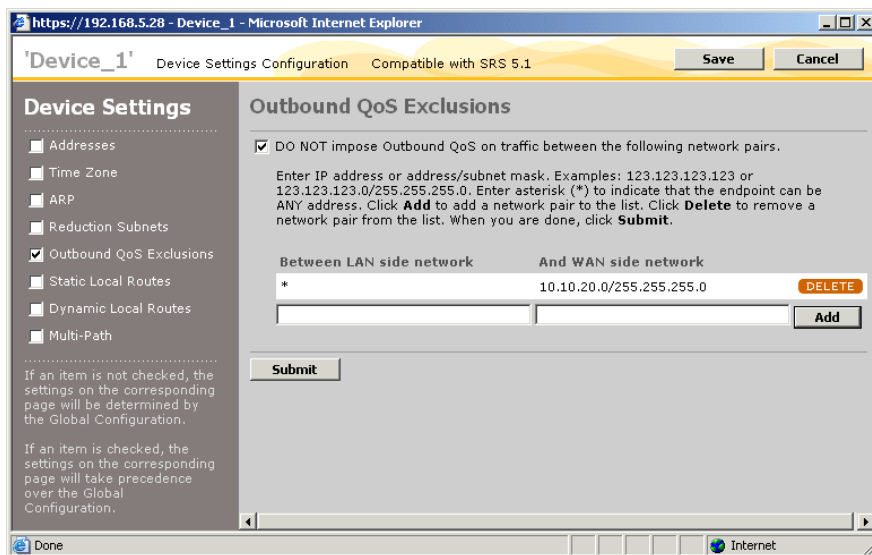


Figure 4-18 Excluding Subnets or Hosts from Bandwidth Management

2. Enter a local IP address or subnet in the **Between LAN side network** field, and enter a remote IP address or a “subnet/mask” in the **And WAN side network** field. Enter an asterisk (*) to indicate any address. Click **Add**.

To remove an entry, click **DELETE** next to the address pair.

If you specify any exclusions, you should also exclude all LAN traffic sent to the device’s local subnet. This ensures that the device manages only the traffic sent across the WAN, and not the traffic addressed to the router. If you do not specify any exclusions, by default each device excludes all LAN traffic sent to the local subnet.

3. Click **Submit** to enter the changes.

Adding Static Routes

Local routes are the routes defined in the device’s routing table. When you first install a device, the routing table contains the local subnet where the device is installed, a route to the default gateway (the default route), and the loopback address. To identify more routes, you can:

- Add static routes manually, as described here
- Add dynamic routes by enabling OSPF and/or RIP (v1 or v2), or by periodically polling the routing table of a Cisco router (refer to “Configuring Dynamic Local Routes” on page 110)
- Import a file of routes from an FTP server (refer to the *Sequence Reducer/Sequence Mirror Operator’s Guide*)

Each device can have a total of 8192 routes (static and dynamic).

If a subnet’s gateway is on the LAN side of the device (as determined by ARP), the subnet is added to the list of reduction subnets. Reduction subnets can then be advertised so that other devices in the community can reduce and accelerate traffic sent to those subnets (refer to “Advertising Reduction Subnets” on page 94).

To manually add static network routes:

1. In the Device Settings partial configuration window, click **Static Local Routes** in the left-hand navigation frame and select the check box.

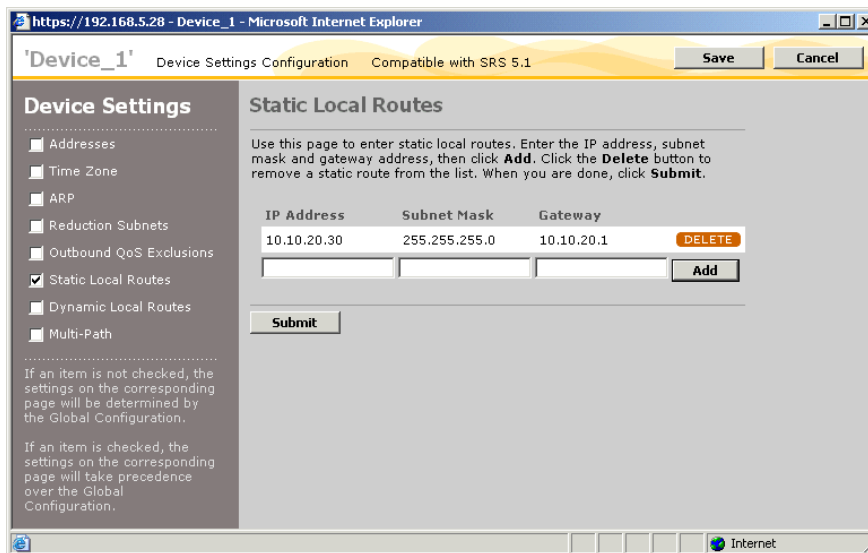


Figure 4-19 Adding a New Local Static Route

2. For each static route you want to add, enter an IP address, subnet mask, and a gateway address for the subnet, and click **Add**. To delete a static route, click **DELETE**.
3. Click **Submit** to enter the new routes.

When you load the configuration, the static routes defined here replace the static routes defined on the device (if any). Also, LAN-side static routes are added to the reduction subnets and advertised automatically to other devices, except when the WAN reduction subnets option is enabled (refer to “Advertising Reduction Subnets” on page 94).

Configuring Router Polling

You can configure a device to discover routes dynamically by periodically polling a Cisco router on the same subnet. All discovered routes are added to the device's routing table. The router must be configured to allow Remote Shell (*rsh*) access. Note that BGP routes are included only if you enable the BGP option using the “configure route-poll set allow-bgp-routes on” CLI command.

NOTE: You cannot poll a Cisco router from an off-path device that uses RIP for packet interception.

Configuring Route Polling

To enable route polling:

1. In the Device Settings partial configuration window, click **Dynamic Local Routes** in the left-hand navigation frame and select the check box.

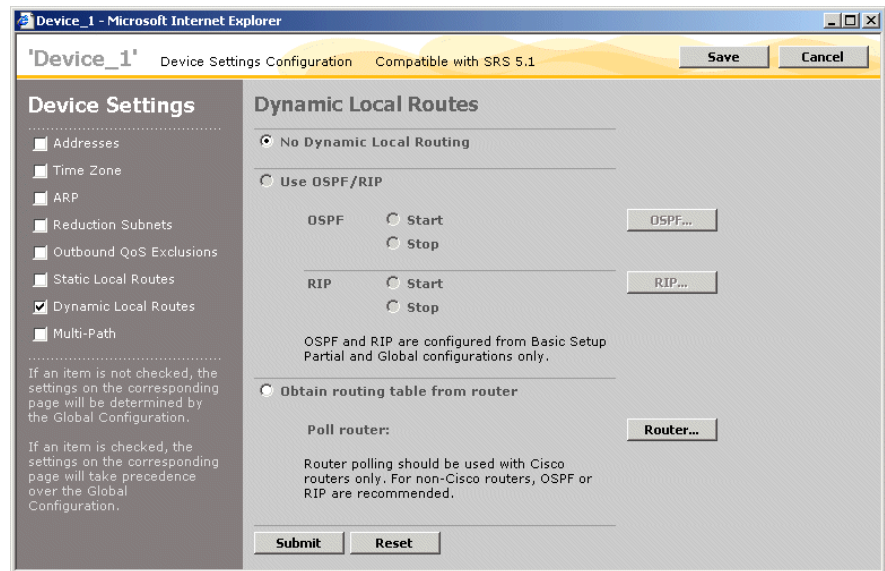


Figure 4-20 Enabling Router Polling

2. Click **Obtain routing table from router** and click **Router**.

3. Specify the following information:

Poll router	Enter the IP address of a Cisco router and the port number used for <i>rsh</i> (the standard port is 514). NOTE: The IP address must be on the same subnet as the device.
Secondary router	Enter the IP address and port of a secondary Cisco router to be used when the primary router is unavailable.
Local user name	Enter a local user name that matches the <i>remote</i> user name specified on the Cisco router.
Remote user name	Enter a remote user name that matches the <i>local</i> user name specified on the Cisco router.
Protocol interval	Enter a polling interval to indicate how often the Cisco router is polled for routing updates. The default is five minutes

- a. Click **Submit** to save the settings and return to the Dynamic Local Routes page.
- b. Click **Submit** to enter the changes, or click **Reset** to discard them.

If a subnet's gateway is on the LAN side of the device (as determined by ARP), the subnet is added to the list of reduction subnets. Reduction subnets can then be advertised so that other devices in the community can reduce and accelerate traffic sent to those subnets (refer to "Advertising Reduction Subnets" on page 94).

Configuring a Cisco Router for Route Polling

The following sample Cisco router commands enable Remote Shell access for the device at IP address 172.16.5.63. The local and remote user names are "lname" and "rname," respectively. On the Sequence Reducer or Sequence Mirror device, the names must be reversed (use "lname" as the remote name, and "rname" as the local name).

```
config terminal
ip rcmd rsh-enable
ip rcmd remote-host lname 172.16.5.63 rname enable
no ip rcmd domain-lookup
end
```

Configuring Multi-Path Addresses

If a pair of devices has two possible WAN paths between them, you can designate one path as the primary and the other as the secondary. You can then route application traffic to the primary or secondary path based on the performance requirements of the application and the actual performance of the path.

To use Multi-Path, you configure both devices so that outgoing packets intended for the secondary path are marked with a secondary source IP address and, optionally, with a specific gateway address or ToS/DSCP value. For more information about Policy-Based Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 235.

To specify a secondary IP and gateway addresses for Multi-Path:

1. In the Device Settings partial configuration window, click **Multi-Path** in the left-hand navigation frame and select the check box.

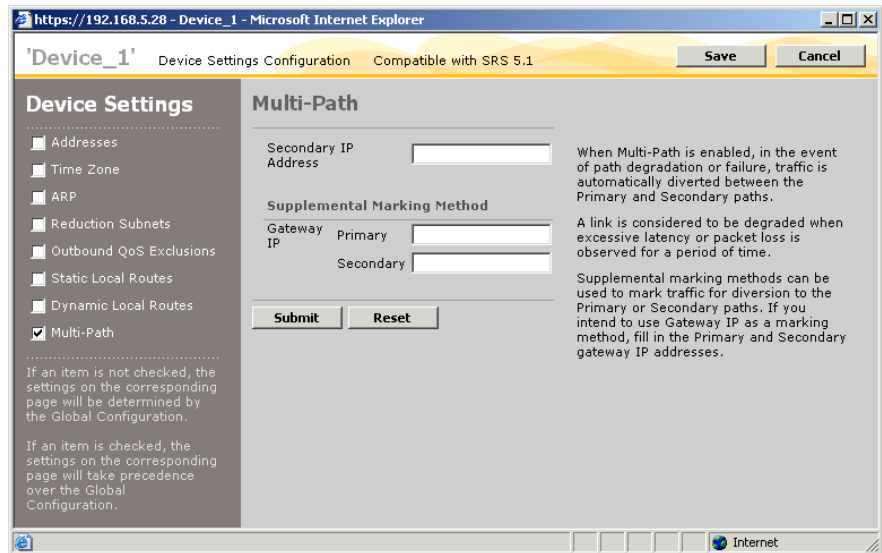


Figure 4-21 Multi-Path Secondary IP Address

2. Specify the following information:

Secondary IP Address	<p>Enter an IP address to be used as the source address on packets to be sent on the secondary path (packets sent on the primary path have the device address). The secondary IP address must be unique, and must be on the same subnet as the device address.</p> <p>Unless the WAN routers for the primary and secondary paths are also on this subnet (see Gateway IP below), the default gateway must be configured to route traffic with this source address to the appropriate WAN link (refer to “Configuring Routers to Support Multi-Path” on page 243).</p> <p>NOTE: If you enter an address assigned to another device, the path will remain inactive.</p>
Gateway IP	<p>If the WAN routers for the primary and secondary paths are on the same subnet as the Sequence Mirror or Sequence Reducer, and the SR or SM is connected to a Layer 2 switch (see Figure 4-22), enter the gateway IP addresses here.</p> <p>ARP is used to obtain the MAC addresses for the two gateways, and then traffic for the primary and secondary paths is marked with the MAC address of the appropriate gateway. In this case, no additional router configuration is needed.</p>

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

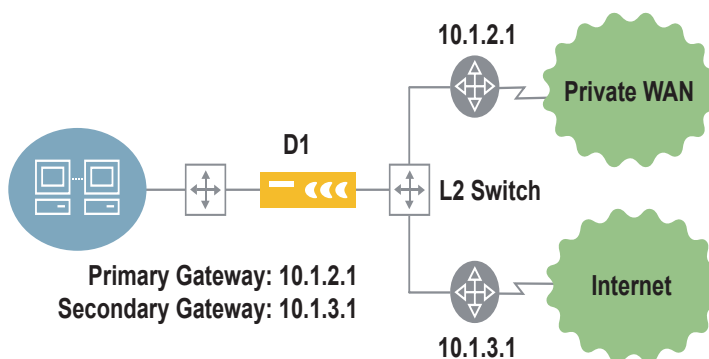


Figure 4-22 Multi-Path with Primary and Secondary Gateways

Configuring the RADIUS Source Address

The Device Settings partial configuration lets you specify an alternate source IP address for the RADIUS client. If you are using RADIUS servers to authenticate users (refer to “Defining RADIUS Servers” on page 121), replies from the RADIUS servers are sent to the specified source address. By default, all replies are sent to the device’s IP address.

To specify the RADIUS source address:

1. In the Device Settings partial configuration window, click **RADIUS** in the left-hand navigation frame and select the check box.
2. Enter the alternate IP address in the **Source IP Address** field.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Basic Setup Parameters

The following sections describe the basic setup configuration settings:

- “Configuring the Interface Settings” on page 104
- “Configuring NTP” on page 107
- “Enabling SNMP” on page 108
- “Enabling Syslog Reporting” on page 109
- “Configuring Dynamic Local Routes” on page 110
- “Enabling Route-Based Router Balancing” on page 112
- “Designating a Registration Server” on page 114
- “Generating NetFlow Records” on page 115

Configuring the Interface Settings

You can configure the two Network Interface Controllers (NICs) for the Local and Remote interfaces. By default, these interfaces are set to auto-negotiate the link speed and mode (half- or full-duplex).

NOTE: The SR-15, SR-20, SR-50, and SM-250 have two 10/100 NICs.
 The SR-55, SR-80, SR-100, and SM-500 have two 10/100/1000 NICs.
 The fiber SR-80 and SR-100 support only 1 Gigabit speeds at full-duplex.

The interface settings let you do the following:

- Manually configure the speed and mode of each interface.
- Enable high-availability support so that a failure detected on one interface is propagated to the other interface
- Enable 802.1Q VLAN support.

If you enable high-availability support, a failure detected on one interface causes the other interface to be turned off for 15 seconds. This allows the switch or router to detect the failure, and ensures that the routing mechanisms work as expected (Figure 4-23).

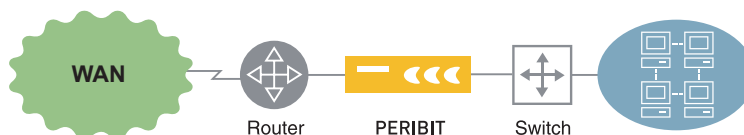


Figure 4-23 Using the High Availability Support Feature

- If the switch fails, the Remote interface is turned off so that the router detects the loss of connectivity with the switch.
- If the router fails, the Local interface is turned off so that the switch detects a loss of connectivity with the router.

On the SR-15, SR-80, and SR-100, you can also disable hardware passthrough so that the router detects the loss of traffic if the device fails (press the **Bypass Disable** button on the back panel).

To configure the interface settings:

1. In the Configuration window, click **Interfaces** in the left-hand navigation frame and select the check box.

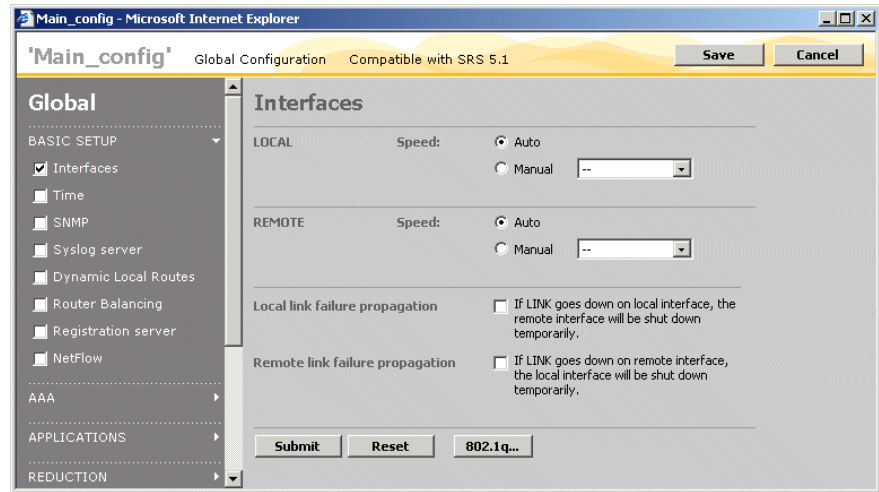


Figure 4-24 Configuring Interface Speed and Duplex Mode Settings

2. By default, the Local and Remote interfaces are set to auto-negotiate. To change the speed and mode for the Local or Remote interfaces, click **Manual**, and select a speed and mode setting (such as **100 half-duplex**).
3. Click the **Local link failure propagation** check box to disable the Remote interface when a switch failure is detected. Click the **Remote link failure propagation** check box to disable the Local interface when a router failure is detected. This allows the switch or router to detect the failure, and ensures that the routing mechanisms work as expected. After 15 seconds, the disabled interface is reactivated.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.
5. To enable reduction of VLAN traffic that conforms to the IEEE 802.1Q specification, click **802.1q**, select **Enable 802.1q**, and specify the following:
 - **Native VLAN ID.** Enter the default VLAN ID (1 through 4095) used for untagged frames in the VLAN environment where the device is installed.

- **VLAN ID.** Enter a VLAN ID (1 through 4095) for the port where the Local interface of the device is connected. On ports that have multiple VLANs, specify the VLAN that has the largest number of hosts. Note that the device resides on one VLAN, but can reduce traffic for all the VLANs.
- **Preserve VLAN ID on output packets.** Select the check box to preserve the VLAN ID in the header of reduced output packets if you have routers that use the VLAN ID for QoS, MPLS, or other functions.

6. Click **Submit** to enter the changes, or click **Reset** to discard them.

Note that when a device issues an ARP for a destination, only the router can respond with the appropriate VLAN tag. Since the router is on the WAN side, the local subnets appear to be WAN-side subnets and, by default, are excluded from the reduction subnets and cannot be advertised for reduction.

To allow WAN-side routes to be advertised for reduction, enter the following CLI commands on the device or in the CLI section of a global configuration or Advanced Setup partial configuration:

```
config reduction-subnet set wan-reduction-subnet on  
commit
```

Since both LAN and WAN-side subnets will be eligible for reduction, be sure to advertise only the true LAN-side subnets (refer to “Advertising Reduction Subnets” on page 94).

Configuring NTP

If your network uses the Network Time Protocol (NTP), you can specify a primary and secondary NTP server to synchronize your device times. If you do not have an NTP server, you can specify the IP address of the CMS server as your primary NTP server.

IMPORTANT: Using an NTP server is highly recommended if you poll the devices hourly for performance statistics. If a device is more than three minutes slow, its hourly data may not be counted in the correct hour, making the hourly reports inaccurate (reports for longer periods will be correct). To change the polling interval, refer to “Configuring Data Collection and Retention” on page 346.

To configure NTP servers:

1. In the Configuration window, click **Time** in the left-hand navigation frame and select the check box.

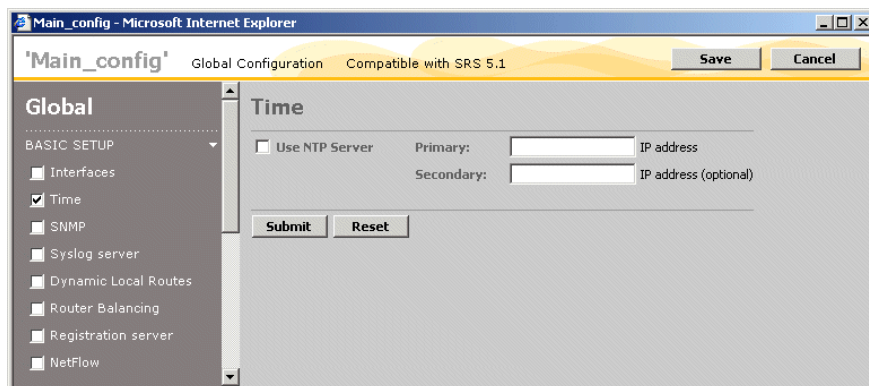


Figure 4-25 Configuring NTP

2. Select **Use NTP Server** and enter the IP address of the NTP server in the **Primary** field. Optionally, enter the address of a secondary NTP server to be used when the primary server is not available.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Enabling SNMP

The following SNMP support is provided:

- SNMP version 2
- Enterprise Management Information Base (MIB)
- MIB II, Interface Group public objects

NOTE: SNMPv2-compatible utilities are needed to query the 64-bit counters in the Enterprise MIB.

The Enterprise MIB can be used to view device performance statistics from a Network Management System (NMS). In addition, the devices can send SNMP traps to the NMS and other network devices. For a description of the SNMP traps, refer to Appendix B of the *Sequence Reducer/Sequence Mirror Operator's Guide*.

To enable SNMP:

1. In the Configuration window, click **SNMP** in the left-hand navigation frame and select the check box.

The screenshot shows the 'Main_config' web interface in Microsoft Internet Explorer. The left-hand navigation pane is expanded, showing 'Global' configuration options. Under 'Global', 'BASIC SETUP' is expanded, and 'SNMP' is selected with a checkmark. The main content area is titled 'SNMP'. It contains the following fields and controls:

- SNMP Enabled:** A checkbox labeled 'Yes' which is checked.
- Read Community String:** A text input field containing '*****'.
- Write Community String:** A text input field containing '*****'.
- Trap Enabled:** A checkbox labeled 'Yes' which is unchecked.
- Trap Community String:** A text input field containing '*****'.
- Trap Destinations:** A list box for entering IP addresses, currently empty.
- Authentication Trap Enabled:** A checkbox labeled 'Yes' which is unchecked.

At the bottom of the main area are 'Submit' and 'Reset' buttons. A note on the right side of the 'Trap Destinations' field says 'Enter IP addresses, one per line.'

Figure 4-26 Enabling SNMP

2. Select the **SNMP Enabled** check box to enable SNMP, and then enter the Read and Write Community Strings used by the NMS to access SNMP data. The defaults are **public** and **private**.

3. Select the **Trap Enabled** check box to generate SNMP traps (version 2 traps only). Next, enter a Trap Community String and the IP addresses (one per line) where the traps are sent. The default community string is **trap community**.
4. Select the **Authentication Trap Enabled** check box to generate traps for incorrect logins and unauthorized user access attempts.
5. Click **Submit** to enter the changes, or click **Reset** to discard them.

Enabling Syslog Reporting

Syslog messages can be sent to up to five Syslog servers. Syslog servers let you centrally log and analyze configuration events and system error messages, such as interface status, security alerts, and environmental conditions.

For a description of Syslog messages generated by CMS and the Sequence Mirror and Sequence Reducer devices, refer to Appendix B of the *Sequence Reducer/Sequence Mirror Operator's Guide*.

To enable Syslog reporting:

1. In the Configuration window, click **Syslog Server** in the left-hand navigation frame and select the check box.

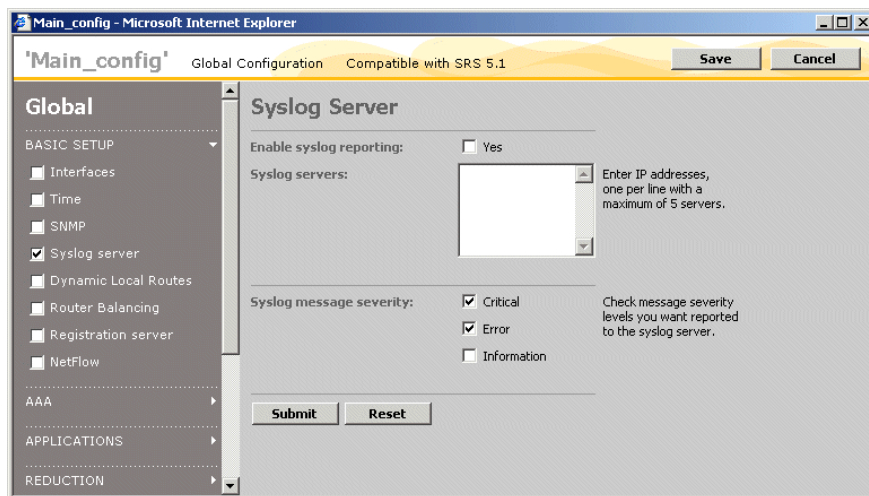


Figure 4-27 Enabling Device Syslog Reporting

2. Select the **Yes** check box to enable Syslog reporting, and then enter the IP addresses of up to five Syslog servers (one per line).

3. Select the severity levels of the messages sent to the Syslog server:
 - **Critical:** Critical error messages about software or hardware malfunctions.
 - **Error:** Error messages, such as License expired.
 - **Informational:** Informational messages, such as reload requests and low-process stack messages.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Dynamic Local Routes

If your network uses OSPF or RIP, you can enable these protocols to discover routes dynamically on the local and remote sides of each device. Alternatively, you can configure a device to periodically poll a Cisco router on the same subnet (refer to “Configuring Router Polling” on page 99).

A total of 8192 IP routes (static and dynamic) are supported (the SR-15 is limited to 1000). All discovered routes are added to the routing table on each device.

NOTE: If RIP or OSPF are enabled, routes added by ICMP redirects are ignored.

If a subnet’s gateway is on the LAN side of the device (as determined by ARP), the subnet is added to the list of reduction subnets. Reduction subnets can then be advertised so that other devices in the community can reduce and accelerate traffic sent to those subnets (refer to “Advertising Reduction Subnets” on page 94).

To configure RIP and/or OSPF:

1. In the Configuration window, click **Dynamic Local Routes** in the left-hand navigation frame and select the check box.

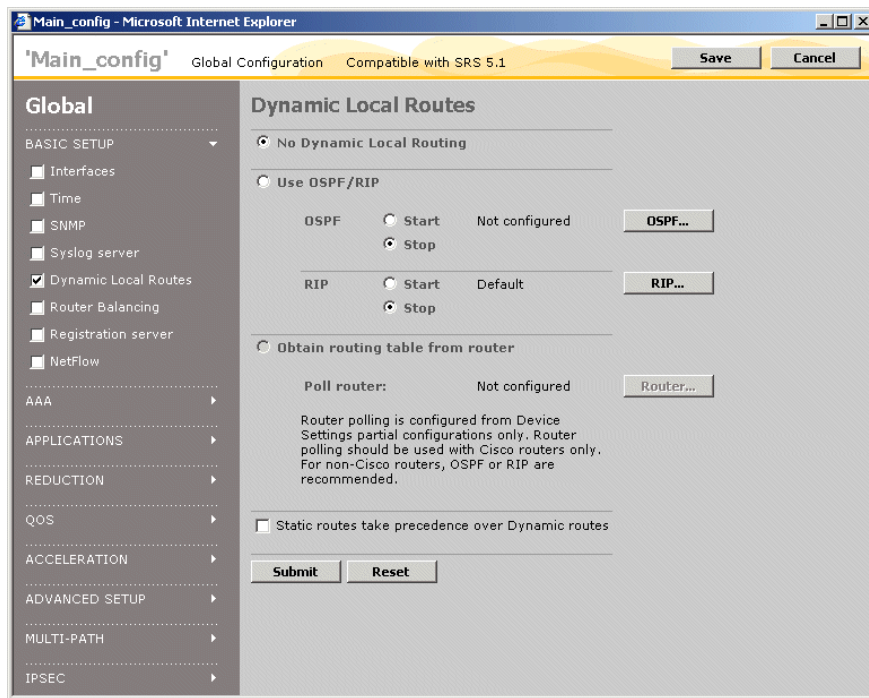


Figure 4-28 Configuring RIP and OSPF

2. To enable OSPF:
 - a. Click **OSPF...** and enter the Area ID for OSPF.
 - b. If your network uses OSPF authentication, select **Password** and enter the password (up to 8 characters), or select **MD5** and enter the key ID (0 to 255) and the MD5 key (up to 16 characters).
 - c. Click **Submit**.
 - d. Click **Use OSPF/RIP**, and select **Start** next to **OSPF**.
 - e. Click **Submit** to enter the changes, or click **Reset** to discard them.

3. To enable RIP:
 - a. Click **RIP...** and select the version of RIP used in your network (1 or 2).
 - b. If your network uses RIP authentication, select **Password** and enter the password (up to 15 characters).
 - c. Click **Submit**.
 - d. Click **Use OSPF/RIP**, and select **Start** next to **RIP**.
 - e. Click **Submit** to enter the changes, or click **Reset** to discard them.
4. By default, dynamic routes take precedence over static routes to the same destination. To give precedence to static routes, click **Static routes take precedence over Dynamic routes**, and click **Submit**.

Enabling Route-Based Router Balancing

For SRS 5.0 configurations, you can configure devices to balance the reduced traffic load across multiple routers that have equal-cost paths to the same destination (route-based balancing). To configure a router to distribute traffic based on ToS values set by a Sequence Mirror or Sequence Reducer (ToS marking for router-based balancing), refer to the “configure route” CLI command in the *Sequence Reducer/Sequence Mirror Operator’s Guide*.

Using route-based balancing, reduced traffic can be distributed across up to four different gateways. In Figure 4-29, device D1 identifies two gateways that have equal cost paths to the network (N2) advertised by D2. D1 can use the two gateways on a per-destination, per-packet (round-robin), or per-flow basis.

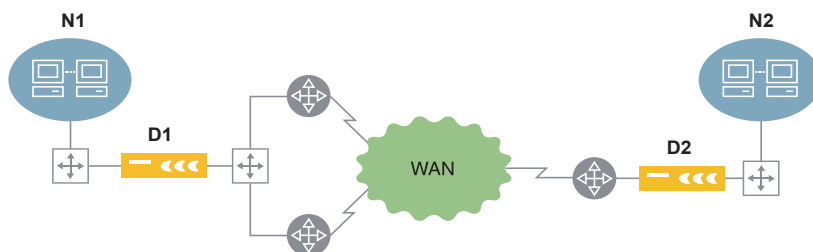


Figure 4-29 Configuring Router Balancing Policies

To identify gateways (up to four) that have equal cost paths to the same IP address, open the device Web console for a device and click **Local Routes**. Equal cost paths are grouped together in the Local Routes page (Figure 4-30).

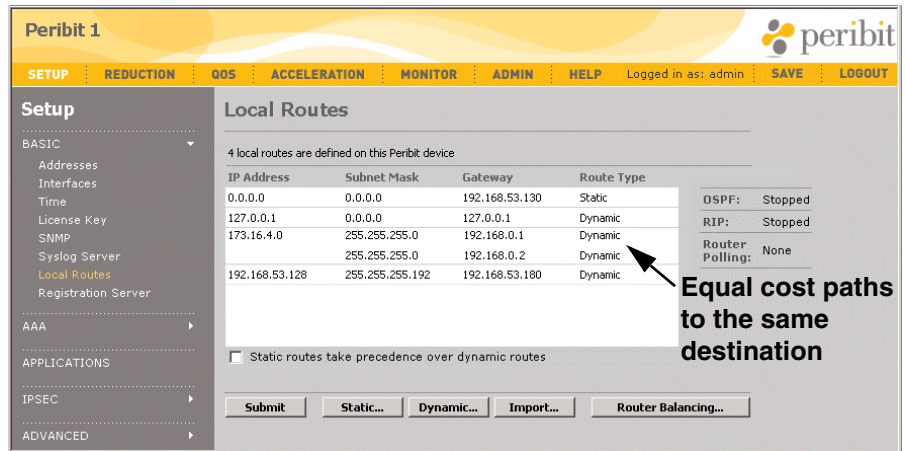


Figure 4-30 Common Routes with Equal Cost Paths

To enable route-based router balancing:

1. In the Configuration window, click **Router Balancing** in the left-hand navigation frame and select the check box.

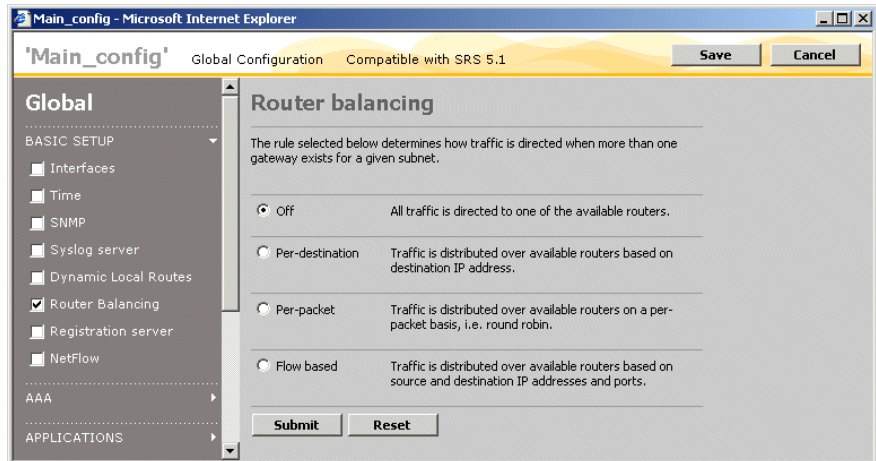


Figure 4-31 Configuring Route-Based Router Balancing

2. Select one of the following router balancing policies:
 - **Off.** (Default) All traffic is directed to one of the available routers. No balancing.
 - **Per-destination.** Traffic is distributed over available routers based on destination IP address.
 - **Per-packet.** Traffic is distributed over available routers on a per-packet basis (round robin).

NOTE: Packets that lack port information, such as ICMP and fragmented packets, are sent to the first gateway, and are not balanced according to the per-packet scheme.

- **Flow based.** Traffic is distributed over available routers based on source and destination IP addresses and ports.

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Designating a Registration Server

A registration server is a device that stores the network information for all the other devices that report to it. Each device contacts the registration server periodically to identify the other devices in the same community. CMS queries the registration server once a day to obtain the latest network information for each device.

In global configurations and Basic Setup partial configurations, a registration server address and password must be specified if the **Registration Server** check box is selected. If you change the password defined on a registration server, you can update the password in CMS, and download the new password to all devices (refer to “Managing Communities” on page 328).

NOTE: If the registration server address is incorrect, any device where you load the configuration will lose access to the other devices in the community, and CMS will lose access to the device within 24 hours.

To specify the registration server:

1. In the Configuration window, click **Registration Server** in the left-hand navigation frame and select the check box.



Figure 4-32 Designating a Registration Server

2. Specify the IP address and password of the registration server. The password must match the one defined on the registration server.

When you save or download a configuration, an error occurs if the password does not match the current or previous password stored in CMS for the specified registration server (refer to “Managing Communities” on page 328).

3. Optionally, can click **Use IP address** and enter the IP address of the secondary (backup) registration server.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Generating NetFlow Records

You can configure a device to send its Top Traffic statistics to a Cisco NetFlow server. Each device collects traffic statistics for the most active traffic flows, including the protocol, source and destination addresses and ports, and the number of packets and bytes sent and received.

For SRS 5.0 devices, if the collected statistics are sent to a Cisco NetFlow server, they cannot be displayed in Top Traffic report in the Web console. NetFlow data is sent in Version 5 format, as described in “NetFlow Version 5 Export” on page 363.

To generate NetFlow records:

1. In the Configuration window, click **NetFlow** in the left-hand navigation frame and select the check box.

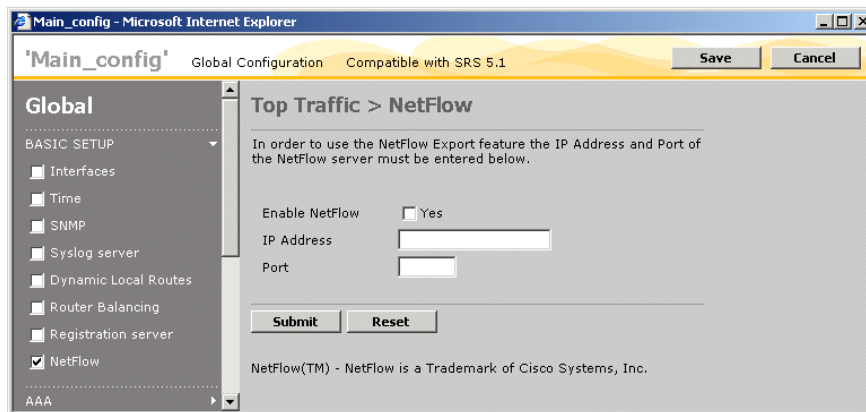


Figure 4-33 Generating NetFlow Records

2. Click **Enable NetFlow**, and enter the IP address and port number of a NetFlow server.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring AAA Settings

AAA stands for authentication, authorization, and accounting. Authentication verifies a user's identity, such as by user name and password or a challenge/response mechanism. Authorization provides access control, such as privilege level assignment and timeout enforcement. Users must be authenticated before they can be authorized. Accounting collects and sends auditing information, such as user traffic statistics and connection times.

Users can be authenticated and authorized using a local database or a remote RADIUS server. RADIUS allows the device to be integrated with existing authentication infrastructures such as Active Directory, NT Domain, LDAP Meta-Directories, and most Token Card and SmartCard servers. The RADIUS server provides the connection to the back-end authentication infrastructure, and existing user entries in the directory can be used for authentication and authorization.

Each device is a standard RFC 2138-compliant RADIUS client. For RADIUS servers that require a client type to be specified, choose the option for a standard client and standard RADIUS dictionary. Two standard RADIUS authorization attributes are supported:

- **Attribute 6: Service-Type.** Indicates a user's access privileges. The valid service types are Administrative (6) and NAS-Prompt (7). Administrative (6) grants read-write access, and NAS-Prompt (7) grants read-only access.
- **Attribute 28: Idle-Timeout.** Indicates the number of consecutive seconds a user session can be idle before the connection is closed.

Multiple RADIUS servers can be configured for redundancy. You can use both the local database and RADIUS, so that some users are authenticated locally and others are authenticated through RADIUS.

The following sections describe the AAA configuration settings:

- “Selecting Authentication Methods” in the next section
- “Enabling Authorization Checking” on page 120
- “Defining RADIUS Servers” on page 121
- “Defining Local Users” on page 123
- “Securing Operator Access” on page 125
- “Securing Front Panel Access” on page 126

Selecting Authentication Methods

You can specify the order in which a device's local database and RADIUS server groups are accessed to authenticate users on the Web, the SSH (CLI), and the console port. You can also specify the number of SSH login attempts allowed before a user is locked out. By default, all users are authenticated locally.

To define RADIUS servers and server groups, refer to “Defining RADIUS Servers” on page 121. To define user accounts locally, refer to “Defining Local Users” on page 123.

To select the authentication methods for each user interface:

1. In the Configuration window, click **AAA** in the left-hand navigation frame, click **Authentication**, and select the check box.

'Main_config' Global Configuration Compatible with SRS 5.1 Save Cancel

Global

- BASIC SETUP
 - Interfaces
 - Time
 - SNMP
 - Syslog server
 - Dynamic Local Routes
 - Router Balancing
 - Registration server
 - ☒ NetFlow
- AAA
 - ☒ Authentication
 - Authorization
 - RADIUS
 - ☒ Local Users
 - Operator Access
 - Front Panel Access
- APPLICATIONS
- REDUCTION

Authentication

Authentication methods are evaluated in order until one responds with a 'pass' or 'fail'. When a method responds, the evaluation is considered final and no other methods are used.

There is one exception to this rule. If the first method is set to 'Local' and the second method is 'RADIUS', then if the Local method does not find a username entry in the local database, instead of issuing a 'fail', the RADIUS method will be used.

The 'Local' method cannot be immediately followed by the 'None' method.

Console	Order	Method
	1	Local
	2	--Select a method--
	3	--Select a method--
	4	--Select a method--

SSH	Order	Method
	1	Local
	2	--Select a method--
	3	--Select a method--
	4	--Select a method--

Disconnect user ☒ After 3 failed attempts ☐ Never

Web	Order	Method
	1	Local
	2	--Select a method--
	3	--Select a method--
	4	--Select a method--

Submit Reset

Figure 4-34 Selecting Authentication Methods

2. Specify the following information:

Console	<p>Select up to four authentication methods for users logging in through a terminal connected to the console port. The options are:</p> <ul style="list-style-type: none"> • RADIUS: <i>group_name</i>. Attempts to authenticate users by accessing the RADIUS servers in the specified group. The servers are accessed in the order specified by the group. If all RADIUS servers are down or do not respond, the next method is tried. • Local. Attempts to authenticate users locally. • None. Login not required. Can be used alone or after the last RADIUS group. Cannot be used directly after Local. <p>Each method is tried in the order specified. Authentication stops with the first success or failure. However, if Local is the first method, the next method is tried if the user is not defined locally.</p>
SSH	<p>Select up to four authentication methods for users logging in using the SSH protocol. Same options as the console, except that None is not available (authentication is required).</p> <p>Select the number of unsuccessful SSH login attempts allowed before a user is disconnected (1 to 10) or select Never.</p>
Web	<p>Select up to four authentication methods for users logging in through the Web. Same options as the console, except that None is not available (authentication is required).</p>

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Enabling Authorization Checking

By default, all authenticated users have read-write access and a 30-minute idle timeout. If you create read-only user accounts or change the default idle timeout, either in RADIUS or in the local user database, you must enable authorization checking for the changes to take effect.

To enable or disable authorization checking:

1. In the Configuration window, click **AAA** in the left-hand navigation frame, click **Authorization**, and select the check box.

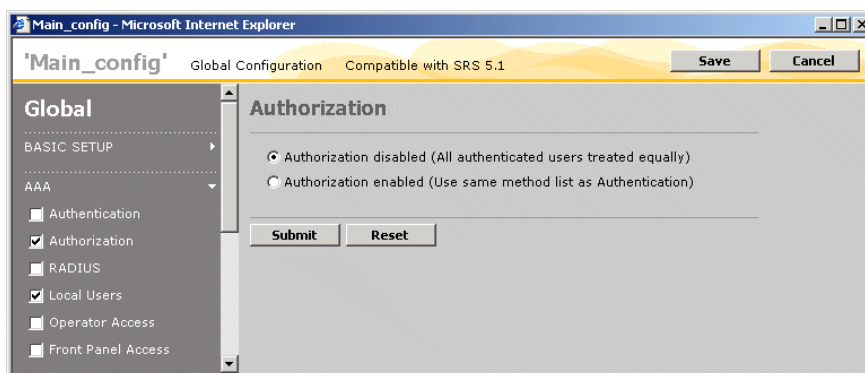


Figure 4-35 Enabling Authorization Checking

2. Select one of the following.
 - **Authorization disabled.** All users have read-write privileges and a 30-minute idle timeout.
 - **Authorization enabled.** User privilege level specified by authentication method. If RADIUS is used for authentication, but does not specify a privilege level or an idle timeout, all users have read-write privileges and a 30-minute idle timeout.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Defining RADIUS Servers

You can use RADIUS servers to authenticate users by defining one or more RADIUS servers and assigning them to at least one server group. The servers in each group are accessed in the order specified. You can define up to four groups of five servers (the same server can appear in multiple groups).

To specify the server groups used for authentication, refer to “Selecting Authentication Methods” on page 117.

To define RADIUS servers and server groups:

1. In the Configuration window, click **AAA** in the left-hand navigation frame, click **RADIUS**, and select the check box.

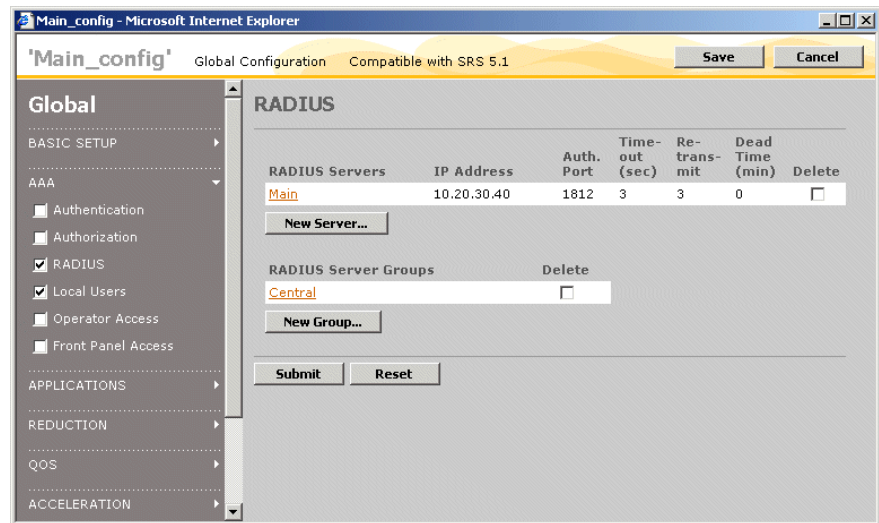


Figure 4-36 Defining RADIUS Servers and Server Groups

From the RADIUS page, you can:

- Add new servers and assign them to groups, as described in Step 2 and Step 3.
- Change a server or server group. Click the server or group name, make any needed changes, and click **Submit**.
- Delete servers or groups. Select the check box next to the servers and groups you want to delete, and click **Submit**. Deleting a server group does not delete the associated servers.

2. To add a new server, click **New Server**, specify the following information, and click **Submit**:

Server Name	Enter the RADIUS server name (up to 32 characters).
IP Address	Enter the IP address of the server.
Authentication Port	Enter the UDP port number used for authentication (default is 1812).
Timeout	Enter the number of seconds (1 to 65535) that the device waits for the server to respond.
Retransmit	Enter the number of times (1 to 100) that requests are retransmitted to a server before trying the next server in the group (if any).
Dead Time	If the server fails to respond to all retransmissions, enter the number of minutes (0 to 1440) that the device waits before trying to access the server again.
Shared Secret Key	Enter the secret key (up to 31 characters) used to access the server. The same key must be configured on the RADIUS server.

3. To add a new server group, click **New Group**, specify the following information, and click **Submit**:

RADIUS Group Name	Enter the server group name (up to 32 characters).
RADIUS Servers	Select the RADIUS servers in the group (up to five). The servers are accessed in the order specified. For example, if the first server does not respond, the second server is accessed.

Defining Local Users

You can define up to 25 users that can be authenticated locally by each device. Each user can have full (admin) or read-only access privileges. The default password (**peribit**) of the predefined **admin** account must be changed for all new global configurations and for new AAA partial configurations where the **Local User** check box is selected.

To specify how users are authenticated (locally and/or through RADIUS), refer to “Selecting Authentication Methods” on page 117.

To define local user accounts:

1. In the Configuration window, click **AAA** in the left-hand navigation frame, click **Local Users**, and select the check box.

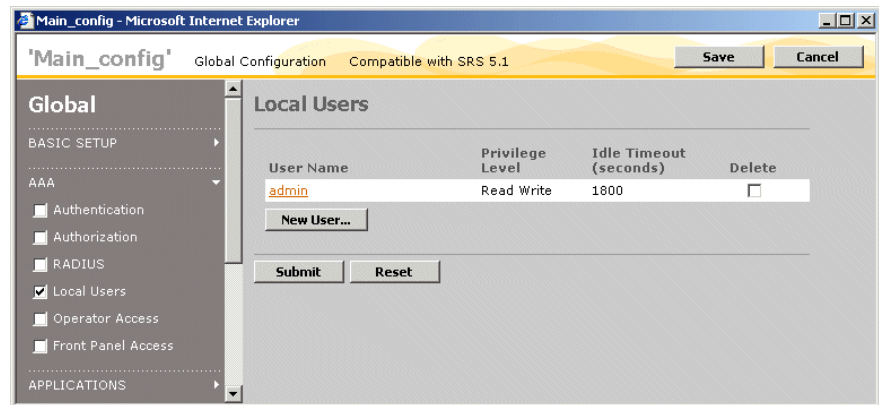


Figure 4-37 Defining Local Users

2. To add a new account, click **New User**, specify the following information, and click **Submit**:

User Name	Enter the account name (up to 32 characters).
Privilege Level	Select administrator (read-write) or read-only privileges.
Idle Timeout	Enter the number of consecutive minutes of inactivity before a user is logged out (the default is 30), or select Never .
Password	Enter the password twice (from 4 to 64 characters).

NOTE: Authorization checking is disabled by default, so that all authenticated users have read-write access and a 30-minute idle timeout. If you create read-only user accounts or change the default idle timeout, you must enable authorization checking (refer to “Enabling Authorization Checking” on page 120)

3. To change a user account, click the user name, make any needed changes, and click **Submit**.
4. To delete user accounts, select the check box next to the accounts you want to delete, and click **Submit**.

Securing Operator Access

You can create an Include or Exclude list to allow or deny access to a device from specific IP addresses or subnets. For example, if you enter one address in the Include list, administrative users can log in only from the specified address. Alternatively, if you enter an address or subnet in the Exclude list, access to the device from that address or subnet is denied.

To restrict operator access:

1. In the Configuration window, click **AAA** in the left-hand navigation frame, click **Operator Access**, and select the check box.

The screenshot shows the 'Main_config' web interface in Microsoft Internet Explorer. The title bar indicates 'Main_config - Microsoft Internet Explorer'. The page has a yellow header with 'Global Configuration' and 'Compatible with SRS 5.1'. There are 'Save' and 'Cancel' buttons in the top right. The left sidebar is titled 'Global' and contains a navigation menu with the following items: BASIC SETUP, AAA (expanded), AUTHENTICATION, AUTHORIZATION, RADIUS, Local Users (checked), Operator Access (checked), Front Panel Access, APPLICATIONS, REDUCTION, QOS, ACCELERATION, ADVANCED SETUP, MULTI-PATH, and IPSEC. The main content area is titled 'Operator Access' and contains the following text: 'The following lists are used to restrict operator access to this Peribit device from designated valid client addresses only. If both lists are empty, then operator access is unrestricted. If an address/subnet is entered in the Include list, then all other addresses/subnets are denied access.' Below this is an example of IP addresses: 'Example: 123.123.123.123 123.123.123.123/255.255.255.0'. There is also a note: 'Also, if you want to preserve the changes, you must save the configuration to flash memory using the 'Save Configuration' page under the 'Maintenance' tab after submit.' Below the text are two text input fields: 'Include list' and 'Exclude list'. The 'Include list' field has a description: 'Enter addresses/subnets which should have access to this Peribit device, one per line.' The 'Exclude list' field has a description: 'Enter addresses/subnets which should be denied access to this Peribit device, one per line.' At the bottom of the main content area are 'Submit' and 'Reset' buttons.

Figure 4-38 Configuring Device Operator Access

2. To allow access to a device only from specific IP addresses or subnets, enter the addresses or subnets in the **Include list** (one per line).

The subnet format is:

<IP address>/<subnet mask>

All other client IP addresses are denied access to the device.

3. To deny access to a device only from specific IP addresses or subnets, enter the addresses or subnets in the **Exclude list** (one per line).

NOTE: IP addresses in both the Include and Exclude lists are denied access.

4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Securing Front Panel Access

You can lock the front-panel of a device to prevent anyone from rebooting the device or making configuration changes through the front panel keypad.

NOTE: The SR-15, SR-20, and SM-250 do not have a front-panel keypad. Also, locking the front panel on the SR-100 does not lock the front panels of the client devices.

To lock the front panel keypad:

1. In the Configuration window, click **AAA** in the left-hand navigation frame, click **Front Panel Access**, and select the check box.

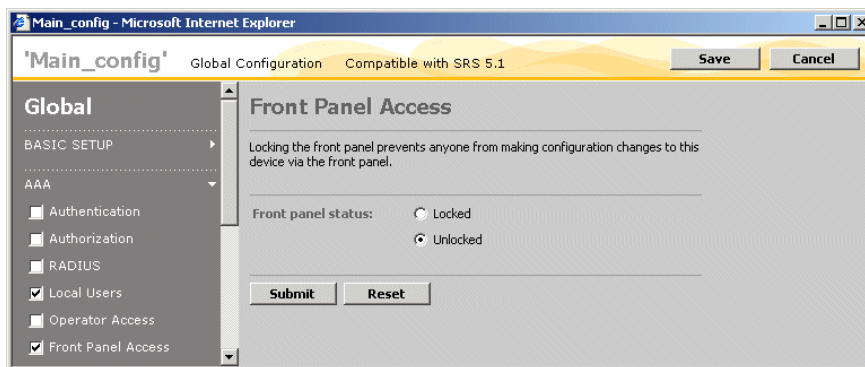


Figure 4-39 Securing Front Panel Access

2. To lock front-panel access, select **Locked**.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Application Settings

Application definitions allow each device to identify the traffic of up to 256 applications (the SR-15 is limited to 100). Definitions are provided for applications with well-known port numbers. All other applications are grouped together as “Undefined” or “Others”.

If you add new application definitions to a global configuration, the applications are included in the Reduction, Acceleration, and QoS sections of the configuration, where you can:

- Enable or disable data reduction, as described in “Reducing Applications” on page 144.
- Enable or disable monitoring of applications for reports, as described in “Monitoring Applications” on page 138.
- Accelerate an application’s traffic (if data reduction is enabled), as described in “Configuring Traffic Acceleration” on page 190.
- Assign the application to a traffic class to manage its outbound bandwidth allocation, as described in “Assigning Applications to Traffic Classes” on page 136. Traffic classes are also used for path optimization, as described in “Configuring Policy-Based Multi-Path” on page 235.

NOTE: Traffic classes are defined under Applications in an SRS 5.1 global configuration, and under QoS in an SRS 5.0 global configuration.

New (or changed) applications also appear in any Reduction, Acceleration, QoS, or Multi-Path partial configurations that reference the global configuration. Similarly, new definitions added to an Applications partial configuration are included in the partial configurations that reference it.

Default Application Definitions for SRS 5.1 Devices

Table 4-4 lists the default application definitions for SRS 5.1 devices. Each definition has rules to match any traffic that has the specified port number(s) as the source or destination.

Table 4-4 Default Application Definitions

Application	Precedence	Port Numbers
AOL	36	5190-5193
CIFS	6	139, 445
Clearcase	23	371
CVS	33	2401
DNS	15	53
Exchange	20	135
		Note: Port 135 is the startup port; other ports are learned dynamically. This definition applies only to Exchange traffic for Windows clients, not Web clients.
Filenet	40	32768-32774
FTP	1	20-21
		Note: Non-default FTP ports are learned dynamically.
Groupwise	29	1677
Hostname Resolution	21	42
HTTP	4	80, 8080
HTTPS	12	443
ICA (Citrix)	9	1494
ICMP	42	Protocol 1 (no ports specified)
Kerberos	17	88
LDAP	16	389
Lotus Notes	7	1352
Mail	3	25,110,143
MS Streaming	30	1755
MS Terminal Services	18	3389
NetApp SnapMirror	39	10566

Table 4-4 Default Application Definitions

Application	Precedence	Port Numbers
NetBios	5	137, 138
NFS	32	2409
Novell NCP	27	524
Oracle	11	1525
PCAnywhere	37	5631-5632
Printer	26	515
RADIUS	31	1812, 1813
RTSP	28	554
SAP	35	3300-3388,3390-3399,3600-3699,3200
Shell	24	514 TCP
SNMP	19	161-162
SNTP	14	123
SQL Server	8	1433
SSH	13	22
Sybase	10	1498
Symantec Anti-Virus	34	2967
Syslog	25	514 UDP
TACACS	22	49
Telnet	2	23
Traceroute	41	33434-33534 UDP
XWindows	38	6000-6063

Viewing the Application Overview

In a global configuration, the Application Overview page shows each application's traffic class, and whether reduction, acceleration, and monitoring are enabled for each application. In an Applications partial configuration, only the traffic class and monitoring status are shown for each application.

To view the application overview:

1. In the Configuration window, click **APPLICATIONS** in the left-hand navigation frame, and click **Overview**.

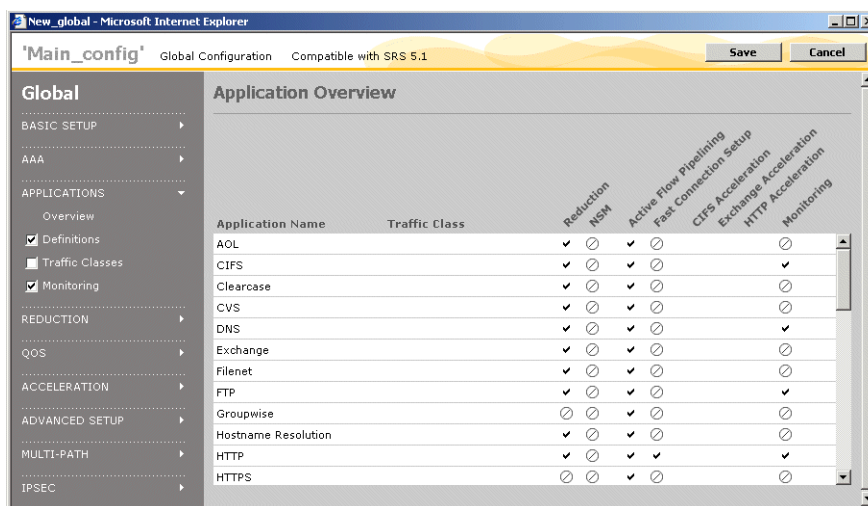


Figure 4-40 Application Overview Page

2. The Application Overview page displays the following information (check marks indicate the enabled features):

Traffic Class	Traffic class assigned to the application. To change the traffic class, refer to “Assigning Applications to Traffic Classes” on page 136.
Reduction	Indicates whether the application’s traffic is reduced (refer to “Reducing Applications” on page 144.).
NSM	Indicates whether Network Sequence Mirroring is used for data reduction (refer to “Reducing Applications” on page 144). NSM requires a hard disk, and applies only to Sequence Mirror devices.
Fast Connection Setup	Indicates whether the application’s traffic is accelerated using Fast Connection Setup (refer to “Enabling Fast Connection Setup by Application” on page 204).

Active Flow Pipelining	Indicates whether the application's traffic is accelerated using Active Flow Pipelining (refer to "Enabling Active Flow Pipelining by Application" on page 203).
CIFS Acceleration	Indicates whether CIFS traffic for the application is accelerated (refer to "Enabling Microsoft CIFS Acceleration" on page 205).
Exchange Acceleration	Indicates whether Exchange traffic for the application is accelerated (refer to "Enabling Microsoft Exchange Acceleration" on page 208).
HTTP Acceleration	Indicates whether HTTP traffic for the application is accelerated (refer to "Enabling HTTP Acceleration" on page 210).
Monitoring	Indicates whether you can view statistics for the application, as described in "Monitoring Applications" on page 138.

Configuring Application Definitions

Each application definition can have up to 10 rules, and each rule can specify a protocol, source and destination port numbers (or range of port numbers), source and destination IP addresses or subnets, a ToS/DSCP value, and a URL or a Citrix client and application name.

A packet matches an application definition if a match occurs on any of its rules. All the values defined in the same rule must be true for a match to occur on that rule. A packet is classified under the first application for which a rule match is found. Packets are compared against the definitions according to the order number (definitions with the lowest order numbers are checked first). The comparison stops on the first match, so if two definitions are similar, the more specific definition must have a lower order number.

NOTE: In the device Web console, you can add new definitions by selecting undefined applications from the Top Traffic report, as described in the *Sequence Reducer/Sequence Mirror Operator's Guide*. You can then extract the configuration settings from the device (refer to "Extracting Configurations" on page 75).

To add or change application definitions:

1. In the Configuration window, click **APPLICATIONS** in the left-hand navigation frame, click **Definitions**, and select the check box.

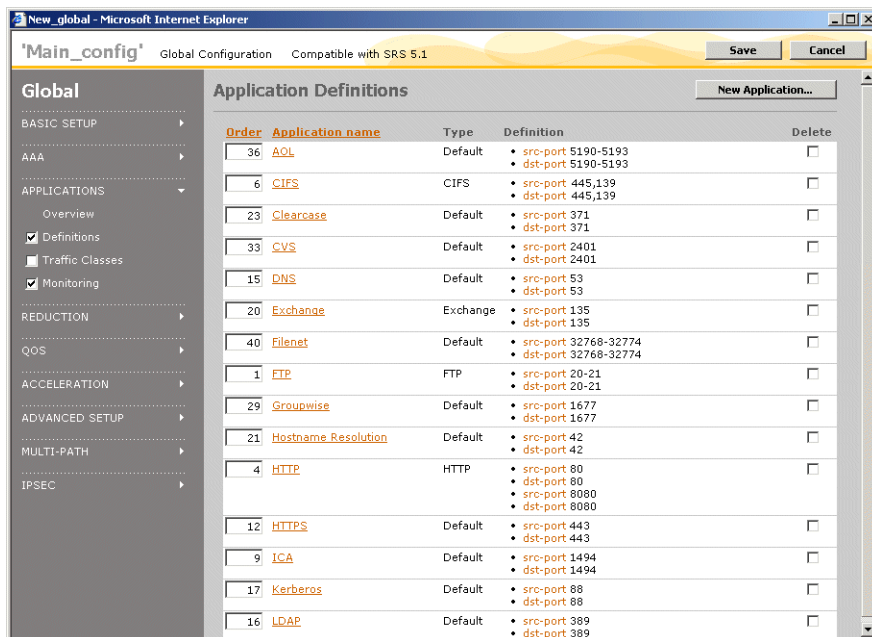


Figure 4-41 Application Management Page

From the Application Management page, you can:

- Add a new application definition, as described in Step 2 through Step 5.
- Change an application definition. Click the application name, make any needed changes, and click **Submit**.
- Change a definition's order number. Type a new value in the **Order** field, and click **Submit** to renumber the definitions. The new value cannot exceed the highest value in the current range. The definitions are compared against the traffic in ascending order.
- Delete application definitions. Select the check box next to the applications you want to delete, and click **Submit**.

2. To add a new application definition, click **New Application**.

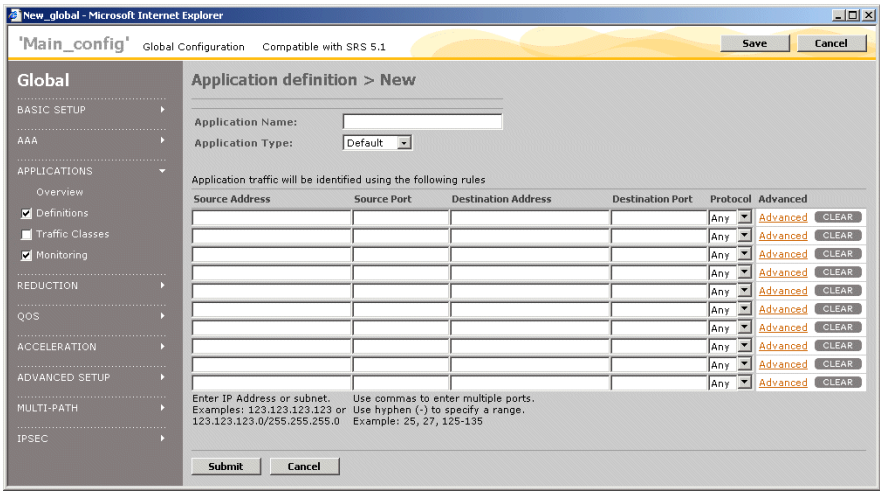


Figure 4-42 Defining New Applications

3. Specify the following information:

- Application name
- Enter a name for the application (up to 63 characters). Use only letters, numbers, blanks, and the following special characters:
: # \$ & _ - + . () '
- Application type
- Select one of the following application types:
 - **Default.** No special processing.
 - **CIFS.** Apply to CIFS application definitions whose traffic you want to accelerate (refer to “Enabling Microsoft CIFS Acceleration” on page 205). The source and destination ports for all CIFS definitions should be “139,145”.
 - **Citrix.** Apply to ICA application definitions for which you want to specify a Citrix client or application name for pattern matching.
 - **Exchange.** Apply to Exchange application definitions whose traffic you want to accelerate (refer to “Enabling Microsoft Exchange Acceleration” on page 208). Also allows Exchange ports to be learned dynamically. The source and destination ports for all Exchange definitions should be 135.

- **FTP.** Apply to the FTP application to allow FTP ports to be learned dynamically. Applies only to active FTP.
- **HTTP.** Apply to HTTP application definitions whose traffic you want to accelerate (refer to “Enabling HTTP Acceleration” on page 210). Also allows a URL to be specified for pattern matching.

Specify up to 10 rules composed of one or more of the following values. A match occurs if any of the rules are true. All values defined in the same rule must be true for a match to occur on that rule. You can specify a total of 512 rules for all applications.

Source Address	<p>Enter a source IP address or subnet. The general format is: <code>address/subnetmask</code></p> <p>A blank or an asterisk (*) with no subnet mask indicates any source IP address.</p>
Source Port	<p>Enter a source port number, a series of comma-separated port numbers, or a range of port numbers separated by a hyphen (-). A blank indicates any port. For a list of common application port numbers, refer to “Common Application Port Numbers” on page 377.</p>
Destination Address	<p>Enter a destination IP address or subnet (same format as the source address). A blank or asterisk (*) indicates any destination IP address. Typically, source and destination addresses are specified in separate rules so that a match occurs on either one. A rule that specifies both source and destination addresses will match only the traffic between those addresses.</p>
Destination Port	<p>Enter one or more destination port numbers (same format as the source port). A blank indicates any port. Typically, source and destination ports are specified in separate rules so that a match occurs on either one. A rule that specifies source and destination ports will match only the traffic between those ports.</p>
Protocol	<p>Select an application protocol or select Any to indicate TCP or UDP. You can also type in a protocol number (0 to 134). By default, a match can occur on any TCP or UDP packet.</p> <p>NOTE: Any protocol defined by number is added to the Any list of defaults that applies to each rule that does not specify a protocol. To use application pattern matching (described below), select TCP.</p>

4. To include a Type of Service (ToS) value, URL, or Citrix name in a rule, click **Advanced** next to the rule and specify the following:

ToS Bits	<p>Select the check box, and then select one of the following:</p> <ul style="list-style-type: none"> • ToS. Select an IP precedence value (0 through 7). • DSCP. Enter a DSCP value (0 through 255). <p>For more information about ToS and DSCP, refer to “Changing Outbound ToS/DSCP Values” on page 184.</p>
Application pattern matching	<p>If the application type is HTTP or Citrix, you can enter a URL or a Citrix client and/or application name.</p> <p>A URL can be up to 127 characters. The general format is:</p> <pre><host>/<uri></pre> <p>Where:</p> <p><host> is up to eight strings separated by periods. You can use an asterisk (*) by itself to indicate any string.</p> <p>For example:</p> <pre>www.juniper.*.net/</pre> <p>The slash is required even when only the host is specified. Consecutive periods, such as "...." are interpreted as ".*.*.*", and will match any host name.</p> <p><uri> is up to eight strings separated by slashes. You can use an asterisk (*) by itself to indicate any string.</p> <p>For example:</p> <pre>www.juniper.*.net/*/index.htm</pre> <p>When an asterisk is part of a string, it is treated as a single character (not a wildcard), such as “www.juniper*.net”.</p>

Click **Continue** to return to the Application Definition page.

5. Click **Submit** to enter the changes, or click **Reset** to discard them. To erase an entire rule, including the advanced settings, click **CLEAR**.

Testing New Application Definitions

Each new definition is assigned the next highest order number (the lowest precedence), and data reduction is enabled automatically. The new application is also monitored automatically if you have not exceeded the maximum number of monitored applications (40).

If you load a new definition on a device and do not see any traffic for the application, check the accuracy of the definition, and verify that the traffic is not being counted against an application with a more general definition and a higher precedence (lower order number).

Assigning Applications to Traffic Classes

Traffic classes are used by outbound QoS to allocate bandwidth to application traffic sent to the WAN, and by Policy-based Multi-Path to send traffic over the primary or secondary path to a remote device. By default, all applications belong to the Default traffic class.

You can define up to 15 additional traffic classes and assign one or more applications to each class. An application can belong to only one traffic class, but it can belong to different classes on different devices.

NOTE: Traffic classes are defined under Applications in an SRS 5.1 global configuration, and under QoS in an SRS 5.0 global configuration.

For more information about outbound QoS and Multi-Path, refer to:

- “Configuring QoS Settings” on page 154
- “Configuring Policy-Based Multi-Path” on page 235

To define traffic classes and assign applications to each class:

1. In the Configuration window, click **APPLICATIONS** in the left-hand navigation frame, click **Traffic Classes**, and select the check box.

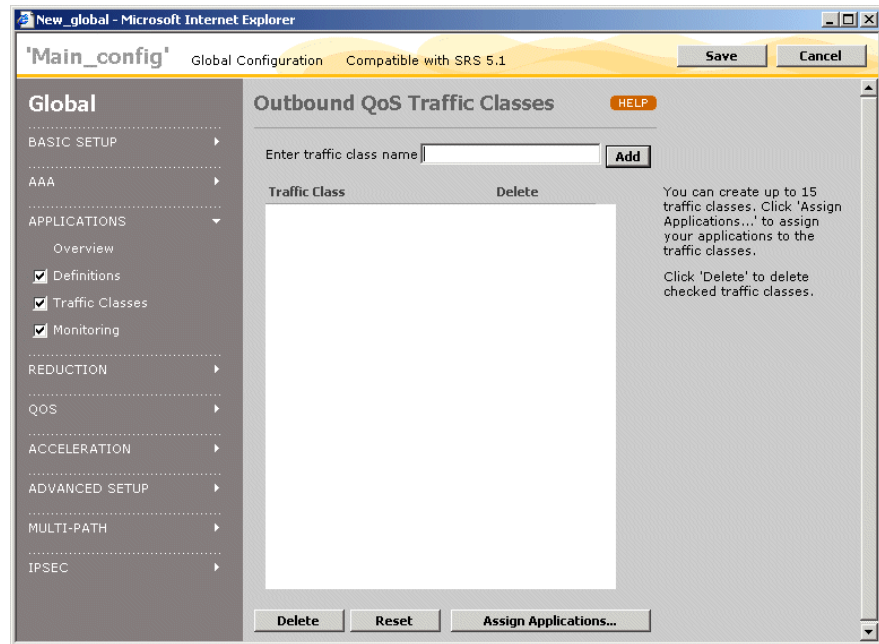


Figure 4-43 Assigning Applications to Traffic Classes

From the Traffic Classes page, you can:

- Add a new traffic class. Enter the class name (up to 20 characters), and click **Add**.

NOTE: Numeric traffic class names are not supported. Names must be alphabetic or alphanumeric.

- Change a class name. Click the class name, enter the new name, and click **Submit**.
 - Delete a traffic class. Click the check box next to the class name, and click **Delete**. Any applications in the deleted class are moved to the Default class. The Default class contains the undefined application traffic, so it cannot be renamed or deleted.
2. To change the applications assigned to each traffic class, click **Assign Applications**, select a traffic class for each application, and click **Submit**.

Monitoring Applications

Monitoring an application lets you view reduction and acceleration statistics for the application. You can monitor up to 40 applications. All unmonitored applications are placed in the “Others” category on reports.

Application definitions are provided for applications with well-known port numbers. All other applications are grouped together as “Undefined,” and are monitored automatically. To define additional applications, refer to “Configuring Application Definitions” on page 131.

NOTE: Application monitoring is defined under Applications in an SRS 5.1 global configuration, and under Reduction in an SRS 5.0 global configuration.

To select applications to be monitored:

1. In the Configuration window, click **APPLICATIONS** in the left-hand navigation frame, click **Monitoring**, and select the check box.

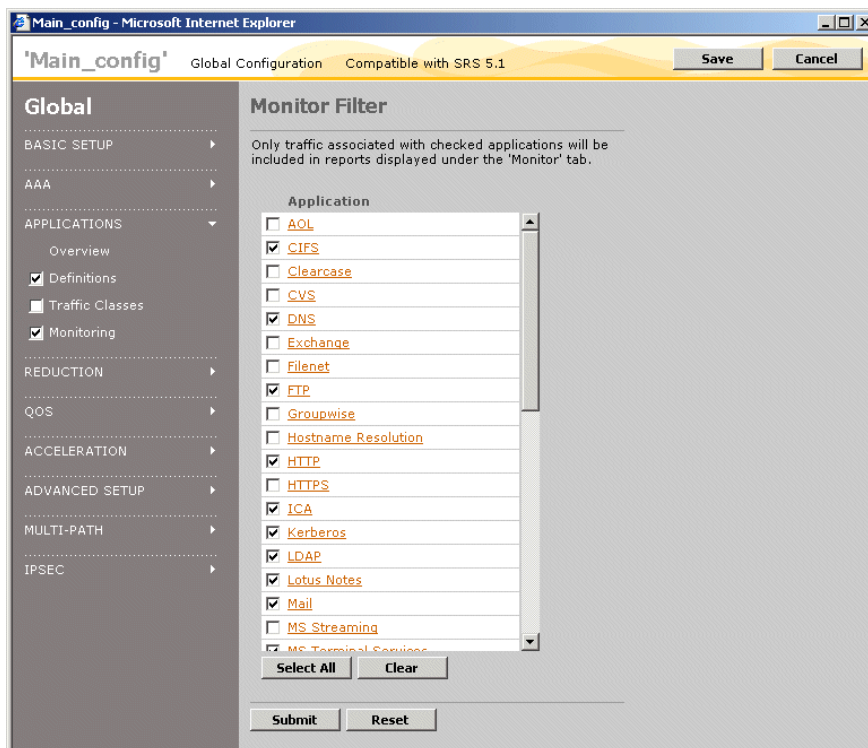


Figure 4-44 Selecting Applications for Monitoring

2. Select the check box next to each application (up to 40) for which you want to view reduction and acceleration statistics. All unreduced or unmonitored applications are placed in the “Others” category on reports.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Reduction Settings

The following sections describe the global reduction parameters:

- “Configuring Endpoints for Reduction Tunnels” on this page
- “Configuring Network Sequence Mirroring” on page 141
- “Reducing Applications” on page 144
- “Configuring Remote Routes” on page 146
- “Configuring Tunnel Load Balancing Policies” on page 147
- “Configuring Default Assemblers” on page 149
- “Defining Preferred Assemblers” on page 151
- “Configuring Tunnel Mode Settings” on page 153

Configuring Endpoints for Reduction Tunnels

By default, each device attempts to form an outbound reduction tunnel with each registered device, or “endpoint,” in the same community. Each device can have two types of tunnels—outbound tunnels that convey reduced data to remote devices, and inbound tunnels that convey the reduced data to be assembled.

Data reduction and assembly begins automatically for the reduction subnets that are advertised (refer to “Advertising Reduction Subnets” on page 94). If necessary, you can disable data reduction or assembly for all remote devices, and/or reduce data only for specific devices in each community. Each device can belong to multiple communities.

To configure the endpoints for reduction tunnels:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Reduction Endpoints**, and select the check box.

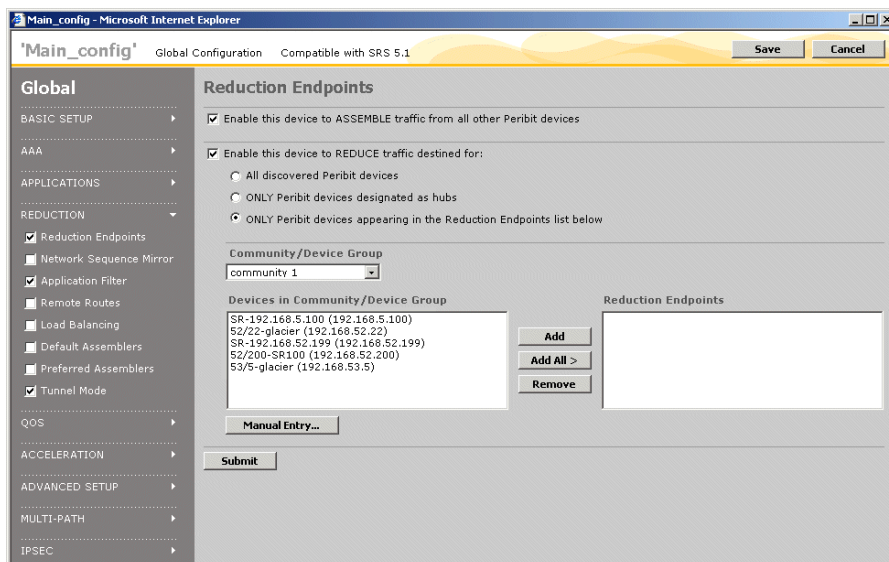


Figure 4-45 Configuring Endpoints for Reduction Tunnels

2. To stop assembling reduced data from other devices, clear the **Enable this device to ASSEMBLE traffic from all other Peribit devices** check box. All devices in the community will stop reducing data for devices that have this setting.
3. To stop reducing data for other devices, clear the **Enable this device to REDUCE traffic destined for:** check box. Otherwise, select one of the following options:
 - **All discovered Peribit devices.** Data is reduced for all other devices in the same community (default).
 - **ONLY Peribit devices designated as hubs.** Data is reduced only for devices in the same community that are designated as a hub.
 - **ONLY Peribit devices appearing in the Reduction Endpoints list below.** Data is reduced only for the devices in the Reduction Endpoints list.

4. To add devices to the **Reduction Endpoints** list:
 - a. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.
 - b. Select the devices you want to enable reduction tunnels for, and click **Add**. To remove devices from the Reduction Endpoints list, select the devices and click **Remove**.
 - c. Repeat Steps **a** and **b** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - d. If one or more devices you want to add are not listed for the community/device group, you can add the devices manually. Click **Manual Entry**, enter the device IP addresses (one per line), and click **Submit**.
5. Click **Submit** to enter the changes.

NOTE: Reduction is required for acceleration and Policy-Based Multi-Path (PBM). When you save a global configuration, an error occurs if reduction is not enabled for all endpoints that use acceleration or PBM. If you remove an endpoint from a Reduction partial configuration, an error occurs if you load the configuration on a device where acceleration or PBM are enabled.

Configuring Network Sequence Mirroring

Network Sequence Mirroring (NSM) is an enhanced data reduction technique available on Sequence Mirror (SM) devices. NSM uses disk storage to identify longer patterns of repeated traffic, and to retain those patterns for longer periods of time (even when a reduction tunnel is down). NSM is most effective where large files are often sent over the WAN, such as for database backups.

To use NSM between two Sequence Mirror devices, you must enable the following on both devices:

- Reduction tunnels (refer to “Configuring Endpoints for Reduction Tunnels” on page 139),
- Active Flow Pipelining (AFP) and outbound QoS (refer to “Enabling Acceleration by Application” on page 203)

Applications that are enabled for AFP can then be enabled for NSM (refer to “Reducing Applications” on page 144).

When you install a new Sequence Mirror, reduction tunnels, outbound QoS, AFP, and NSM are enabled automatically between the new device and all other Sequence Mirrors in the community. At any time, you can disable NSM for selected endpoints and applications.

To configure NSM for remote Sequence Mirror devices:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Network Sequence Mirror**, and select the check box.

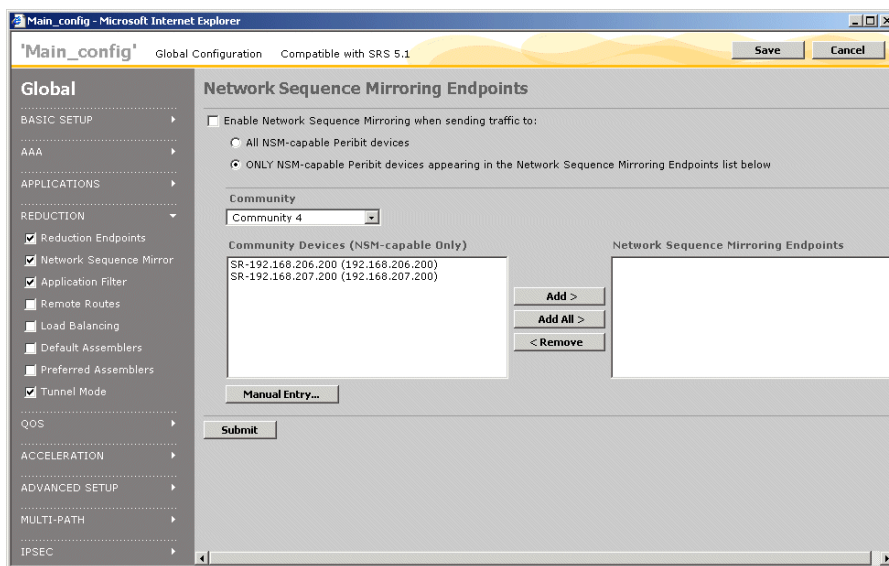


Figure 4-46 Configuring Endpoints for Network Sequence Mirroring

2. To disable NSM on this device so that standard data reduction is used for all remote devices, clear the **Enable Network Sequence Mirroring...** check box. Otherwise, select one of the following options:
 - **All NSM-capable Peribit devices.** NSM is used for all remote Sequence Mirror devices in the community (default).
 - **ONLY NSM-capable Peribit devices checked below.** NSM is used only for devices in the Network Sequence Mirroring Endpoints list.

NOTE: NSM takes effect between two SM devices only if it is enabled on both devices. NSM settings are ignored if you download the configuration to a Sequence Reducer.

3. To add devices to the **Network Sequence Mirroring Endpoints** list:
 - a. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each SM device in the selected community/device group. The IP address is enclosed in parentheses.
 - b. Select the devices you want to enable NSM for, and click **Add**. To remove devices from the list, select the devices and click **Remove**.
 - c. Repeat Steps **a** and **b** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - d. If one or more devices you want to add are not listed for the community/device group, you can add the devices manually. Click **Manual Entry**, enter the device IP addresses (one per line), and click **Submit**.
4. Click **Submit** to enter the changes.

Reducing Applications

For each application, you can enable or disable data reduction and Network Sequence Mirroring (NSM). To conserve system processing capacity, you should disable reduction for applications whose traffic is encrypted or already compressed. However, you must reduce all TCP applications that you want to accelerate.

Application definitions are provided for applications with well-known port numbers. All other applications are grouped together as “Undefined”. To define additional applications, refer to “Configuring Application Settings” on page 127.

To select the applications to be reduced and monitored:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Application Filter**, and select the check box.

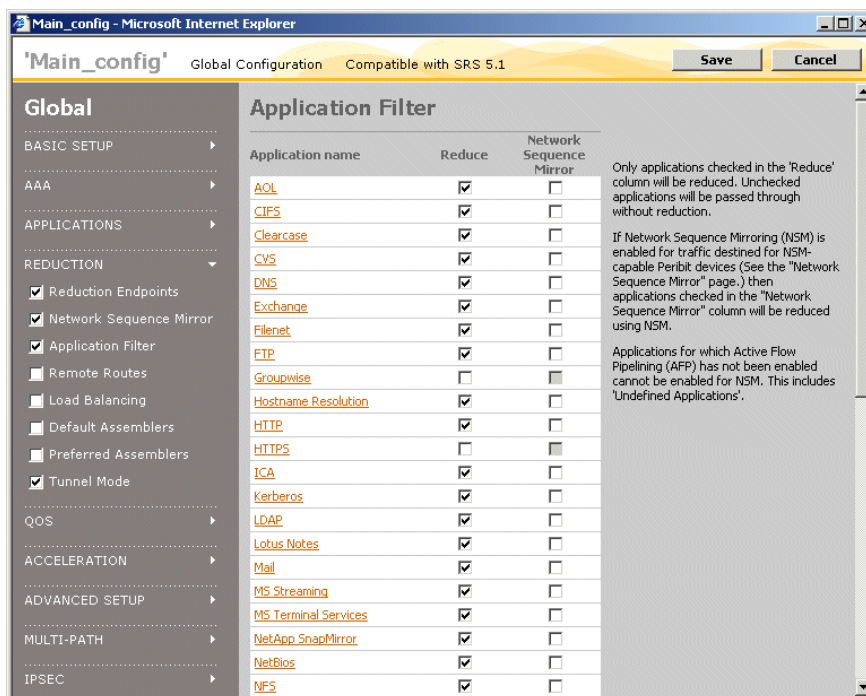


Figure 4-47 Selecting Applications for Reduction and Monitoring

2. To view or change an application's definition, click an application name, make any needed changes, and click **Submit** (global configurations only).

3. Enable or disable the following options for each application:.

Reduce	<p>Select the check box next to each application to be reduced. By default, all applications are reduced (except Groupwise, HTTPS, SNMP, SSH, and Traceroute). If an application is not reduced, its traffic passes through the device without reduction. To reduce all applications, click Reduce All.</p> <p>To conserve processing capacity, disable reduction for applications whose traffic is encrypted or already compressed. However, you must reduce all TCP applications that you want to accelerate (refer to “Configuring Traffic Acceleration” on page 190).</p>
Network Sequence Mirror	<p>On Sequence Mirror devices, you can enable Network Sequence Mirroring (NSM) for reduced applications. If NSM is enabled for one or more remote Sequence Mirror devices (refer to “Configuring Network Sequence Mirroring” on page 141), then NSM is used to reduce the application traffic sent to those devices.</p> <p>NSM uses disk storage to identify longer patterns of repeated traffic (including entire files), and is most effective for applications that do large data transfers. Standard reduction is used for traffic sent to Sequence Reducers or to Sequence Mirrors where NSM is disabled.</p> <p>To enable NSM for all reduced applications, click NSM All. To use NSM for an application, the application must be enabled for reduction and for Active Flow Pipelining (refer to “Enabling Active Flow Pipelining by Application” on page 203).</p>

4. Click **Submit** to enter the changes, or click **Reset** to discard them.

NOTE: NSM settings are ignored if you download the configuration to a Sequence Reducer.

Configuring Remote Routes

Remote routes are the reduction subnets advertised by the other devices in the community. Each device can reduce only the traffic that is destined for a remote route advertised by another device. You can specify how often remote routes are fetched from the other devices, and enable a test to validate each remote route.

NOTE: Enable the test only if the validity of the remote routes is in question. You should not use this option if load balancing is enabled (refer to “Configuring Tunnel Load Balancing Policies” on page 147).

To configure the remote route settings:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Remote Routes**, and select the check box.

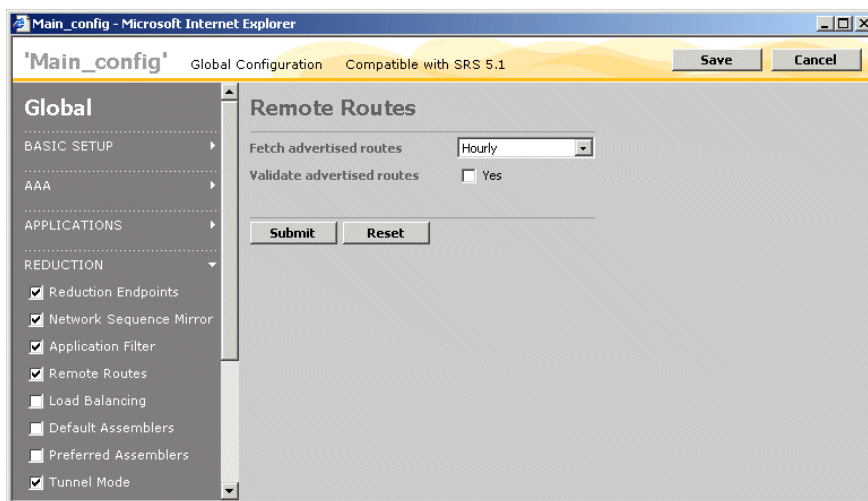


Figure 4-48 Configuring Remote Routes Parameters

2. To change how often the remote routes are fetched from the other devices in the community, select a frequency from the drop-down menu.

Remote routes are advertised each time a device starts, and route changes are advertised when they occur. Fetching routes periodically helps ensure the consistency of routing information across all the devices in the community.

3. To test the validity of each route, click **Validate advertised routes**. Each time remote routes are advertised or fetched, three probe packets are sent to three representative IP addresses in each advertised subnet. If the remote device receives any of the probes, it discards the probes without forwarding them, and returns a report to the sending device (over TCP). If a report is not received in one minute, the route is dropped from the remote routes.

NOTE: Enable this test only if the validity of the remote routes is in question. Route validation is not supported for off-path devices using packet interception or when load balancing is enabled (refer to “Configuring Tunnel Load Balancing Policies” in the next section).

4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Tunnel Load Balancing Policies

If two or more devices in the same community have equal cost paths to the same subnet, you can use load balancing to share the load of assembling the reduced data. Alternatively, you can specify preferred assemblers, as described in “Defining Preferred Assemblers” on page 151. If neither load balancing nor preferred assemblers are used, the path selection is arbitrary.

For example, in Figure 4-49, devices D2 and D3 advertise a local route to Subnet 2. On D1, the two routes to Subnet 2 have equal cost paths.

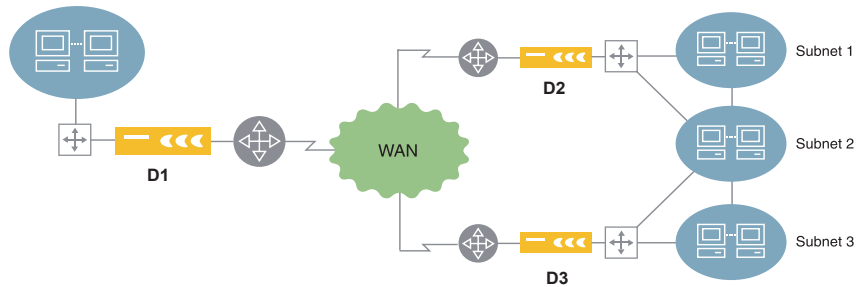


Figure 4-49 Sample Load Balancing Scenario

NOTE: If you enable load balancing policies, you should not enable the validate advertised routes feature (refer to “Configuring Remote Routes” on page 146).

To configure load balancing policies:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Load Balancing**, and select the check box.

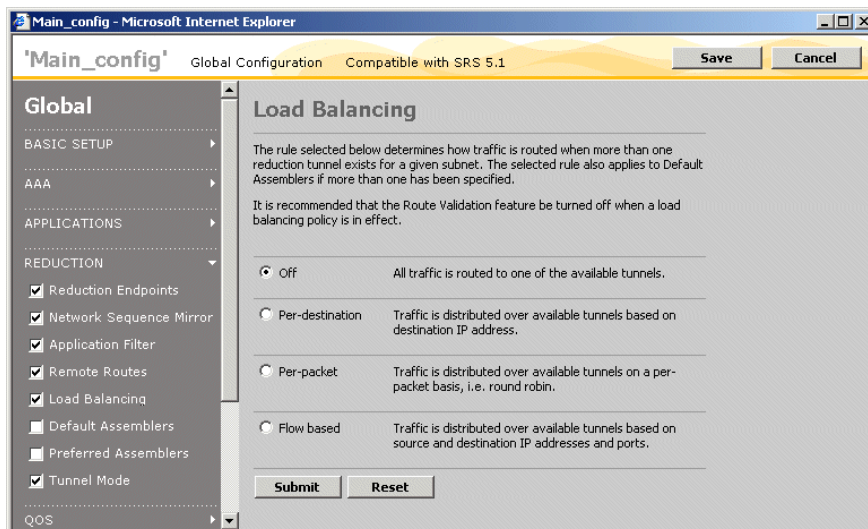


Figure 4-50 Configuring Load Balancing

2. Select one of the following load balancing policies when multiple equal cost paths exist:
 - **Off.** (Default) All traffic is routed to one of the available tunnels. No load balancing.
 - **Per-destination.** Traffic is distributed over available tunnels based on destination IP address.
 - **Per-packet.** Traffic is distributed over available tunnels on a per-packet basis (round robin).
 - **Flow based.** Traffic is distributed over available tunnels based on source and destination IP addresses and ports. If there are two or more paths in both directions, the outgoing traffic may not use the same path as the return traffic.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Default Assemblers

You can sometimes simplify route administration by designating a device as the default assembler for one or more remote devices. The default assembler need not discover and advertise all of its local routes because the remote devices automatically reduce and forward any traffic that uses the default route. In general, the default route is used when no other route is available. Note that outbound QoS and IPsec encryption also use default assemblers, regardless of whether reduction is enabled.

For example, in a Hub and Spoke topology, on each spoke device you might designate the hub as the default assembler. This ensures that all traffic goes to the hub, including the traffic destined for other spokes.

Note that traffic sent to the default assembler is not reduced when:

- The sending device has a static or dynamic route to one of the default assembler's local subnets that the default assembler has not advertised. In some cases, you may want to disable dynamic routing on the remote device.
- The sending device excludes a specific address or subnet, either through the exclusion list (see below) or through the source/destination filter defined on the device.

Figure 4-51 shows a simple example of a remote site with one outbound connection to the corporate network. If device D1 is the default assembler for D2, all traffic that uses the default route on D2 is reduced and sent to D1. To disable reduction for traffic sent to subnet S4, you can add S4 to the exclusion list on D2.

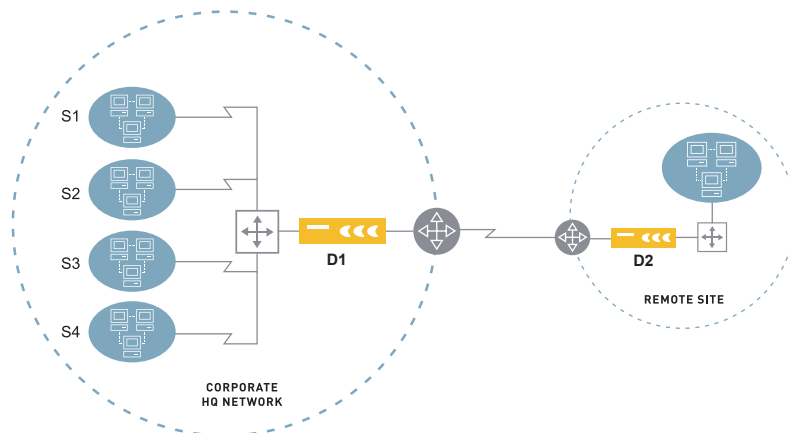


Figure 4-51 Sample Default Assembler Scenario

You can specify up to six default assemblers on a device. If you specify more than one default assembler, the current load balancing policies are applied (refer to “Configuring Tunnel Load Balancing Policies” on page 147).

To create a list of default assemblers:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Default Assemblers**, and select the check box.

The screenshot shows the 'Main_config' web interface in Microsoft Internet Explorer. The title bar says 'Main_config - Microsoft Internet Explorer'. The page has a yellow header with 'Main_config', 'Global Configuration', 'Compatible with SRS 5.1', and 'Save' and 'Cancel' buttons. The left sidebar is titled 'Global' and lists various configuration sections: BASIC SETUP, AAA, APPLICATIONS, REDUCTION (expanded), QOS, ACCELERATION, ADVANCED SETUP, MULTI-PATH, and IPSEC. Under REDUCTION, several options are checked: Reduction Endpoints, Network Sequence Mirror, Application Filter, Remote Routes, Load Balancing, Default Assemblers, Preferred Assemblers, and Tunnel Mode. The main content area is titled 'Default Assemblers'. It contains two text boxes: 'Default Assemblers' and 'Exclude List'. The 'Default Assemblers' box has a text area and a 'Submit' button. The 'Exclude List' box has a text area and a 'Reset' button. Instructions for each box are provided on the right side of the main area.

Figure 4-52 Configuring Default Assemblers

2. In the **Default Assemblers** box, enter the IP address of up to six default assemblers (one per line). If load balancing is disabled, the precedence of the default assemblers is based on their order in the list.
3. In the **Exclude List** box, enter an IP address or an IP address and subnet mask separated by a slash (/) for the hosts or subnets whose traffic is not reduced before being sent to the default assembler. If you enter an address or subnet that belongs to some other SR or SM, the exclusion is ignored.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.

5. Do the following for each default assembler (log in to the device Web console or load new Device Settings partial configurations from CMS):
 - If dynamic routing is not used, add a static route to each device in the community. The gateway for each route is the default gateway on the Remote interface (the WAN side).
 - Change the default gateway to the IP address of the next-hop router on the Local interface (the LAN side).

Defining Preferred Assemblers

If two or more devices in the same community have equal cost paths to the same subnet, you can control the selected path by specifying a preferred assembler. Alternatively, you can use load balancing to vary the selected path, as described in “Configuring Tunnel Load Balancing Policies” on page 147. If neither load balancing nor preferred assemblers are used, the path selection is arbitrary.

NOTE: Preferred assemblers are ignored if load balancing is enabled.

For example, in Figure 4-53, data from Subnet 1 has two network paths to Subnet 2. If device D1 designates D2 as a preferred assembler, all reduced data destined to Subnet 2 is sent to D2. If D2 is unavailable, D3 is used.

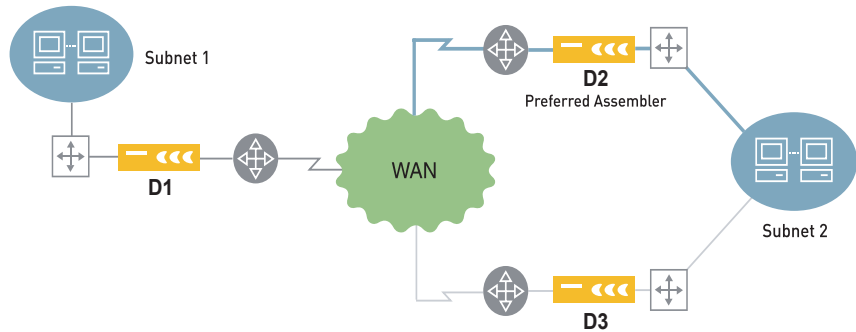


Figure 4-53 Designating a Preferred Assembler

Note that a preferred assembler is used even for routes that have a lower cost on an alternate device.

To create a list of preferred assemblers:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Preferred Assemblers**, and select the check box.

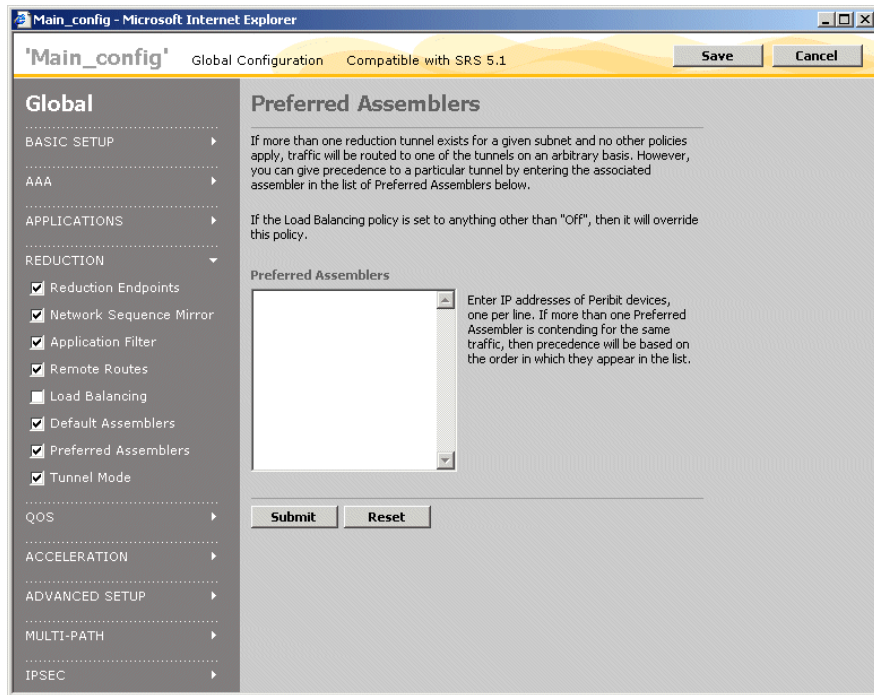


Figure 4-54 Defining Preferred Assemblers

2. Enter the IP address of a remote preferred assembler. You can specify up to 80 preferred assemblers (one per line).

If you specify more than one preferred assembler, the precedence of the preferred assemblers is based on their order in the list.

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Tunnel Mode Settings

The tunnel mode determines how reduced traffic is sent to the remote devices in the same community. By default, reduced packets are enclosed in meta packets and sent over a reduction tunnel as a single traffic flow. The tunnel modes provide varying degrees of visibility for the individual packets and traffic flows.

To configure the tunnel mode settings:

1. In the Configuration window, click **REDUCTION** in the left-hand navigation frame, click **Tunnel Mode**, and select the check box.

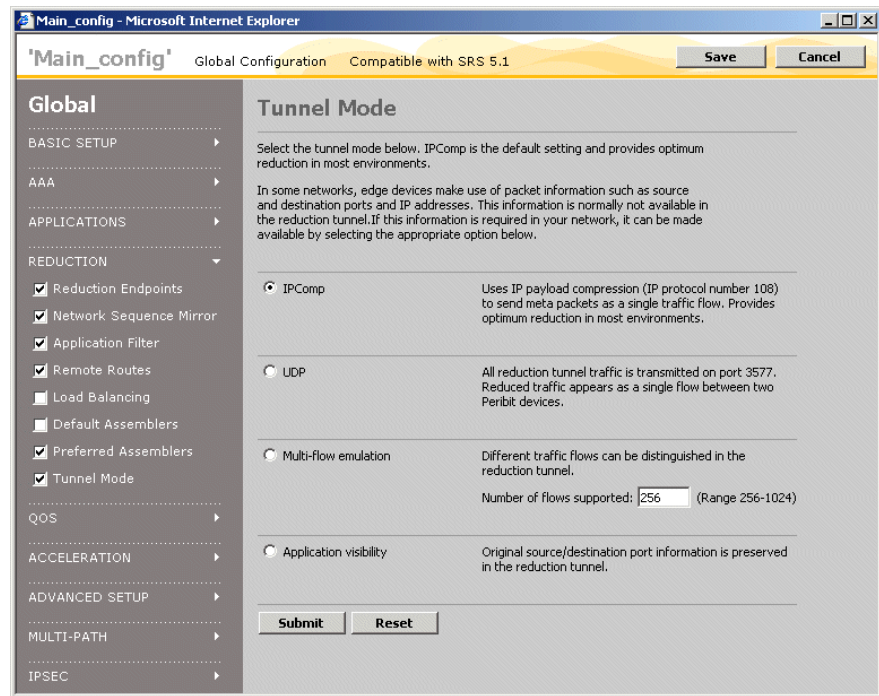


Figure 4-55 Configuring Tunnel Mode Settings

2. Select one of the following tunnel modes:

IPComp	Uses the IP payload compression protocol (protocol number 108) to send meta packets as a single traffic flow. Provides optimum reduction in most environments.
UDP	Uses UDP (port 3577) to send meta packets as a single traffic flow.

Multi-flow emulation	Uses UDP and arbitrarily assigns source port numbers to each traffic flow so that routers using Weighted Fair Queueing (WFQ) can distribute WAN bandwidth among the various flows. Enter the maximum number of flows expected (256 through 1024) to help allocate resources efficiently (not a hard limit).
Application visibility	Uses UDP and preserves the source and destination ports of all packets so that performance monitoring tools can identify the devices responsible for the traffic in the reduction tunnel. Your tools must be configured to monitor UDP traffic.

NOTE: The multi-flow emulation and application visibility options reduce packet aggregation, thus affecting the reduction in the number of packets.

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring QoS Settings

The following sections describe the global outbound QoS parameters:

- “Using Outbound QoS to Enhance Performance” on page 155
- “Understanding Outbound QoS” on page 156
- “Procedure for Configuring Outbound QoS Policies” on page 166
- “Using the Outbound QoS Setup Wizard” on page 168
- “Defining Outbound QoS Settings by Endpoint” on page 176
- “Defining Outbound QoS Templates” on page 179
- “Defining Outbound QoS Endpoints” on page 180
- “Changing Outbound ToS/DSCP Values” on page 184
- “Starting and Stopping Outbound QoS” on page 187
- “Configuring Inbound QoS Policies” on page 188

Using Outbound QoS to Enhance Performance

Outbound QoS provides two key benefits:

- **Basic bandwidth allocation.** Data reduction performance is automatically optimized based on the local WAN speed, and is particularly effective for low-speed links. Only minimal QoS settings are required.
- **Advanced bandwidth allocation.** Application performance across the WAN is optimized by specifying guaranteed bandwidths for critical applications.

NOTE: Basic bandwidth allocation is highly recommended to optimize performance on all devices.

The advanced QoS policies let you guarantee bandwidths by traffic class, and define templates of QoS policies that can be easily applied to multiple endpoints. ToS and DSCP markings can be used for QoS scheduling and/or preserved for use by upstream devices. Special bandwidth policies can be configured to handle “oversubscribed” WANs where the local WAN bandwidth is less than the sum of the remote endpoint bandwidths.

To enable basic bandwidth allocation:

1. Specify the outbound WAN speed, as described in “Defining Outbound QoS Endpoints” on page 180. Adding the remote devices and specifying the WAN circuit speed for each device is also recommended.

For guidance on adjusting the WAN speeds to account for router overhead, refer to “WAN Circuit Speeds and Router Overhead” on page 159.

2. Start outbound QoS using Weighted Fair Queuing (WFQ) or Weighted Strict Priority (WSP), as described in “Starting and Stopping Outbound QoS” on page 187. Unless you need strict priority treatment for traffic classes, WFQ is recommended.

Understanding Outbound QoS

Outbound QoS policies control how WAN bandwidth is allocated to the various types of application traffic that traverse the device. These policies apply to both reduced and unreduced traffic. Outbound bandwidth management lets you:

- Guarantee a minimum bandwidth for your most critical applications.
- Set priorities to determine how the “excess” bandwidth is allocated. The excess bandwidth is the unguaranteed bandwidth, plus the guaranteed bandwidth that is not currently in use.
- Set maximum bandwidths to limit (or drop) low-priority traffic.
- Change the ToS/DSCP values on selected traffic for use by other QoS devices in the network.

A Setup Wizard is provided to simplify the creation of QoS templates that specify the priorities and bandwidths by traffic class. Templates created by the wizard can be modified manually.

NOTE: Outbound bandwidth management is not effective for an off-path device unless all outbound WAN traffic is routed through the device.

The following topics provide an overview of outbound QoS:

- “Traffic Classes and Bandwidths” on page 157
- “QoS Templates and Endpoints” on page 158
- “WAN Circuit Speeds and Router Overhead” on page 159
- “Dedicated, Oversubscribed, and Variable Rate WANs” on page 160
- “Direct Setup Versus Wizard Configuration Results” on page 162
- “Class Priorities and Excess Bandwidth Allocation” on page 164
- “ToS/DSCP Values” on page 166
- “Unadvertised Subnets” on page 166

Traffic Classes and Bandwidths

Priorities and bandwidths are specified by traffic class, and each class can have one or more applications. Initially, all applications belong to the Default class. To guarantee a minimum bandwidth for one application, assign the application to its own class, and then specify the guaranteed bandwidth. Figure 4-56 shows the default settings for the standard traffic classes created by the Setup Wizard. You can have up to 16 traffic classes.

Traffic Class	Priority	Guaranteed Bandwidth	Maximum Bandwidth
Default	0 (Lowest)	0.00 %	100.00 %
Business Critical	0 (Lowest)	40.00 %	100.00 %
Business Standard	0 (Lowest)	20.00 %	100.00 %
Low-Latency	7 (Highest)	20.00 %	100.00 %
Prohibited	0 (Lowest)	0.00 %	0.00 %

Figure 4-56 Predefined Traffic Classes

You can guarantee up to 80% of the total bandwidth across all classes. Traffic is dropped when the maximum bandwidth is exceeded or when the guaranteed bandwidth is exceeded while the circuit is fully utilized, such as during a burst of high-priority traffic. The 20% of unguaranteed bandwidth ensures that bandwidth is always available for local system resources, such as SNMP updates and management traffic.

The priority value (0 to 7) assigned to each traffic class is used to allocate the excess bandwidth to each class as the traffic load fluctuates (refer to “Class Priorities and Excess Bandwidth Allocation” on page 164).

Note that the Default class, which cannot be deleted, includes all undefined traffic. You must create an application definition for any traffic whose bandwidth you want to manage separately (refer to “Configuring Application Definitions” on page 131).

QoS Templates and Endpoints

The priorities and bandwidths defined for each traffic class constitute a template. On each device, you can manage the outbound bandwidth by assigning a template to each remote device (endpoint). You can create a different template for each endpoint, or create a single template and customize it for specific endpoints.

NOTE: QoS templates let you vary the priorities and bandwidths for each traffic class, but all templates (and all endpoints) have the same traffic classes, and the same applications in each class.

The Setup Wizard creates two identical templates and assigns them to the selected endpoints:

- **Wizard-PrimeTime.** Applies to prime time hours, or to all hours if prime time is not defined. To specify the prime time, refer to “Defining the Prime Time” on page 218.
- **Wizard-NonPrimeTime.** Applies to non-prime time hours (if prime time hours are defined), and can be modified to allocate more bandwidth to applications that run during off-peak hours, such as database backups. To view bandwidth reports for prime time vs. non-prime time hours, use the device Web console.

You can also assign a template to the predefined “Other Traffic” endpoint to manage outbound traffic that does not have a remote device or for which the remote device is not enabled for outbound QoS. In addition, to more closely manage “Other Traffic”, you can create virtual endpoints for specific remote subnets.

WAN Circuit Speeds and Router Overhead

On each device that supports outbound QoS, you must specify the following WAN circuit speeds:

- **Outbound speed.** The sum of the WAN circuit speeds on the adjacent router (the aggregate local WAN speed). Alternatively, you can use the device's Ethernet speed if you enable congestion control for all remote endpoints.
- **Endpoint circuit speeds.** The maximum WAN circuit speed associated with each remote WX device for which you want to manage the outbound bandwidth. You can use the device's Ethernet speed if you enable congestion control for the remote endpoint.

NOTE: To effectively manage the WAN bandwidth, the WX device must be the sole source of the WAN traffic.

If congestion control is NOT enabled, all WAN circuit speeds specified for outbound QoS must be set slightly lower than the WAN router's full interface speed to allow for router overhead (Frame Relay LMI updates, CDP, SNMP, routing updates, and so on). Setting the bandwidth about 2% below the link speed should work well in most cases. However, the router overhead is highly variable, and depends on the network configuration.

The following table provides some recommended adjustments to the WAN interface speeds. Note that failure to account for router overhead will effectively shift bandwidth management to the router, and may cause the router to drop traffic.

Table 4-5 Recommended WAN Circuit Speed Adjustments

WAN Interface	Recommended QoS Speed	Description
Frame Relay	CIR minus 2%	Reduce the Committed Information Rate (CIR) by 2%. Higher speeds, up to the Peak Information Rate (PIR), may be acceptable, depending on the traffic load and whether "discard eligible" traffic is actually discarded. If the device exceeds the CIR, and discard eligible traffic is dropped, the QoS behavior may be unpredictable.
1.544 Mbps (T1)	1500 Kbps	The T1 line rate is 1.544 Mbps, but the data rate is 1.536 Mbps. The 8 Kbps difference is used for framing and encapsulation. Subtracting 2% from 1.536 yields about 1.5 Mbps.
512 Kbps (Fractional T1)	500 Kbps	Use one third of the T1 setting.
64 Kbps	60 Kbps	On low-speed links, router overhead may take up a greater percentage of the WAN link speed. Using 60 Kbps assumes that 6% of the link is used for router control traffic.

Dedicated, Oversubscribed, and Variable Rate WANs

In point-to-multi-point configurations, the guaranteed bandwidth percentages assigned to each traffic class can be adjusted automatically, depending on whether the WAN is “dedicated” or “oversubscribed,” and whether the available bandwidth is variable:

- **Dedicated.** The sum of the WAN circuit speeds on the adjacent router (the outbound speed) is equal to or greater than the sum of the remote WAN speeds. In this case, no adjustments to the bandwidth percentages are needed. In Figure 4-57, the aggregate WAN speed for device D1 (the outbound speed) is 1.5 Mbps, which equals the total speed of the three remote endpoints—D2, D3, and Other Traffic.

If D1 specifies a guaranteed bandwidth of 60% for all traffic classes for each endpoint, the guaranteed capacity is 300 Kbps for D2, D3, and Other Traffic ($.6 \times 500$ Kbps).

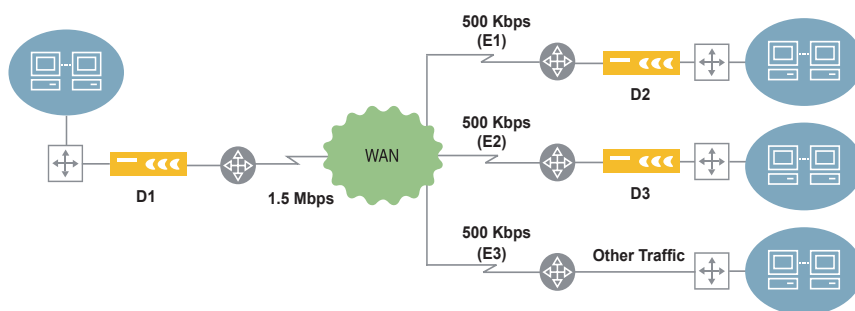


Figure 4-57 Dedicated WAN

- **Oversubscribed.** The local outbound WAN speed is less than the sum of the remote WAN speeds. In this case, the total guaranteed bandwidth across all classes *and endpoints*, cannot exceed 80% of the outbound speed. In Figure 4-58, the WAN is oversubscribed from the perspective of device D1 because the outbound speed is 1.5 Mbps and the sum of the remote speeds is 2060 Mbps.

On D1, if you manually specify a guaranteed bandwidth of 60% for all traffic classes for each endpoint, an error occurs because the sum of the guaranteed bandwidths for all endpoints ($300 + 900 + 36 = 1236$ Kbps) exceeds 80% of the outbound speed (1200 Kbps). However, the Setup Wizard lets you enter guarantees of up to 80%, and then automatically adjusts the guaranteed bandwidths for each traffic class to proportionately distribute the total guaranteed bandwidth.

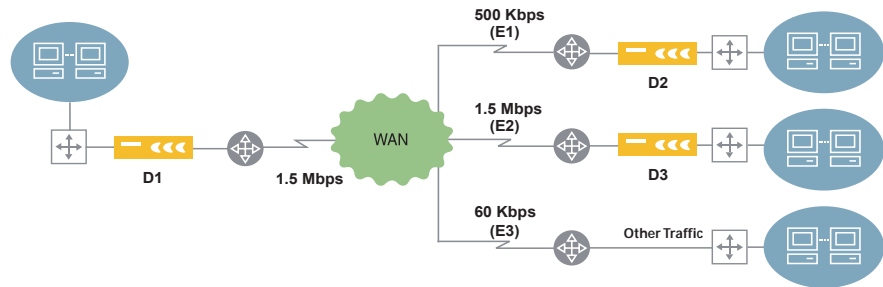


Figure 4-58 Oversubscribed WAN

- Variable WAN Bandwidth Support.** Some networks have variable WAN bandwidths, such as Frame Relay networks, which support a sustained CIR rate and bursts to a peak rate, MPLS networks, which are inherently "connectionless," and shared satellite uplink environments where several routers may share a single satellite connection. The congestion control feature dynamically alters the bandwidth allocation per-endpoint based on the measured real-time available WAN bandwidth.

Since congestion control dynamically adjusts to the available bandwidth, the specified WAN speeds are not critical. For example, you can specify the Ethernet speed as the outbound speed, and then enable congestion control for all remote endpoints. However, in the case of oversubscribed WANs, the displayed guaranteed bandwidths may not be accurate (refer to “Defining Outbound QoS Settings by Endpoint” on page 176).

NOTE: Congestion control manages only traffic sent to other WX endpoints. If you have substantial passthrough traffic for other destinations, you may want to reduce the maximum speed for the “Other traffic” endpoint to limit the bandwidth allocated to passthrough traffic (refer to “Defining Outbound QoS Endpoints” on page 180).

Direct Setup Versus Wizard Configuration Results

For a dedicated WAN, if you apply the same bandwidths and priorities to each endpoint, the Setup Wizard produces the same results as entering the QoS settings directly. However, for an oversubscribed WAN, the Wizard adjusts the template percentages so that the guaranteed portion of the outbound speed is distributed fairly across all classes and endpoints.

For example, Table 4-6 shows the Wizard and direct setup results when D1 in Figure 4-58 is configured with two traffic classes and the same guaranteed bandwidths for each endpoint.

Table 4-6 Direct Setup Versus Wizard Results for a Simple Oversubscribed WAN for Device D1

Endpoint	Remote Circuit Speed	Traffic Class	Class Guaranteed Percentage	Direct Guaranteed Percentage	Direct Guaranteed Rate	Wizard Guaranteed Percentage	Wizard Guaranteed Rate
E1	500 Kbps	Default	15%	15%	75 Kbps	10.92%	54 Kbps
		Business	40%	40%	200 Kbps	29.12%	145 Kbps
E2	1500 Kbps	Default	15%	15%	225 Kbps	10.92%	163 Kbps
		Business	40%	40%	600 Kbps	29.12%	436 Kbps
E3	60 Kbps	Default	15%	15%	9 Kbps	10.92%	6 Kbps
		Business	40%	40%	24 Kbps	29.12%	17 Kbps
Totals	2060 Kbps		55%	55%	1133 Kbps	40.04%	821 Kbps

Direct Setup Results

If you enter the QoS settings directly, the **Direct Guaranteed Rate** column in Table 4-6 shows the guaranteed bandwidth in Kbps allocated to each traffic class on each endpoint. The guaranteed rate is calculated as follows:

(Remote Circuit Speed) * (Class Guaranteed Percentage)

For example, the guaranteed rate for the Default class at endpoint D1 is:

$(500) * (.15) = 75 \text{ Kbps}$

Since the total guaranteed bandwidth (1133 Kbps) does not exceed 80% of the D1 outbound speed ($.8 * 1500 = 1200 \text{ Kbps}$), you can enter all the QoS settings directly without having to adjust the guaranteed percentages. Figure 4-59 shows the “Oversubscribed” template specifying the 15% and 40% guarantees, and Figure 4-60 shows the guaranteed bandwidths in Kbps displayed on the Outbound QoS Overview page when the template is applied to each endpoint.

Template Name		Oversubscribed	
Traffic Class	Priority	Bandwidth Limit (%)	
		Guaranteed	Maximum
Default	0 (Lowest)	15.00	100.00
Business	0 (Lowest)	40.00	100.00

Figure 4-59 Oversubscribed Template for Device D1

Endpoint	Template	Circuit Speed (Kbps)	Traffic Classes		Total Guaranteed Bandwidth	
			Default	Business		
Other traffic	EDIT	Oversubscribed	60	9	24	33
192.168.53.5	EDIT	Oversubscribed	500	75	200	275
192.168.52.22	EDIT	Oversubscribed	1500	225	600	825
Total			309	824		1133

Figure 4-60 Direct Setup Results on the Outbound QoS Overview Page for Device D1

Wizard Results

If you use the Setup Wizard, the 15% and 40% guarantees entered in the Wizard are adjusted in the resulting Wizard template, as shown in the **Wizard Guaranteed Percentage** column in Table 4-6. The Wizard template guarantees are calculated as follows:

(Class Guaranteed Percentage) * (Outbound Speed/Total Remote Circuit Speeds)

For example, the 15% guarantee entered for the Default class becomes:

$$(.15) * (1500/2060) = .1092 = 10.92\%$$

The **Wizard Guaranteed Rate** column shows the adjusted guaranteed rates for each class on each endpoint. For example, the guaranteed rate for the Default class at endpoint D1 is:

$$(500) * (.1092) = 54 \text{ Kbps}$$

Note that the Wizard total guaranteed bandwidth (821 Kbps) is 55% (15% + 40%) of the outbound speed (1500 Kbps) for SR1. Figure 4-61 shows the guaranteed bandwidths in Kbps generated by the Setup Wizard and displayed on the Outbound QoS Overview page.

Endpoint		Template	Circuit Speed (Kbps)	Traffic Classes		Total Guaranteed Bandwidth
				Default	Business	
Other traffic	EDIT	Wizard-PrimeTime	60	6	17	23
192.168.53.5	EDIT	Wizard-PrimeTime	500	54	145	199
192.168.52.22	EDIT	Wizard-PrimeTime	1500	163	436	599
Total				223	598	821

Figure 4-61 Wizard Results on the Outbound QoS Overview Page for Device D1

The Wizard adjusts the bandwidths for oversubscribed WANs only when there are multiple remote endpoints. For example, in Figure 4-58 on page 161, the WAN is oversubscribed from the perspective of D2, but the bandwidths defined on D2 would not be adjusted because D1 is the only remote endpoint.

Class Priorities and Excess Bandwidth Allocation

Excess bandwidth is the unguaranteed bandwidth, plus the guaranteed bandwidth that is not currently in use. As the traffic load varies, the excess bandwidth is allocated dynamically to each traffic class based on the class priority (0 to 7) and the selected queuing model. The two queuing models are Weighted Fair Queuing and Weighted Strict Priority (the selected model applies to all classes).

NOTE: The priorities assigned to each traffic class are used only by the WX device, and are not related to ToS priorities.

- **Weighted Strict Priority (WSP).** Queues are created for each priority, and the excess bandwidth is allocated by processing the queues based only on priority. That is, the class with the highest priority gets all the excess bandwidth it needs before any excess bandwidth is allocated to the class with the next highest priority.
- **Weighted Fair Queuing (WFQ).** Queues are created for each traffic class, and the excess bandwidth is allocated as described in Table 4-7. The allocation depends on whether the WAN is dedicated or oversubscribed.

Table 4-7 WFQ Allocation of Excess Bandwidth

WAN Type	Excess Bandwidth Allocation
Dedicated	<p data-bbox="654 314 1278 425">To calculate the percentage of excess bandwidth allocated to a traffic class for a specific remote endpoint (since priorities start with zero, they must be incremented by one for this calculation):</p> $\text{(Class Priority + 1)} / \text{(Sum of active class priorities + 1 for each class)}$ <p data-bbox="654 513 1278 624">For example, for the five standard classes where four classes have priority zero and the Low Latency class has priority 7, the Low Latency class receives the following minimum percentage of excess bandwidth:</p> $\text{Excess\%} = 8/12 = 66\%$ <p data-bbox="654 682 1278 734">Note that if only one class has traffic, then that class receives 100% of the bandwidth.</p> <p data-bbox="654 751 1278 803">To calculate the minimum excess bandwidth for a class in Kbps:</p> $\text{(Excess\%)}(\text{Remote WAN speed} - \text{Total class guarantee in Kbps})$ <p data-bbox="654 892 1278 972">For example, if the Excess% is 66%, the remote WAN speed is 500 Kbps, and the guaranteed bandwidth for all classes is 80%, the minimum excess bandwidth is:</p> $(.66)(500 - 500 \times .8) = 66 \text{ Kbps}$
Oversubscribed	<p data-bbox="654 1038 1278 1149">The excess bandwidth percentage for a class on a specific endpoint is calculated in the same manner as a dedicated WAN, except that the priorities must be totaled across all remote endpoints.</p> <p data-bbox="654 1166 1278 1246">For example, if you have three endpoints using the same classes and priorities as in the dedicated example, the minimum excess bandwidth for the Low Latency class is:</p> $\text{Excess\%} = 8/(12 + 12 + 12) = 22\%$ <p data-bbox="654 1303 1278 1355">To calculate the minimum excess bandwidth for a class in Kbps:</p> $\text{(Excess\%)}(\text{Outbound speed} - \text{All endpoint class guarantees in Kbps})$ <p data-bbox="654 1444 1278 1588">Note that you must calculate the sum of the guaranteed bandwidths for each class on each remote endpoint. For the example in Table 4-6 on page 162, the sum of the bandwidths is 1133 Kbps using direct setup or 821 Kbps using the Wizard.</p>

ToS/DSCP Values

The ToS/DSCP values in the packet headers can be set by traffic class for use by other devices in your network. You can also preserve the incoming ToS/DSCP values in the “meta-packets,” so that each meta-packet encapsulates only packets that have the same ToS/DSCP value. This allows other QoS devices in the path to manage the meta-packets in the same manner as the individual packets. For more information about setting ToS/DSCP values, refer to “Changing Outbound ToS/DSCP Values” on page 184.

Unadvertised Subnets

All subnets that are not advertised by a WX device will be managed by the QoS settings for the “Other traffic” endpoint. To ensure that the appropriate QoS policies are applied to all traffic, each device should advertise all the subnets it can access. The source/destination filter can be used to prevent data reduction for specific destinations, as needed (refer to “Configuring Source/Destination Filters” on page 216).

By default, each device dynamically adjusts its advertised subnets to exclude any hosts or gateways that become unreachable. Traffic to these “carved out” addresses is also attributed to the “Other traffic” endpoint.

Procedure for Configuring Outbound QoS Policies

Use the following procedure to configure outbound QoS policies on each device:

1. For best results, verify that each device advertises all the subnets it can access. Unadvertised subnets are managed by the QoS settings for the “Other traffic” endpoint. If necessary, use the source/destination filter to prevent data reduction for specific destinations (refer to “Configuring Source/Destination Filters” on page 216).
2. Run the Setup Wizard or specify the outbound QoS policies directly:
 - To run the Setup Wizard in CMS, refer to “Using the Outbound QoS Setup Wizard” in the next section). The Setup Wizard creates and applies the **Wizard-PrimeTime** and **Wizard-NonPrimeTime** templates to the selected endpoints.

CAUTION: Each time you run the Setup Wizard the two existing Wizard templates are overwritten and all customized settings are lost, including the customized settings for each endpoint. To preserve custom settings, use the Setup Wizard for the initial configuration, and then make all subsequent changes directly.

- To specify the outbound QoS policies directly:
 - a. Specify the traffic classes and the applications in each class (refer to “Assigning Applications to Traffic Classes” on page 136).
 - b. Define one or more templates to specify the priorities and bandwidths for each traffic class (refer to “Defining Outbound QoS Templates” on page 179).
 - c. Specify the local outbound speed and the maximum circuit speeds for each remote endpoint (refer to “WAN Circuit Speeds and Router Overhead” on page 159 and “Defining Outbound QoS Endpoints” on page 180).
 - d. Assign a template to each endpoint (refer to “Defining Outbound QoS Settings by Endpoint” on page 176).
 - e. Enable QoS and select a queuing model (refer to “Starting and Stopping Outbound QoS” on page 187).
- 3. Note that the following changes must be made directly:
 - Change a template for a specific endpoint (refer to “Defining Outbound QoS Settings by Endpoint” on page 176).
 - Change traffic class names (refer to “Assigning Applications to Traffic Classes” on page 136).
 - Add new templates, change a template name, or change just one of the Wizard templates (refer to “Defining Outbound QoS Templates” on page 179).
 - Define virtual endpoints or exclude address or subnet pairs from bandwidth management (refer to “Defining Outbound QoS Exclusions” on page 96).
 - Change the ToS/DSCP values for one or more traffic classes (refer to “Changing Outbound ToS/DSCP Values” on page 184).

Using the Outbound QoS Setup Wizard

Use the Setup Wizard the first time you define outbound QoS policies. The Setup Wizard creates two identical templates and assigns them to the selected endpoints:

- **Wizard-PrimeTime.** Applies to the prime time hours (critical business hours). To specify the prime time, refer to “Defining the Prime Time” on page 218.
- **Wizard-NonPrimeTime.** Applies to nonprime time hours. To view QoS reports for prime time or nonprime time hours, use the device Web console.

Each time you run the Setup Wizard, both of the templates and all customized settings are overwritten. To change just one of the templates, refer to “Defining Outbound QoS Templates” on page 179.

To run the outbound QoS Setup Wizard:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, and then click **Setup Wizard**.
2. Click **Enable Outbound QoS** and click **Next**.

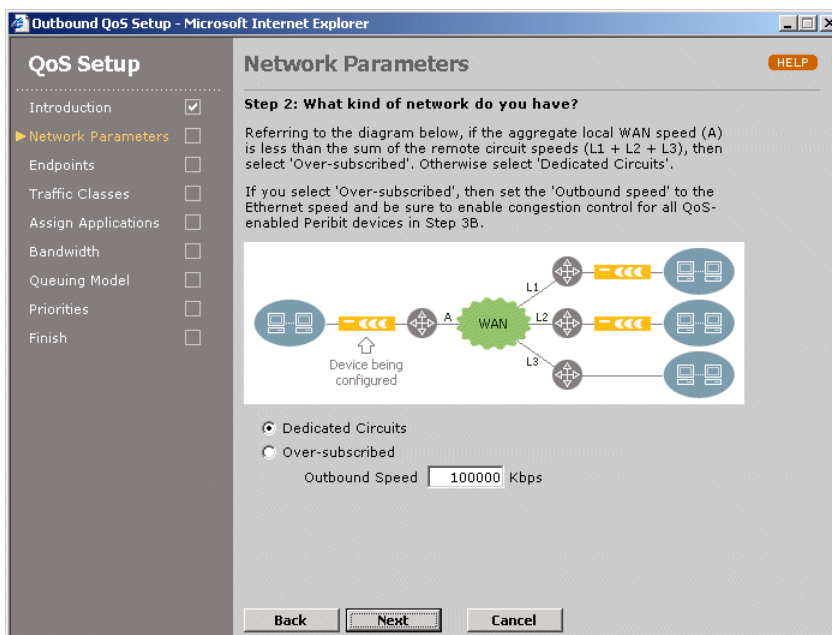


Figure 4-62 Configuring Outbound QoS Network Parameters

3. Calculate the local outbound WAN speed by adding up the speeds of all the WAN interfaces on the router adjacent to the device(s) where you intend to load the configuration, and then select one of the following and click **Next**:

Dedicated Circuits Indicates that the local outbound WAN speed equals or exceeds the sum of the remote WAN speeds whose bandwidths you want to manage (the default).

Over-subscribed Indicates that the local outbound WAN speed is less than the sum of the remote WAN speeds. If you plan to enable congestion control for all endpoints, you can enter the device's Ethernet speed in the **Out-bound Speed** field; otherwise, enter the local WAN speed and be sure to account for router overhead (refer to "WAN Circuit Speeds and Router Overhead" on page 159).

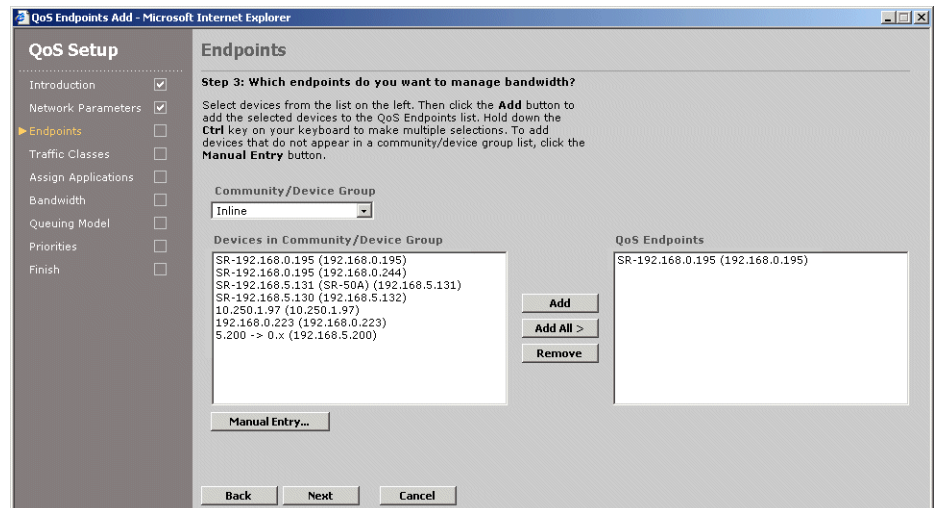


Figure 4-63 Configuring QoS Endpoints

4. To enable outbound QoS to one or more remote endpoints:
 - a. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable QoS for one or both paths. To configure Multi-Path, refer to "Configuring Multi-Path Addresses" on page 101.

- b. Select the devices you want to enable outbound QoS for, and click **Add**. To remove devices from the QoS Endpoints list, select the devices and click **Remove**.
- c. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
- d. If one or more devices are not listed, click **Manual Entry** and enter the device IP addresses manually (one per line), and click **Submit**.

NOTE: Outbound QoS is required for Packet Flow Acceleration (PFA). When you save a global configuration, an error occurs if QoS is not enabled for all endpoints using PFA. If you remove an endpoint from a QoS partial configuration, an error occurs if you load the configuration on a device where PFA is enabled for that endpoint.

- e. When you are done, click **Next**.

Outbound QoS Setup - Microsoft Internet Explorer

QoS Setup

- Introduction ☒
- Network Parameters ☒
- Endpoints** ☒
- Traffic Classes ☐
- Assign Applications ☐
- Bandwidth ☐
- Queuing Model ☐
- Priorities ☐
- Finish ☐

Endpoints HELP

Step 3: For which endpoints do you want to manage bandwidth?

Enter the circuit speeds (in kbps) for all endpoints listed here. Then select the endpoints that you want to participate in Bandwidth Management.

Endpoint	Name	Circuit Speed
<input type="checkbox"/> Other traffic		1000000
<input checked="" type="checkbox"/> No remote Peribit	Branch1	256
<input type="checkbox"/> 192.168.54.22	54/22	
<input type="checkbox"/> 192.168.55.100	55/100-SR20	
<input type="checkbox"/> 192.168.5.101	SR-192.168.5.101	
<input type="checkbox"/> 192.168.52.149	52/149	

Endpoints enabled for acceleration cannot be disabled for QoS.

Figure 4-64 Configuring Endpoint Circuit Speeds

5. Enter the remote WAN circuit speed (in Kbps) for each endpoint.

CAUTION: If you do not enable congestion control (see Step 6), be sure to adjust the WAN speed to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 159). Exceeding the actual WAN speed effectively shifts bandwidth management to the router, and may cause the router to drop traffic.

Note the following:

- The “Other traffic” endpoint is used to manage the bandwidth for all traffic that is not sent to one of the selected devices. For oversubscribed WANs, the “Other traffic” endpoint is shown here, and the two generated templates are applied to it. The circuit speed for “Other traffic” defaults to the outbound speed.
- If any “No Remote Peribit” virtual endpoints have been defined to manage the “Other traffic” sent to specific remote subnets (refer to “Defining Outbound QoS Endpoints” on page 180), you can change their circuit speeds or disable them. The settings for “Other traffic” and virtual endpoints can be changed in the same manner as other endpoints (refer to “Defining Outbound QoS Settings by Endpoint” on page 176).

Click **Next**.

QoS Setup

- Introduction ☒
- Network Parameters ☒
- Endpoints ☐
- Traffic Classes ☐
- Assign Applications ☐
- Bandwidth ☐
- Queuing Model ☐
- Priorities ☐
- Finish ☐

Endpoints HELP

Step 3B: Do your circuit speeds vary?

For some endpoints, actual circuit speeds may vary. In order to optimize Bandwidth Management for these endpoints, you should enable Congestion Control and indicate the minimum speed (in kbps) for the relevant endpoints. If you don't know the minimum speed, enter '0'.

☒ Enable Congestion Control when sending traffic to:

☐ All QoS-enabled Peribit devices

☒ ONLY checked Peribit devices below

IP Address	Device Name	Minimum Speed
<input type="checkbox"/> 192.168.54.22	54/22	<input type="text"/>
<input type="checkbox"/> 192.168.55.100	55/100-SR20	<input type="text"/>
<input type="checkbox"/> 192.168.5.101	SR-192.168.5.101	<input type="text"/>
<input type="checkbox"/> 192.168.52.149	52/149	<input type="text"/>

Figure 4-65 Configuring Congestion Control

6. If the WAN bandwidth to a remote WX device is variable, such as for MPLS, Frame Relay, or shared satellite links, enable congestion control for traffic sent to that device. Also, if you entered the device's Ethernet speed as the outbound speed, enable congestion control for all endpoints.

Congestion control dynamically adjusts the bandwidth allocation for each endpoint based on the latency measured for the ACKs returned for each reduced meta packet. Throughput is lowered as latency increases, and increased as latency decreases. In this way, congestion control can set the speed to slightly below the level where packet loss starts to occur.

To enable congestion control:

- a. Select **Enable Congestion Control...** and select one of the following options:
 - **All QoS-enabled Peribit devices.** Applies congestion control to all remote devices for which QoS is enabled (default).
 - **ONLY checked Peribit devices below.** Select the check box for one or more QoS-enabled endpoints.
- b. Enter a minimum circuit speed for each endpoint. For Frame Relay, use the CIR; for a shared satellite link, use a percentage of the total speed, depending on how many devices share the link. For MPLS networks, use the service level guarantee. If you do not know the minimum speed, enter a zero.

NOTE: Congestion control manages only traffic sent to other WX endpoints for which reduction tunnels are enabled. If you have substantial passthrough traffic for other destinations, you may want to reduce the maximum speed for the "Other traffic" endpoint to limit the bandwidth allocated to passthrough traffic. After you complete the Wizard configuration, refer to "Defining Outbound QoS Endpoints" on page 180.

7. Click **Next**. To define your own traffic classes, click **Custom**, and then click **Next**.

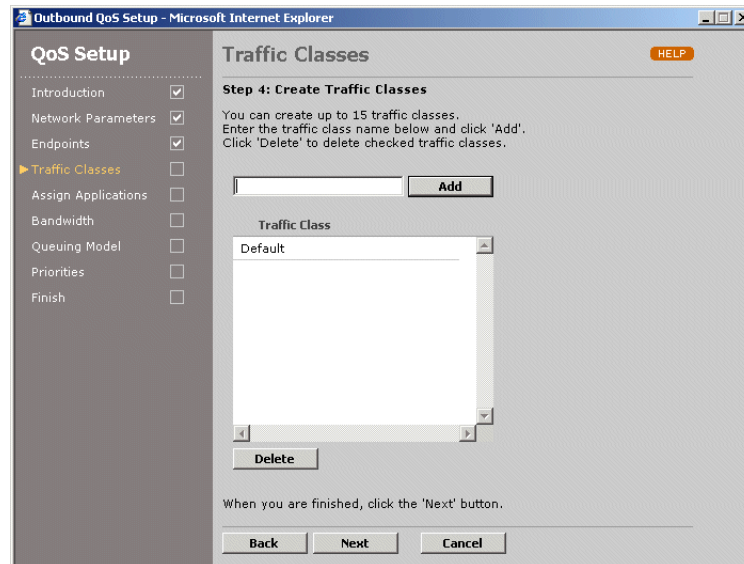


Figure 4-66 Configuring Traffic Classes

8. To add a new traffic class, enter the class name (up to 20 characters) and click **Add**. You can add up to 15 classes. To delete a traffic class, click the check box next to the class name and click **Delete**. The Default class is reserved for undefined application traffic and cannot be deleted. Click **Next**.

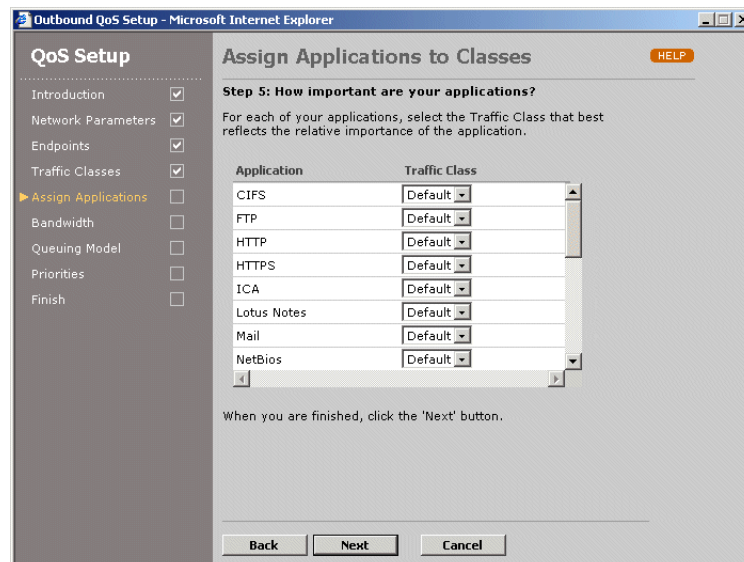


Figure 4-67 Assigning Applications to Traffic Classes

9. Select the appropriate traffic class for each application. If one of your network applications is not shown, you must create an application definition for it, as described in “Configuring Application Settings” on page 127. Click **Next**.

Figure 4-68 Defining Guaranteed and Maximum Bandwidths

10. Enter the bandwidth information for each traffic class, and click **Next**. Note that these bandwidths have no effect until you complete the configuration by running the Setup Wizard in the device Web console.

Guaranteed
Bandwidth

Percentage of the bandwidth that is guaranteed to be allocated to the applications in the traffic class. Lower values indicate that the traffic in the class is more likely to be delayed. Traffic may be dropped when the guaranteed bandwidth is exceeded, such as during a burst of higher-priority traffic.

The total guaranteed bandwidth across all traffic classes cannot exceed 80%. Also, the total guaranteed bandwidth across all endpoints cannot exceed 80% of the local outbound WAN speed.

Maximum
Bandwidth

Maximum percentage of the bandwidth that can be allocated to the applications in the traffic class. Traffic is dropped when the maximum bandwidth is exceeded. A zero indicates that all traffic in the class is dropped.

NOTE: If more than one application is assigned to a class, the specified bandwidths are distributed evenly among the applications.

11. Select one of the following queuing models to allocate the available bandwidth as load conditions change. The available bandwidth is the unguaranteed bandwidth, plus the guaranteed bandwidth that is not currently in use.

Weighted Fair Queuing Queues are created for each traffic class, and the available bandwidth is allocated by processing the queues based on their priority and guaranteed bandwidth.

Weighted Strict Priority Queues are created for each priority, and the available bandwidth is allocated by processing the queues based on their priority. Processing is weighted equally for traffic classes that have the same priority.

You can later change the queuing method, as described in “Starting and Stopping Outbound QoS” on page 187. Click **Next**.

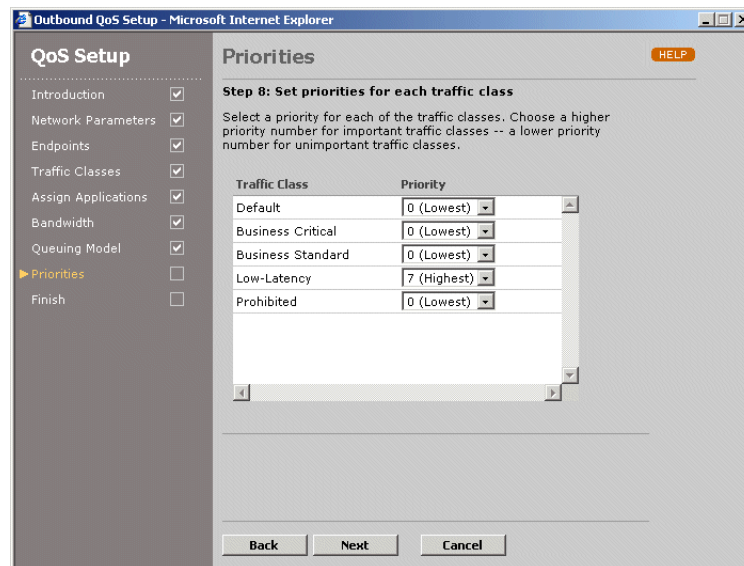


Figure 4-69 Defining Priorities by Traffic Class

12. Select a priority value (0 to 7) for each traffic class, where 7 is the highest priority. These values are used by the Weighted Fair Queuing and Weighted Strict Priority queuing models to allocate available (unguaranteed) bandwidth to the competing traffic classes. These priorities are used only by the WX device, and are not related to ToS priorities.

13. Click **Next**, click **Submit**, and then click **Close**.

You can now customize the outbound QoS settings for each endpoint, as described in “Defining Outbound QoS Settings by Endpoint” in the next section.

Defining Outbound QoS Settings by Endpoint

After you run the Setup Wizard to create the initial outbound QoS settings, you can manually change the prime-time or nonprime-time template assigned to each endpoint or override the template values (class priorities or bandwidths) for a single endpoint. To change the WAN circuit speed for an endpoint, refer to “Defining Outbound QoS Endpoints” on page 180.

To view or change the outbound QoS settings by endpoint:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, and then click **Overview**.

The screenshot shows the 'Main_config' web interface in Microsoft Internet Explorer. The left navigation pane is expanded to 'QOS', and 'Overview' is selected. The main content area is titled 'Outbound QoS Overview'. It includes configuration options for Display (Guaranteed Bandwidth), Show bandwidth as (% of circuit speed), Time Frame (Prime Time), Outbound Speed (500 Kbps), and Outbound QoS Setting (Off). Below these options is a table of endpoints with their templates, circuit speeds, and guaranteed bandwidths across different traffic classes.

Endpoint	Template	Circuit Speed (Kbps)	Default	Traffic Classes					Total Guaranteed Bandwidth
				Business Critical	Business Standard	Low-Latency	Prohibited		
Other traffic	Wizard-PrimeTime	10000	0.00	13.33	6.66	6.66	0.00	26.65	
192.168.5.131	Wizard-PrimeTime	10000	0.00	13.33	6.66	6.66	0.00	26.65	
192.168.5.200	Wizard-PrimeTime	10000	0.00	13.33	6.66	6.66	0.00	26.65	

Figure 4-70 Outbound QoS Overview

The Outbound QoS Overview page shows the outbound speed of the device, the selected queuing model, and the template name, circuit speed, and guaranteed bandwidths for each remote endpoint.

Note that the “Other traffic” endpoint lets you manage the bandwidth for all traffic that is not sent to one of the other endpoints shown here.

2. To change the data shown for each endpoint, select one or more of the following and click **Update**.
 - Select **Maximum Bandwidth** from the **Display** menu to view the maximum bandwidth values for each endpoint.
 - Select **Kbps** from the **Show bandwidth as** menu to view the bandwidth percentages as circuit speeds.

NOTE: If congestion control is enabled, the guaranteed bandwidths shown in Kbps will not be accurate. For oversubscribed WANs, the guaranteed percentages will be accurate only if the outbound speed and remote speeds are the true WAN speeds (not the Ethernet speeds).

- Select **Non Prime Time** from the **Time Frame** menu to view the nonprime-time templates associated with each endpoint. This menu is displayed only if prime time is enabled (refer to “Defining the Prime Time” on page 218).
3. To change an endpoint’s template or override a template setting, click **EDIT** next to the endpoint name. To override a template, be sure to select the appropriate time frame from the **Time Frame** menu (**Prime Time** or **Non Prime Time**).

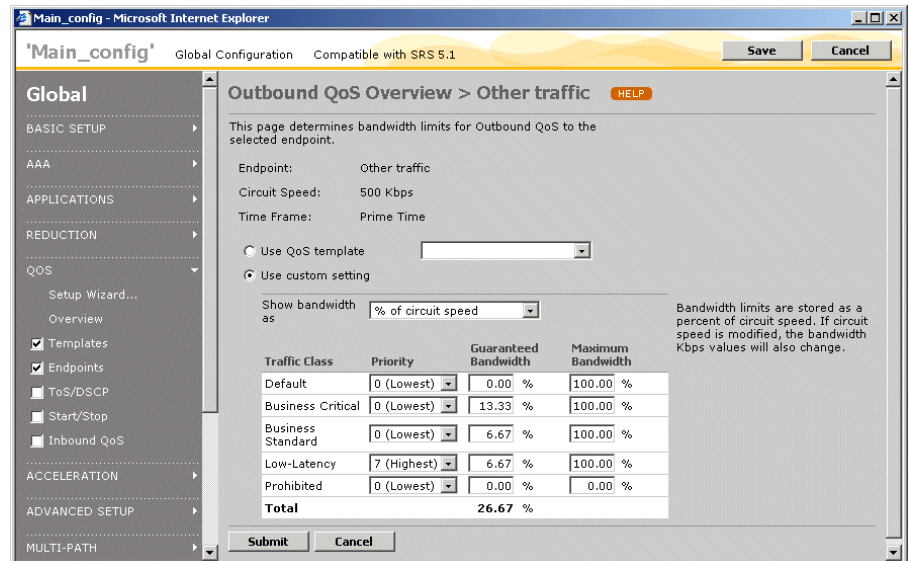


Figure 4-71 Changing Endpoint Templates or Template Settings

4. Do one of the following:

- To change the template for this endpoint, select a template from the drop-down menu, and click **Submit**. To create new templates, refer to “Defining Outbound QoS Templates” on page 179.
- To override the current template settings for this endpoint, click **Use custom setting** and change the priority or bandwidth settings for one or more traffic classes, and click **Submit**.

Note that to increase the guaranteed bandwidth for a traffic class on an oversubscribed WAN, you must first decrease the bandwidth on another class (on the same endpoint or a different endpoint), reduce the circuit speed, or increase the outbound speed. The Setup Wizard adjusts the guaranteed bandwidths for you (refer to “Using the Outbound QoS Setup Wizard” on page 168).

5. Click **Submit** to enter the changes, or click **Reset** to discard them.

The Outbound QoS Overview page is displayed. When you override the template settings for an endpoint, the template name is changed to **None**. You can later reapply the template to restore the original settings.

Defining Outbound QoS Templates

Templates specify the priority, and guaranteed and maximum bandwidths for each traffic class. You can change the templates created by the Setup Wizard or create new templates. To apply a template to an endpoint, refer to “Defining Outbound QoS Settings by Endpoint” on page 176.

To define outbound QoS templates:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, click **Templates**, and select the check box.

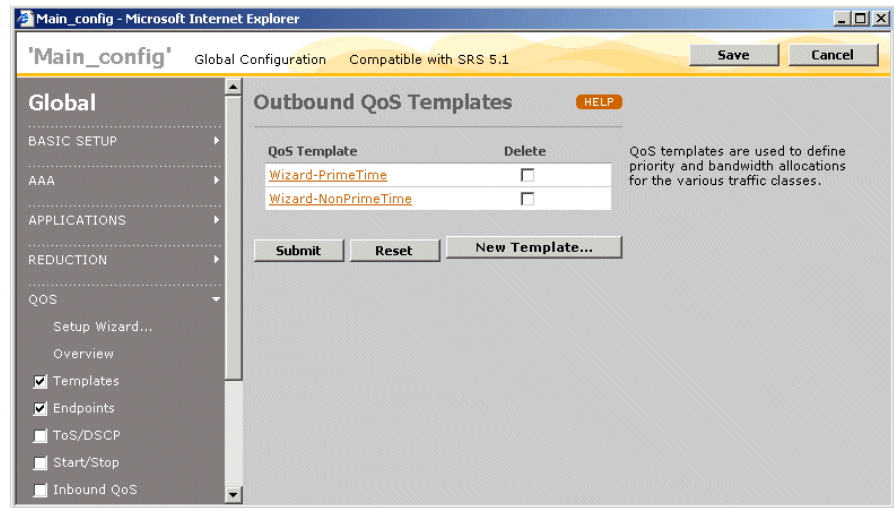


Figure 4-72 Defining Outbound QoS Templates

From the Outbound QoS Templates page, you can:

- Add a new template, as described in Step 2 through Step 3.
- Change a template name or settings. Click the template name, change the template name and/or the settings for each traffic class, and click **Submit**.
- Delete a template. Click the check box next to the template name, and click **Submit**. If the template is applied to an endpoint, all priority and guaranteed bandwidth values are set to zero for that endpoint. Maximum bandwidth values are set to 100%.

2. To add a new template, click **New Template** and enter the following information:

Template Name	Enter the name of the template (up to 20 characters).
Priority	Select a priority value (0 to 7), where 7 is the highest priority. These values are used by the Weighted Fair Queuing and Strict Priority queuing models to allocate excess bandwidth to the competing classes of applications.
Guaranteed Bandwidth	<p>Enter a percentage of the bandwidth that is guaranteed to be allocated to the applications in the traffic class. Lower values indicate that the traffic in the class is more likely to be delayed. Traffic may be dropped when the guaranteed bandwidth is exceeded, such as during a burst of higher-priority traffic.</p> <p>The total guaranteed bandwidth across all traffic classes cannot exceed 80%. Also, the total guaranteed bandwidth across all endpoints cannot exceed 80% of the outbound speed.</p>
Maximum Bandwidth	Enter the maximum percentage of the bandwidth that can be allocated to the applications in the traffic class. Traffic is dropped when the maximum bandwidth is exceeded. A zero indicates that all traffic in the class is dropped.

NOTE: If more than one application is assigned to a class, the bandwidths defined for the class are distributed evenly among the applications.

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Defining Outbound QoS Endpoints

Each device can manage the outbound bandwidth for one or more remote WX devices or other (virtual) endpoints. After you run the Setup Wizard, you can:

- Add or remove endpoints for bandwidth management.
- Define virtual endpoints to manage the traffic to specific remote subnets that do not have a WX device.
- Change the local outbound WAN speed or remote WAN circuit speeds.
- Enable congestion control for one or more endpoints.

To exclude specific LAN/WAN address or subnet pairs from bandwidth management, refer to “Defining Outbound QoS Exclusions” on page 96.

For oversubscribed WANs, you may have to decrease some speeds or guaranteed percentages before increasing others. If you use the Setup Wizard to change QoS settings, all percentages are adjusted automatically (refer to “Using the Outbound QoS Setup Wizard” on page 168).

To define the outbound QoS endpoints:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, click **Endpoints**, and select the check box.

Main_config Global Configuration Compatible with SRS 5.1 Save Cancel

Global

- BASIC SETUP
- AAA
- APPLICATIONS
- REDUCTION
- QOS**
 - Setup Wizard...
 - Overview
 - ☒ Templates
 - ☒ Endpoints
 - ☐ ToS/DSCP
 - ☐ Start/Stop
 - ☐ Inbound QoS
- ACCELERATION
- ADVANCED SETUP
- MULTI-PATH
- IPSEC

Outbound QoS Endpoints HELP

You can enable Outbound QoS by using the Setup Wizard or from the 'Start/Stop' page. If Outbound QoS is enabled, then only outbound traffic destined for the checked endpoints below will be affected.

Perbit devices are automatically included in the list below. If you want to enable QoS to endpoints that are NOT reachable through a Perbit device, you can manually add the endpoint to the list by clicking **ADD**. To [view a list of remote networks NOT accessed through a remote Perbit device](#), click this link.

For some endpoints, actual circuit speeds may vary. In order to optimize Bandwidth Management for these endpoints, you should enable Congestion Control and indicate the minimum speed (in Kbps) for the relevant endpoints. If you don't know the minimum speed, enter '0'. NOTE: Endpoints enabled for Congestion Control must also be enabled for Reduction. (See the **Endpoints** page under **REDUCTION**.)

Aggregate WAN Speed Kbps

☒ Enable Congestion Control when sending traffic to:

- ☐ All QoS-enabled Perbit devices
- ☒ ONLY Perbit devices checked under "Congestion Control"

Endpoint	Name	Circuit Speed (Kbps)	Congestion Control	Min. Speed (Kbps)
Other traffic		default	<input type="button" value="ADD..."/>	
<input checked="" type="checkbox"/> No Remote Perbit	Non-SR-Chicago	<input type="text" value="500"/>	<input type="button" value="DELETE"/>	
<input type="checkbox"/> 192.168.207.200	192.168.207.200	<input type="text" value="0"/>	<input checked="" type="checkbox"/>	<input type="text" value="512"/>
<input type="checkbox"/> 192.168.206.200	SR-192.168.206.200	<input type="text" value="0"/>	<input type="checkbox"/>	<input type="text" value=""/>

Figure 4-73 Enabling Bandwidth Management by Endpoint

2. To change the local outbound speed associated with the device(s) where you intend to load the configuration, or the maximum circuit speed associated with each remote endpoint, enter the new values (in Kbps) and click **Submit**. For a description of the outbound WAN speed, refer to “Dedicated, Oversubscribed, and Variable Rate WANs” on page 160. The outbound and remote circuit speeds are required.

Note that the “Other traffic” endpoint is used to manage the bandwidth for traffic that is not sent to one of the other endpoints.

CAUTION: If congestion control is not enabled (see Step 4), be sure to adjust the WAN speed to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 159).

3. To add or remove remote endpoints for outbound QoS:

- a. Click **Add/Remove Endpoints**.
- b. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable QoS for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 101.

- c. Select the devices you want to enable outbound QoS for, and click **Add**. To remove devices from the QoS Endpoints list, select the devices and click **Remove**.
- d. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
- e. If one or more devices are not listed, click **Manual Entry** and enter the device IP addresses manually (one per line), and click **Submit**.
- f. When you are done, click **Submit**.

NOTE: Outbound QoS is required for Packet Flow Acceleration (PFA). When you save a global configuration, an error occurs if QoS is not enabled for all endpoints using PFA. If you remove an endpoint from a QoS partial configuration, an error occurs if you load the configuration on a device where PFA is enabled for that endpoint.

- g. Enter the remote WAN circuit speed (in Kbps) for each endpoint that you added, and click **Submit**.

When you add a new endpoint, all the endpoint’s traffic classes have a priority and guaranteed bandwidth of zero, and a maximum bandwidth of 100%. To change the default settings, refer to “Defining Outbound QoS Settings by Endpoint” on page 176.

4. If the WAN bandwidth to a remote WX device is variable, such as for MPLS, Frame Relay, or shared satellite links, enable congestion control for traffic sent to that device. Also, if you entered the device's Ethernet speed as the outbound speed, enable congestion control for all endpoints.

Congestion control dynamically adjusts the bandwidth allocation for each endpoint based on the latency measured for the ACKs returned for each reduced meta packet. Throughput is lowered as latency increases, and increased as latency decreases. In this way, congestion control can set the speed to slightly below the level where packet loss starts to occur.

To enable congestion control:

- a. Select **Enable Congestion Control...** and select one of the following options:
 - **All QoS-enabled Peribit devices.** Applies congestion control to all remote devices for which QoS is enabled (default).
 - **ONLY Peribit devices checked under “Congestion Control”.** Select the **Congestion Control** check box for one or more QoS-enabled endpoints.
- b. Enter a minimum circuit speed for each endpoint. For Frame Relay, use the CIR; for a shared satellite link, use a percentage of the total speed, depending on how many devices share the link. For MPLS networks, use the service level guarantee. If you do not know the minimum speed, enter zero.

NOTE: Congestion control manages only traffic sent to other WX endpoints (reduction tunnels are required). If you have substantial passthrough traffic for other destinations, you may want to reduce the maximum speed for the “Other traffic” and virtual endpoints to limit the bandwidth allocated to passthrough traffic.

5. Virtual endpoints let you manage the traffic to specific remote subnets that do not have a WX device. By default, all such traffic is managed by the “Other traffic” endpoint. To view the subnets associated with the current virtual endpoints, click **view a list of remote networks...**

To add a virtual endpoint, click **ADD**, specify the following information, and click **Submit**. The maximum number of virtual endpoints depends on the device type (2 for the SR-15, 5 for the SR-20 and SM-250, 60 for the SM-500, 120 for the SR-5x, and 320 for the SR-80/SR-100).

Name	Enter the endpoint name (up to 20 characters).
Circuit Speed	Enter the WAN circuit speed associated with this endpoint (in Kbps).
Subnets	<p>Enter the IP addresses or subnets associated with this endpoint (one per line). The subnet format is:</p> <p><IP address>/<subnet mask></p> <p>Subnets specified here are ignored if they are also advertised by a WX device.</p>

To change a virtual endpoint’s name or subnets, click the endpoint name, make the changes, and click **Submit**. To delete a virtual endpoint, click **DELETE** next to the endpoint. Traffic to deleted virtual endpoints is managed by the “Other-traffic” endpoint.

Changing Outbound ToS/DSCP Values

The ToS/DSCP values on incoming traffic from the LAN can be modified to support other QoS devices in your network. For each traffic class, you can specify a Type of Service (ToS) IP precedence value or a Differentiated Services Code Point (DSCP) value, depending on the QoS scheme in use. The specified ToS/DSCP values apply to all traffic in the class, regardless of whether the traffic is reduced or outbound QoS is enabled.

You can also preserve the incoming ToS/DSCP values in the “meta-packets,” so that each meta-packet encapsulates only packets that have the same ToS/DSCP value. This allows other QoS devices in the path to manage the meta-packets in the same manner as individual packets. By default, meta-packets have a ToS/DSCP value of zero and can encapsulate packets with varying ToS/DSCP values.

ToS IP precedence values (0 to 7) use the upper three bits of the Diffserv field; DSCP values (0 to 63) use the upper six bits. The upper three bits of DSCP are used like ToS to indicate the priority (7 is the highest priority). Table 4-8 lists the

equivalent DSCP and ToS IP precedence values for the class selector (CSx) names often used to describe each setting, and the DSCP values for the per-hop behaviors (PHBs) defined by RFCs 2597 and 2598.

Table 4-8 ToS and DSCP Values

Name	DSCP	IP Precedence
Default or BE (best effort)	0	0
CS1	8	1
CS2	16	2
CS3	24	3
CS4	32	4
CS5	40	5
CS6	48	6
CS7	56	7
AF11	10	—
AF12	12	—
AF13	14	—
AF21	18	—
AF22	20	—
AF23	22	—
AF31	26	—
AF32	28	—
AF33	30	—
AF41	34	—
AF42	36	—
AF43	38	—
EF	46	—

To set ToS/DSCP values by traffic class:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, click **ToS/DSCP**, and select the check box.

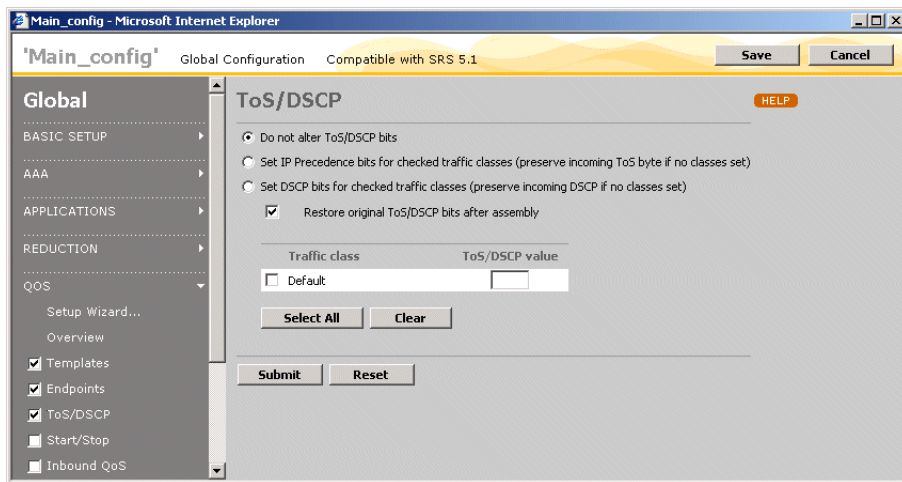


Figure 4-74 Setting ToS/DSCP Values

2. To set ToS/DSCP values by traffic class, select **Set IP Precedence bits...** or **Set DSCP bits...** to specify whether you want to enter ToS or DSCP values.

The default selection, **Do not alter ToS/DSCP bits**, indicates that meta-packets have a ToS/DSCP value of zero. If you want to preserve all the incoming values, and have each meta-packet reflect the ToS/DSCP value of its encapsulated packets, select **Set IP Precedence bits...** or **Set DSCP bits...** and do not check any of the traffic classes.

3. Select the check boxes next to the traffic classes whose ToS/DSCP values you want to set (or click **Select All**).
4. Enter a ToS value (0 to 7) or a DSCP value (0 to 63) in the **ToS/DSCP value** field for each of the selected classes. The value specified for each class is applied to the traffic for all applications in the selected class. To assign applications to a traffic class, refer to “Assigning Applications to Traffic Classes” on page 136.

NOTE: These ToS/DSCP values are overridden by the ToS/DSCP settings defined for Multi-Path (refer to “Enabling Policy-Based Multi-Path” on page 237), and by ToS values set for router-based balancing (refer to the “configure route” CLI command).

5. After reduced traffic from remote devices is assembled, the **Restore original ToS/DSCP bits after assembly** option resets the ToS/DSCP value to its original value (if the remote device changed it).
6. Click **Submit** to enter the changes, or click **Reset** to discard them.

Starting and Stopping Outbound QoS

You can start or stop outbound QoS or change the queuing method without using the Setup Wizard. The queuing method determines how the available (unguaranteed) bandwidth is allocated among the contending applications. The selected queuing model applies to all the managed endpoints.

To stop the outbound QoS service or change the queuing model:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, click **Start/Stop**, and select the check box.

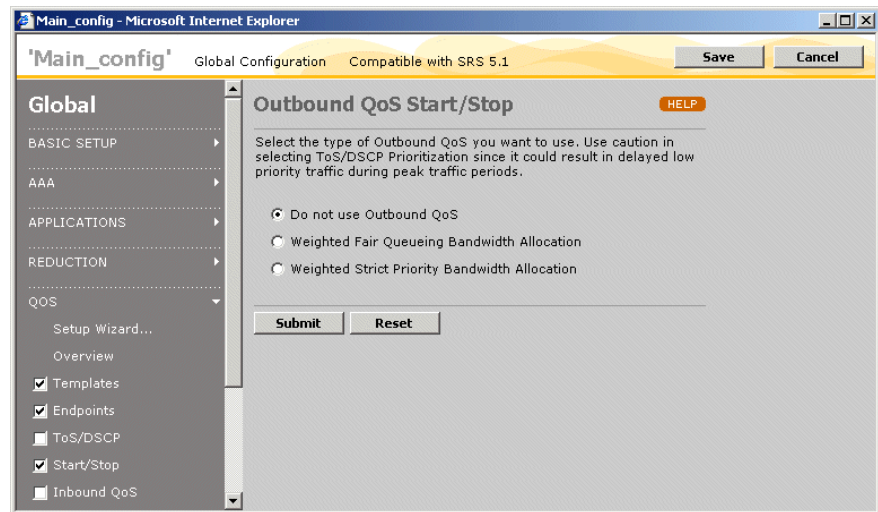


Figure 4-75 Starting and Stopping Outbound QoS

2. To stop the outbound QoS service, click **Do not use Outbound QoS**.
3. To restart the service or change the queuing method used for each endpoint, select one of the following:
 - **Weighted Fair Queueing Bandwidth Allocation.** Queues are created for each traffic class, and the available bandwidth is allocated by processing the queues based on their priority and guaranteed bandwidth.

- **Weighted Strict Priority Bandwidth Allocation.** Queues are created for each priority, and the available bandwidth is allocated by processing the queues based only on priority.

NOTE: When you save a global configuration, an error occurs if PFA is enabled and outbound QoS is off.

4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Inbound QoS Policies

For an SRS 5.0 configuration, you can specify the maximum bandwidths for four classes of incoming WAN traffic destined for the Local Area Network (LAN). Setting maximum bandwidths for each class ensures that low-priority traffic, such as Web traffic, does not interfere with mission-critical applications. Bandwidths are specified as percentages of the inbound WAN speed, and traffic that exceeds the maximum bandwidths is dropped.

NOTE: Inbound QoS applies only to traffic received on the Remote interface. Off-path devices use only the Local interface. In hierarchical deployments where both the Local and Remote interfaces are connected to a WAN router, inbound QoS has no effect on incoming WAN traffic on the Local interface.

The following table describes the four traffic classes for inbound bandwidth management.

Table 4-9 Inbound Bandwidth Management Classes

Class	Description
Reduced	Reduced traffic from other WX devices.
Intranet	Unreduced TCP traffic from a specified list of IP subnets. Use the Top Traffic report to help create the list of subnets (refer to “Top Traffic Statistics” on page 317).
TCP	TCP traffic that is not in the Reduced or Intranet class.
Default	All traffic that is not in the Reduced, Intranet, or TCP class.

For example, to enable inbound bandwidth management on D1 in Figure 4-76, set the inbound speed to 1500 Kbps (1.5 Mbps). You then set maximum bandwidth percentages for one or more of the four traffic classes. In this example, you might set the maximum bandwidth percentage for the Default class to 10% to limit low-priority traffic from the public Internet.

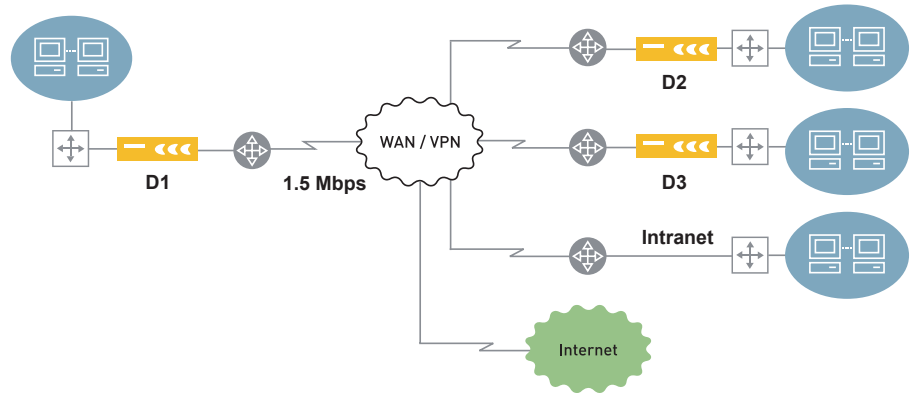


Figure 4-76 Configuring Inbound Bandwidth Management

To configure inbound QoS:

1. In the Configuration window, click **QOS** in the left-hand navigation frame, click **Inbound QoS**, and select the check box.

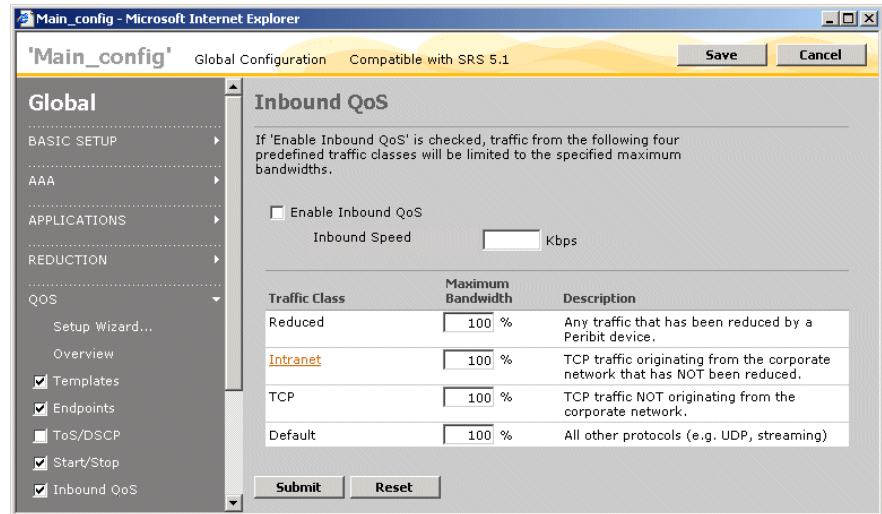


Figure 4-77 Configuring Maximum Inbound QoS Bandwidths

2. To start the inbound QoS service, click **Enable Inbound QoS**.
3. Add up the speeds of all the WAN interfaces on the router that conducts WAN traffic to the device where you intend to load the configuration, and enter the value (in Kbps) in the **Inbound Speed** field.

4. Enter the maximum bandwidth of each traffic class as a percentage of the inbound speed.
5. Click **Submit** to enter the changes, or click **Reset** to discard them.
6. To specify the remote subnets whose traffic belongs to the Intranet class, click **Intranet** and enter the remote subnets (one per line) whose traffic belongs to the Intranet traffic class, and click **Submit**. The subnet format is:

<IP address>/<subnet mask>

7. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Traffic Acceleration

The following topics describe how to configure traffic acceleration:

- “Overview of Packet Flow Acceleration” in the next section
- “Overview of Application Flow Acceleration” on page 193
- “Enabling Acceleration by Endpoint” on page 198
- “Enabling Acceleration by Application” on page 203

Overview of Packet Flow Acceleration

While data reduction effectively increases the available WAN bandwidth, application performance may be further constrained by network latency. Packet Flow Acceleration (PFA) improves the performance of reduced TCP application flows across high-speed, high-latency WAN links. For devices that support Multi-Path, you can enable PFA for the primary and/or secondary paths.

NOTE: PFA is most effective in networks with high-speed connections and high latency. However, PFA may have no effect if the traffic must cross low-speed or high-latency connections that are one or more hops beyond the receiving WX device.

Active Flow Pipelining

Active Flow Pipelining (AFP) is generally the most effective method of TCP acceleration, and is intended primarily for high-latency environments, such as satellite connections, and long-haul high-bandwidth links, such as E3 and T3. AFP is also beneficial when the reduction percentage is very high.

NOTE: AFP is required to use Network Sequence Mirroring (NSM) on Sequence Mirror devices, or to accelerate Microsoft CIFS, Microsoft Exchange, and HTTP traffic using Application Flow Acceleration.

In WAN environments, TCP may restrict the transmission of data (reduces the receive window) because it interprets long wait times for acknowledgements (ACKs) as a sign of network congestion. AFP solves this problem by terminating each TCP session locally. The result is three independent sessions—between the TCP source and the sending device, between the two WX devices, and between the receiving device and the destination.

Since all transmissions are acknowledged locally, more data can be put “in flight” at once. The ACKs are returned to the sender at a rate governed by the speed of the link.

To avoid the TCP congestion mechanism, which is very inefficient over the WAN, a reliable transport protocol ensures in-order delivery between the two WX devices, and provides retransmission when necessary. Congestion is managed by outbound QoS.

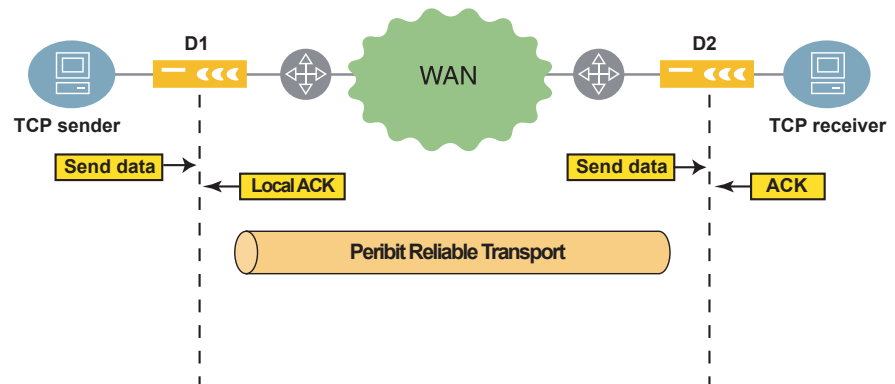


Figure 4-78 Active Flow Pipelining

AFP is intended for applications that do large data transfers. In general, AFP improves performance if the product of the effective bandwidth and latency (the maximum window size) exceeds the TCP window size. Note that 64 KB is the typical TCP window size for Windows 2000 and later (16 KB for Windows 98).

For example, on a T1 link (1.5 Mbps) where the latency is 200 ms, and a 50% data reduction doubles the effective bandwidth, the maximum window size is:

$$(3,088,000 \text{ bps} * 0.200 \text{ seconds})/8 = 77,200 \text{ bytes}$$

In this case, AFP will improve performance if the host's TCP window size is 64 KB or less.

NOTE: Like high bandwidth and latency, high reduction rates also increase the maximum window size, which increases the benefit of AFP.

Asymmetric Routing for AFP

For AFP to accelerate a traffic flow, the flow in both directions must be handled by the same two WX devices. In a load-balancing environment, the two TCP setup packets for a new flow (SYN and SYN ACK) may be seen by different devices. In this case, you can define clusters of devices that advertise their SYN packets so that any device in the cluster that sees the SYN ACK can establish the flow to the sending device.

For more information about using asymmetric routing support, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.

Forward Error Correction

Forward Error Correction (FEC) enables the sending device to send recovery packets along with all data packets, so that the receiving device can reconstruct lost packets without requesting a retransmission. You can specify the number or recovery packets per block of data packets.

FEC is intended for use in high-loss, high-latency environments, such as satellite connections. However, FEC should be disabled if the satellite modem also provides forward error correction. Note that when FEC is enabled for a device, recovery packets are generated for all traffic sent to that device.

After you enable FEC, check the monitoring report periodically. If losses are not persistent, disable FEC to avoid the overhead of processing recovery packets.

Fast Connection Setup

With Fast Connection Setup (FCS), the sending device locally acknowledges the initial session request (the SYN packet) for each new TCP session if the destination is known to be active. FCS saves one round-trip time (RTT) for each

session, and is intended for applications that have many short sessions, such as HTTP 1.0. Short sessions are those that last less than ten times the round-trip time.

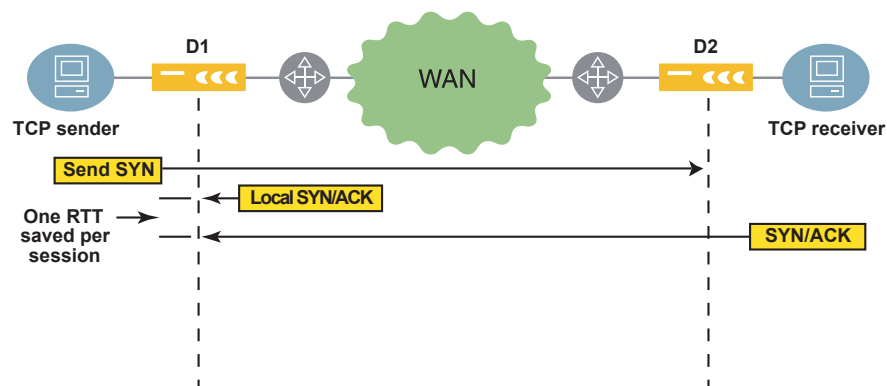


Figure 4-79 Fast Connection Setup

FCS is particularly useful in pre-Windows 2000 environments, where NetBios (not CIFS) is used for file transfer. FCS is also beneficial for HTTP 1.0 traffic (pre-Windows 2000) as it creates more short-lived TCP connections than HTTP 1.1. Some custom enterprise WAN applications may also benefit from FCS.

FCS is most effective in high latency environments, because each RTT that is saved per session represents a larger slice of time as the latency increases. If latency is very low (LAN latencies for example), FCS will not provide much benefit.

FCS is applied only to sessions that last less than ten times the round-trip time (RTT). If a specific application traffic flow has five consecutive short sessions, FCS is applied to all subsequent identical traffic flows. The average session acceleration is calculated as follows:

$$100 - [100 (\text{Accelerated session time}) / (\text{Session time without acceleration})]$$

Note that performance improvements will be more noticeable to users as the percentage of accelerated sessions increases.

Overview of Application Flow Acceleration

Though technologies such as compression (MSR and NSM) and TCP Acceleration (AFP) can greatly increase the performance for applications across the WAN, these benefits may be undermined by inefficient protocols above TCP. To achieve the best end-user performance, specific protocols need to be optimized for the WAN.

The primary purpose for Application Flow Acceleration (AppFlow) is to improve end-user performance for specific business-critical protocols that traverse the WAN. Application Flow Acceleration not only improves performance for existing WAN applications but also facilitates the centralization of branch servers to central data centers.

Currently three business-critical, but WAN-inefficient protocols are accelerated: Microsoft Common Internet File System (CIFS), which is the underlying protocol for Microsoft File Services, traffic between Microsoft Exchange servers and Outlook clients (MAPI over RPC), and Web traffic (HTTP).

If Active Flow Pipelining is enabled for one or more remote endpoints, you can enable application-level acceleration for Microsoft CIFS, Microsoft Exchange, and HTTP traffic sent to those endpoints. You can accelerate all such traffic, or you can create application definitions that let you accelerate traffic to specific servers. Application Flow Acceleration must be enabled on the devices closest to the clients.

NOTE: AppFlow and tunnel switching cannot be enabled on the same device. When AppFlow is enabled, an error occurs if the tunnel switching CLI commands are added to the CLI section of the configuration, or if you load the configuration on a device that has tunnel switching enabled.

Microsoft CIFS and Microsoft Exchange Acceleration

Microsoft CIFS and Microsoft Exchange traffic is accelerated by having the WX device locally acknowledge each block of traffic sent during bulk read/write operations, such as copying files (for CIFS) and sending or receiving Emails with attachments. This allows many data blocks to be in flight at the same time, which speeds up the data transfer. Acceleration benefits begin at relatively low latencies (about 30 ms. round-trip time).

CIFS and Exchange are TCP protocols that transfer bulk data (files or attachments) by breaking up the object into smaller data blocks. CIFS and Exchange write or read one block of data at a time before proceeding to the next block. This serial transmission of small data blocks is a major contributor to slow performance over the WAN.

In read operations (Figure 4-80), the client requests one block of data at a time. The device closest to the client (D1) requests the next *N* blocks. The device closest to the server (D2) locally acknowledges each block from the server and sends them to D1. D1 serves each block to the client as requested.

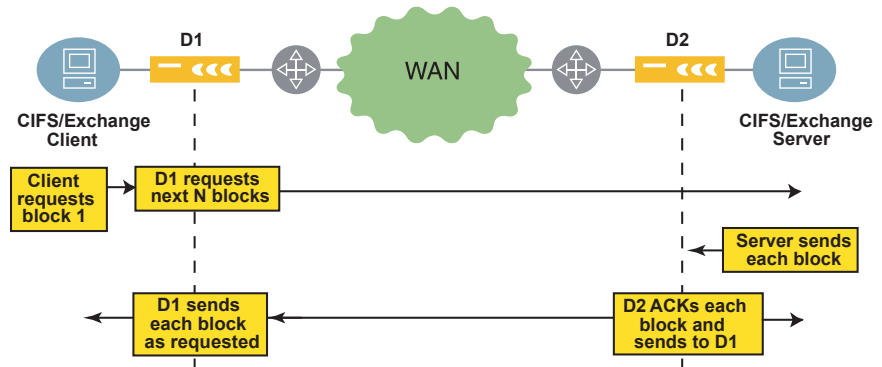


Figure 4-80 Microsoft CIFS/Exchange Read Operations

In write operations (Figure 4-81), the client writes one block at a time. The device closest to the client (D1) acknowledges each block locally, and discards the acknowledgements from the server.

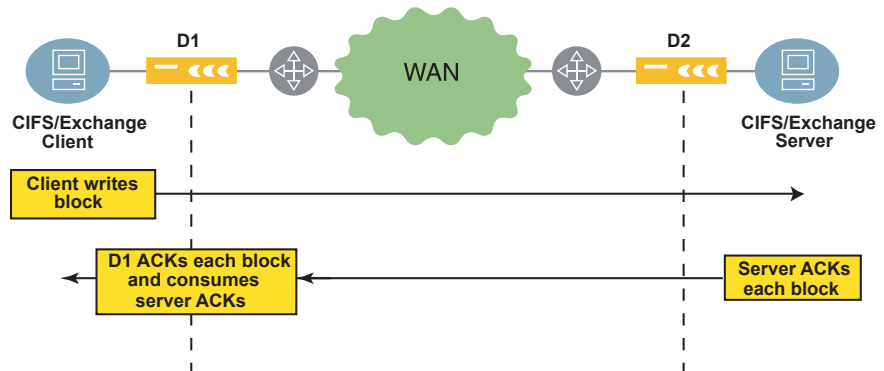


Figure 4-81 Microsoft CIFS/Exchange Write Operations

NOTE: CIFS acceleration is not effective if Server Message Block (SMB) signing is enabled. Signing should be disabled on all servers and clients, and on all domain controllers that are also used as file servers (refer to “[Enabling Microsoft CIFS Acceleration](#)” on page 205). For more information about SMB signing, go to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;887429>

HTTP Acceleration

Two types of application acceleration are available for HTTP traffic:

- **Caching.** Maintains a cache of HTTP responses from HTTP GET requests for the following static objects:

- Cascading style sheets (*.css)
- Static images (*.gif and *.jpeg)
- Java scripts (*.js)

The response cache can contain just response header information (header-only mode) or response headers plus the associated static objects (header-and-body mode). Sequence Reducer devices can cache only HTTP response headers, but Sequence Mirror devices can cache both HTTP response headers and static objects.

In header-only mode, when the browser reloads a Web page and issues a GET IF-MODIFIED-SINCE request to verify that a static object in its cache is still valid, the WX device responds as follows:

- If the object is fresh, a 304 NOT MODIFIED is sent, which saves a round-trip time.
- If the object is not fresh, the request is forwarded to the originating HTTP server.

In header-and-body mode, a Sequence Mirror locally responds to both GET and GET IF-MODIFIED-SINCE requests for the static objects in its cache. Serving cached objects saves at least one round-trip time for each object.

- **Pre-fetching.** After a page is requested once (by any client), a request for the first static object on a page triggers requests for all the page's static objects, which saves one round-trip time for each pre-fetched object. On Sequence Mirror devices, only objects that are considered "stale" by the cache are pre-fetched.

In HTTP cache header-only mode (Figure 4-82), the client sends HTTP GET IF-MODIFIED-SINCE queries before reloading an object from the browser cache. Based on its own caching timer, the device closest to the client (D1) indicates the object has not changed or forwards the query to the server. Even if the query is not forwarded, D1 sends its own query to verify the object's last-modified date.

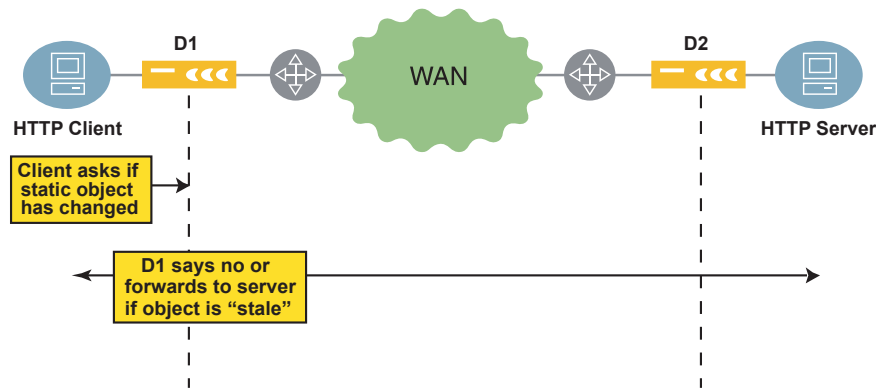


Figure 4-82 HTTP Caching—Header-Only Mode

In HTTP cache header-and-body mode (Figure 4-83), the client sends HTTP GET requests for static objects on a page that has been visited before. The Sequence Mirror device closest to the client (D1) serves the objects directly from its own cache (if they are still fresh) or forwards the requests to the HTTP server.

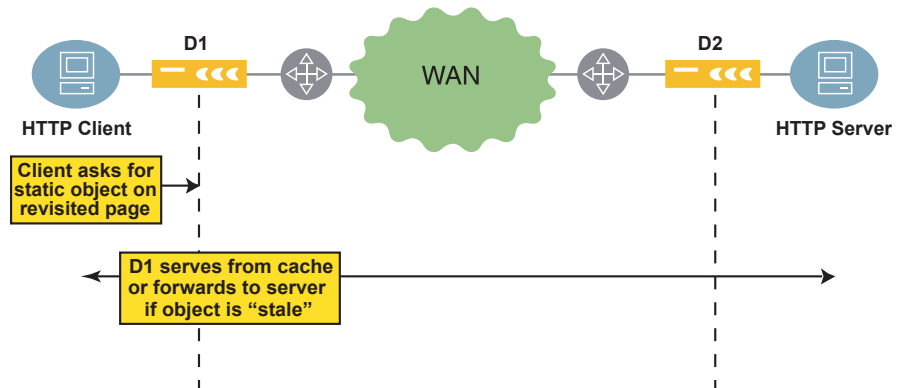


Figure 4-83 HTTP Caching—Header-and-Body Mode (Sequence Mirror Devices)

If pre-fetch is enabled (Figure 4-84), the static objects associated with each page (".css", ".gif", ".jpeg", and ".js") are recorded when the page is first requested. When the first object of a previously seen page is requested again, the device (D1) requests all the static objects that are considered stale. The objects returned by the server are acknowledged locally by D2.

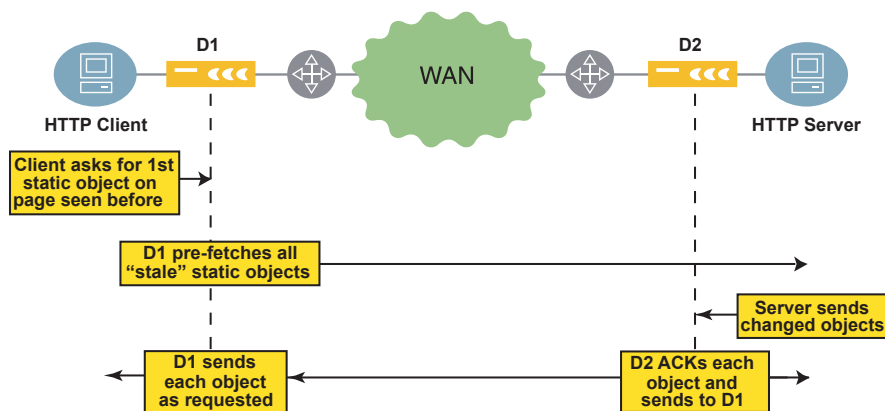


Figure 4-84 HTTP Pre-Fetch

NOTE: HTTP traffic is not accelerated if a proxy server exists between the server-side device and the actual HTTP server. However, if the proxy server is between the Web client and the client-side device, HTTP traffic will be accelerated.

Enabling Acceleration by Endpoint

You can enable each method of Packet Flow Acceleration for all remote devices (endpoints), or for specific endpoints. Active Flow Pipelining must be enabled on both the sending and receiving devices. For other methods, if most of the traffic is in one direction, you can enable just the sending device.

To enable acceleration for a remote endpoint, you must:

- Enable reduction tunnels in both directions between the devices (refer to “Configuring Endpoints for Reduction Tunnels” on page 139).
- Enable reduction for the applications you want to accelerate (refer to “Reducing Applications” on page 144).
- Enable outbound QoS, and specify the WAN circuit speed for each remote device for which you want to accelerate traffic (refer to “Using Outbound QoS to Enhance Performance” on page 155).
- To use Active Flow Pipelining, enable AFP on the Features/Topology page (refer to “Configuring the Feature/Topology Settings” on page 212).

If you enable Active Flow Pipelining or Fast Connection Setup, you must also select the applications that each method is applied to (refer to “Enabling Acceleration by Application” on page 203).

To enable Packet Flow Acceleration by endpoint:

1. In the Configuration window, click **ACCELERATION** in the left-hand navigation frame, click **Overview**, and select the check box.

Main_config - Microsoft Internet Explorer

'Main_config' Global Configuration Compatible with SRS 5.1 [Save] [Cancel]

Global

BASIC SETUP >
AAA >
APPLICATIONS >
REDUCTION >
QoS >
ACCELERATION >
 Overview
 Active Flow Pipelining
 Fast Connection Setup
 CIFS
 HTTP
 Exchange
ADVANCED SETUP >
MULTI-PATH >
IPSEC >

Acceleration Overview [HELP]

Step 1: Enable desired Acceleration capabilities
☒ Active Flow Pipelining ☐ Fast Connection Setup* ☐ Forward Error Correction†

Step 2: Specify how enabled Acceleration capabilities are applied to endpoints
☐ Accelerate all QoS enabled endpoints using default settings
☒ Accelerate checked endpoints using custom settings

Name	IP Address	Active Flow Pipelining	Fast Connection Setup	Forward Error Correction	Data Packets
<input checked="" type="checkbox"/> SR-192.168.71.10	192.168.71.10	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/> SR-192.168.72.10	192.168.72.10	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/> SR-192.168.73.11	192.168.73.11	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/> SR-192.168.74.11	192.168.74.11	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> SR-192.168.75.10	192.168.75.10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Note: QoS must be enabled on an endpoint before it can be accelerated

* Should only be used for connections with applications that generate many very short-lived TCP connections (e.g., HTTP 1.0) across high latency links.

† Should only be used for connections that are subject to high loss and do not have FEC enabled on the satellite modem or CSU/DSU.

[Select All] [Clear]
 [Add/Remove Endpoints] ☒ Show Advanced Settings
 [Submit] [Reset]

Figure 4-85 Enabling Packet Flow Acceleration

2. At the top of the page, select the check box next to each of the PFA methods that you want to use for one or more of the remote endpoints.
3. Select one of the following options:
 - **Accelerate all QoS enabled endpoints using default settings.** Traffic is accelerated to all remote devices for which a reduction tunnel exists and outbound QoS is configured correctly. The PFA methods you select apply to all qualifying endpoints, and to all qualifying endpoints added to the same community in the future.
 - **Accelerate checked endpoints using custom settings.** Traffic is accelerated only to the selected devices, and different PFA methods can be used for each endpoint. Click the check box next to the IP address of the appropriate devices. An endpoint is greyed out if QoS is not enabled for the device.

To add or remove specific remote endpoints for PFA:

- a. Click **Add/Remove Endpoints**.
- b. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable PFA for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 101.
- c. Select the devices you want to accelerate traffic for, and click **Add**. To remove devices from the Acceleration Endpoints list, select the devices and click **Remove**.
- d. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
- e. If one or more devices are not listed, click **Manual Entry** and enter the device IP addresses manually (one per line), and click **Submit**.
- f. When you are done, click **Submit**.

NOTE: When you save a global configuration, an error occurs if QoS and reduction are not enabled for all endpoints using PFA. If you enable PFA for an endpoint in an Acceleration partial configuration, an error occurs if you load the configuration on a device where QoS or reduction is not enabled for that endpoint.

4. Select the PFA methods to be used for each endpoint or for all endpoints:.

Active Flow Pipelining	<p>Intended for high-latency environments, such as satellite connections, long-haul high-bandwidth links, such as E3 and T3, and networks where reduction rates are very high.</p> <p>AFP must be enabled on both the sending and receiving device, and cannot be used simultaneously on the same path with Fast Connection Setup. AFP is required for Network Sequence Mirroring and Application Flow Acceleration.</p> <p>NOTE: In some cases, you may need to do one or more of the following (refer to the “configure acceleration” CLI command in the operator’s guide):</p> <ul style="list-style-type: none"> • Adjust the buffer size for optimum performance. • Increase the number of lost heartbeat packets allowed on high-loss links (reduction may stop when consecutive heartbeat packets are lost). • Enable clustering if the outbound and return traffic does not always traverse the same two WX devices. • If tunnel load balancing is enabled, verify that it is “Flow based” or “Per-destination” (refer to “Configuring Tunnel Load Balancing Policies” on page 147) • For device speeds of 20 Mbps or more, enable fast reduction tunnels for greater throughput if acceleration is more important than reduction (refer to the “config reduction set fast-reduction-tunnel” CLI command).
Fast Connection Setup	<p>Intended for applications that have many short sessions, such as HTTP 1.0 and NetBios. The sending device locally acknowledges session requests for destinations known to be active. Short sessions are those that last less than ten times the round-trip time (RTT).</p>
Forward Error Correction	<p>Intended for high-loss environments. The sending device sends recovery packets with the data to reduce the number of retransmissions required when data packets are lost. By default, one recovery packet is sent for every nine data packets. To change the number of data and recovery packets, click Show Advanced Settings at the bottom of the page.</p> <p>After you enable FEC, check the monitoring report periodically. If losses are not persistent, disable FEC to avoid the overhead required to process recovery packets.</p>

**Recovery Packets
and Data Packets**

Select the number of recovery packets (1 through 5) for the number of data packets (4 through 25). The settings should be based on the WAN error rate, as shown in Table 4-10.

Note the following:

- Increasing the ratio of recovery packets to data packets reduces retransmissions, but requires more overhead. May be useful for losses caused by congestion.
- Data packets must be a multiple of the recovery packets. For one recovery packet, the data packets can be 4 through 25; for 2 recovery packets, the data packets can be 4, 6, 8, and so on through 24.

Table 4-10 Recommended Data and Recovery Packets for FEC

Error Rate	Recovery Packets	Data Packets	Recovery Packet Overhead
6.25%	1	4	25%
5.00%	1	5	20%
4.25%	1	6	17%
3.50%	1	7	14%
3.00%	1	8	13%
2.75%	1	9	11%
2.50%	1	10	10%
2.25% or less	1	11	9%

5. Click **Submit** to enter the changes, or click **Reset** to discard them. You can now enable PFA for specific applications, as described in the next section.

Enabling Acceleration by Application

The following topics describe how to accelerate specific applications.

- “Enabling Active Flow Pipelining by Application” in the next section
- “Enabling Fast Connection Setup by Application” on page 204
- “Enabling Microsoft CIFS Acceleration” on page 205
- “Enabling Microsoft Exchange Acceleration” on page 208
- “Enabling HTTP Acceleration” on page 210

Enabling Active Flow Pipelining by Application

After you enable AFP as described in “Enabling Acceleration by Endpoint” on page 198, you can select the applications whose outgoing traffic you want to accelerate. To enable AFP for one or more applications:

1. In the Configuration window, click **ACCELERATION** in the left-hand navigation frame, click **Active Flow Pipelining**, and select the check box.

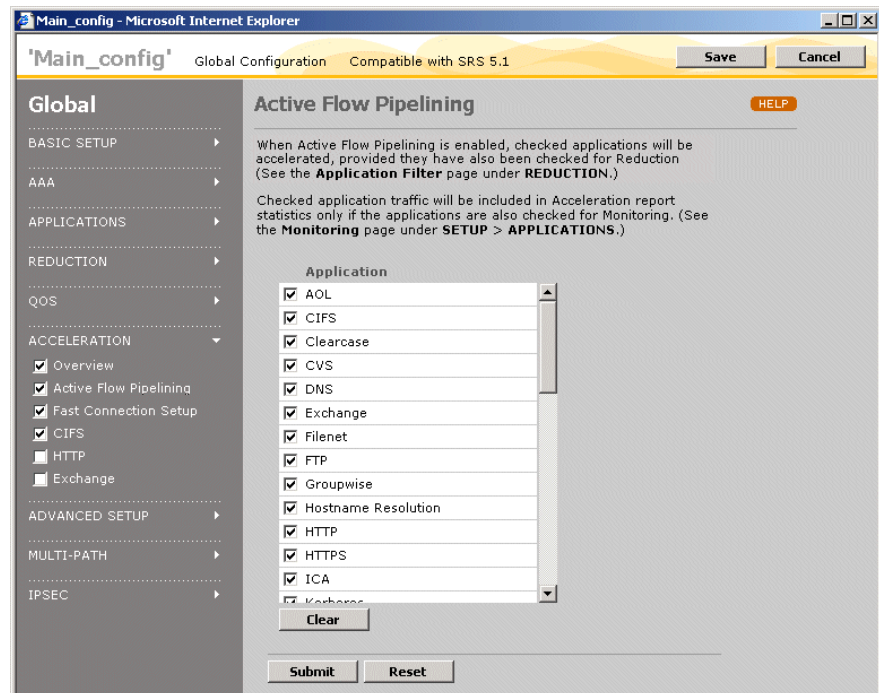


Figure 4-86 Enabling Active Flow Pipelining by Application

2. Select the check box next to each application that you want to accelerate using AFP. To disable AFP for all applications, click **Clear**. The selected applications are accelerated only if they are also being reduced (refer to “Reducing Applications” on page 144).

NOTE: AFP must be enabled on both the sending and receiving devices.

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Enabling Fast Connection Setup by Application

After you enable Fast Connection Setup, as described in “Enabling Acceleration by Endpoint” on page 198, you can select the applications whose outgoing traffic you want to accelerate. Fast Connection Setup is intended for applications that have many short sessions, such as HTTP 1.0.

To enable Fast Connection Setup for one or more applications:

1. In the Configuration window, click **ACCELERATION** in the left-hand navigation frame, click **Fast Connection Setup**, and select the check box.

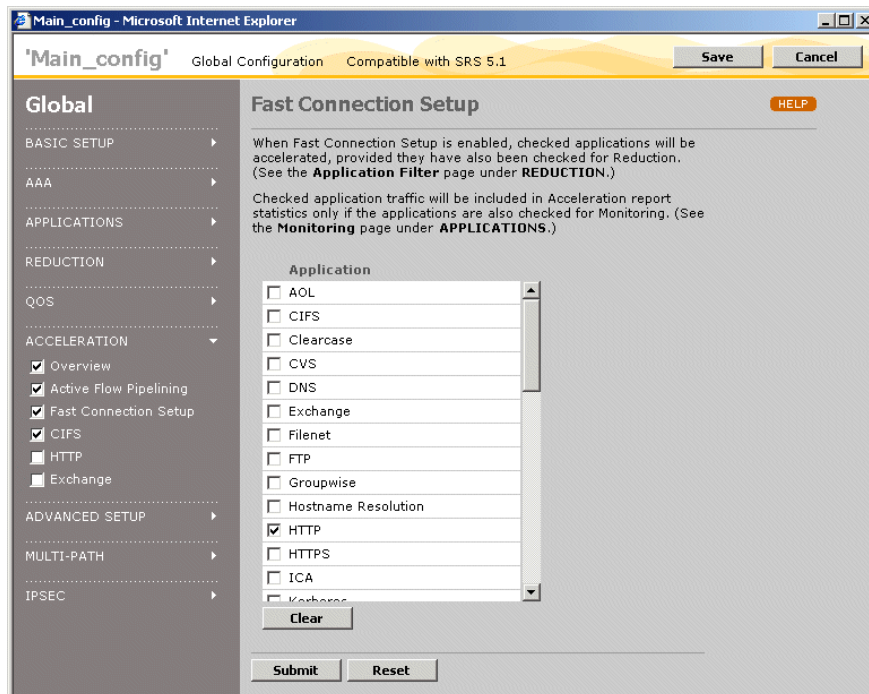


Figure 4-87 Enabling Fast Connection Setup by Application

2. Select the check box next to each application that you want to accelerate using Fast Connection Setup. To disable Fast Connection Setup for all applications, click **Clear**. The selected applications are accelerated only if they are also being reduced (refer to “Reducing Applications” on page 144)
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

The selected applications have no effect if you load the configuration on a device where Fast Connection Setup is not enabled.

Enabling Microsoft CIFS Acceleration

You can accelerate all CIFS traffic using the default CIFS application definition, or you can create multiple application definitions to accelerate selected CIFS traffic, such as the traffic to or from a specific server.

Microsoft CIFS traffic between Windows 2000 or XP clients and Windows 2000 or 2003 servers is accelerated. Enable CIFS acceleration on the devices closest to the clients (not required on the devices closest to the servers).

Note the following:

- Any new CIFS application definitions created must have an application type of CIFS and port numbers 139 and 445 (refer to “Configuring Application Definitions” on page 131)
- Active Flow Pipelining must be enabled on both the client- and server-side devices (refer to “Enabling Active Flow Pipelining by Application” on page 203).
- Application Flow Acceleration must be enabled on the client-side device (select **All features** on the Features/Topology page, as described in “Configuring the Feature/Topology Settings” on page 212). On the server-side device, you can conserve system resources by selecting **All features except Application Flow Acceleration** on the Features/Topology page.

To enable CIFS acceleration for one or more applications:

1. To add new CIFS application definitions to accelerate specific CIFS traffic:
 - a. Under **APPLICATIONS** in a global or Applications partial configuration, click **Definitions** in the left-hand navigation frame, select the check box, and then click **New Applications**.
 - b. Select the **CIFS** application type, and be sure to specify port numbers 139 and 445. Complete the definition, and click **Submit**.

- c. On the Application Definitions page, the new definition receives the order number of the generic CIFS definition. For example, if the order number of the generic definition was 6, the new definition becomes 6 and all subsequent definitions are incremented.
2. To enable acceleration for CIFS applications, under **ACCELERATION** in a global or Acceleration partial configuration, click **CIFS** in the left-hand navigation frame, and select the check box.

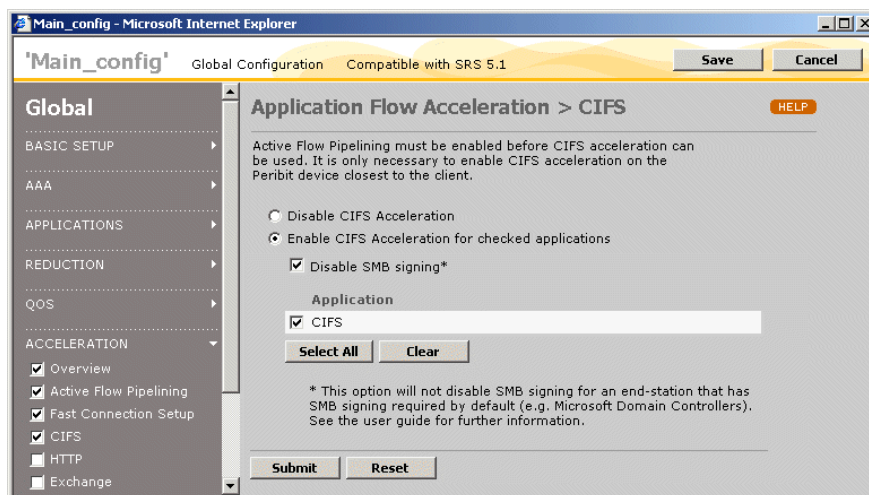


Figure 4-88 Enabling CIFS Acceleration

3. Select **Enable CIFS Acceleration for checked applications** and click the check box next to the appropriate applications, or click **Select All**. You can select only applications that have an application type of CIFS and are enabled for Active Flow Pipelining.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.
5. Verify that SMB signing is disabled on Windows 2000 and Windows 2003 domain controllers that are also file servers. On a Windows 2000 domain controller:
 - a. Open **Active Directory Users and Computers** on the domain controller.
 - b. Right click **Domain Controllers** and select **Properties**.
 - c. Click the **Group Policy** tab.
 - d. Click **Default Domain Controllers Policy** and select **Edit**.

- e. Click **Default Domain Controllers Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options**.
 - f. Disable the four signing options:
 - **Digitally sign client communication (always)**
 - **Digitally sign client communication (when possible)**
 - **Digitally sign server communication (always)**
 - **Digitally sign server communication (when possible)**
 - g. Reboot all domain controllers, member servers, and clients for which you want to accelerate CIFS traffic.
6. To verify that SMB signing is disabled on Windows 2000 clients:
- a. Click **Start/Settings/Control Panel** and select **Administrative Tools**.
 - b. Select **Local Security Policy**, and then select **Local Policies/Security Options**.
 - c. Disable the four signing options:
 - **Digitally sign client communication (always)**
 - **Digitally sign client communication (when possible)**
 - **Digitally sign server communication (always)**
 - **Digitally sign server communication (when possible)**

Enabling Microsoft Exchange Acceleration

You can accelerate all Exchange traffic using the default Exchange application definition, or you can create multiple application definitions to accelerate selected Exchange traffic, such as the traffic to or from a specific server.

Microsoft Exchange traffic between the following platforms is accelerated:

- Windows 2000 or XP clients and Windows 2000 or 2003 servers
- Outlook 2000, 2002 or 2003 clients and Exchange 5.5, 2000 or 2003 servers

NOTE: Traffic between an Outlook 2003 client and Exchange 2003 server is not accelerated, but Sequence Mirror devices using NSM disk-based compression provide substantial benefits for such traffic without acceleration. Also, Exchange 2003/Outlook 2003 use Microsoft Exchange compression by default, which should be disabled (refer to <http://support.microsoft.com/?kbid=825371>).

Enable Exchange acceleration on the devices closest to the clients (not required on the devices closest to the servers).

Note the following:

- Any new Exchange application definitions created must have an application type of Exchange and port number 135 (refer to “Configuring Application Definitions” on page 131)
- Active Flow Pipelining must be enabled on both the client- and server-side devices (refer to “Enabling Active Flow Pipelining by Application” on page 203).
- Application Flow Acceleration must be enabled on the client-side device (select **All features** on the Features/Topology page, as described in “Configuring the Feature/Topology Settings” on page 212). On the server-side device, you can conserve system resources by selecting **All features except Application Flow Acceleration** on the Features/Topology page.

To enable Exchange acceleration for one or more applications:

1. To add new Exchange application definitions to accelerate specific Exchange traffic:
 - a. Under **APPLICATIONS** in a global or Applications partial configuration, click **Definitions** in the left-hand navigation frame, select the check box, and then click **New Applications**.
 - b. Select the **Exchange** application type, and be sure to specify port number 135. Complete the definition, and click **Submit**.
 - c. On the Application Definitions page, the new definition receives the order number of the generic Exchange definition. For example, if the order number of the generic definition was 20, the new definition becomes 20 and all subsequent definitions are incremented.
2. To enable acceleration for Exchange applications, under **ACCELERATION** in a global or Acceleration partial configuration, click **Exchange** in the left-hand navigation frame, and select the check box.

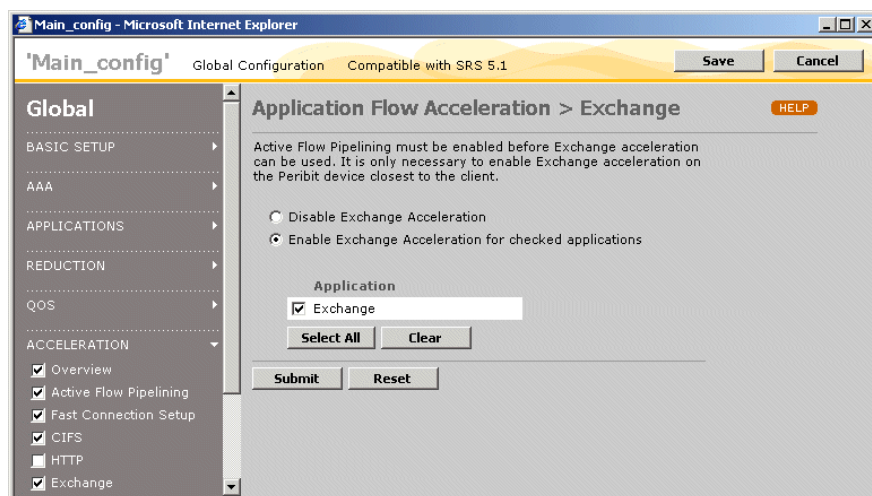


Figure 4-89 Enabling Exchange Acceleration

3. Select **Enable Exchange Acceleration for checked applications** and click the check box next to the appropriate applications, or click **Select All**. You can select only applications that have an application type of Exchange and are enabled for Active Flow Pipelining.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Enabling HTTP Acceleration

You can accelerate all HTTP traffic using the default HTTP application definition, or you can create multiple application definitions to accelerate selected HTTP traffic, such as the traffic to or from a specific server.

Enable HTTP acceleration on the devices closest to the clients (not required on the devices closest to the servers).

Note the following:

- Any new HTTP application definitions created must have an application type of HTTP and the correct port number (refer to “Configuring Application Definitions” on page 131)
- Active Flow Pipelining must be enabled on both the client- and server-side devices (refer to “Enabling Active Flow Pipelining by Application” on page 203).
- Application Flow Acceleration must be enabled on the client-side device (select **All features** on the Features/Topology page, as described in “Configuring the Feature/Topology Settings” on page 212). On the server-side device, you can conserve system resources by selecting **All features except Application Flow Acceleration** on the Features/Topology page.

NOTE: HTTP traffic is not accelerated if a proxy server exists between the server-side device and the actual HTTP server. However, if the proxy server is between the Web client and the client-side device, HTTP traffic will be accelerated.

To enable HTTP acceleration for one or more applications:

1. To add new HTTP application definitions to accelerate specific HTTP traffic:
 - a. Under **APPLICATIONS** in a global or Applications partial configuration, click **Definitions** in the left-hand navigation frame, select the check box, and then click **New Applications**.
 - b. Select the **HTTP** application type, and be sure to specify the HTTP port number (usually 80). Complete the definition, and click **Submit**.
 - c. On the Application Definitions page, the new definition receives the order number of the generic HTTP definition. For example, if the order number of the generic definition was 4, the new definition becomes 4 and all subsequent definitions are incremented.

2. To enable acceleration for HTTP applications, under **ACCELERATION** in a global or Acceleration partial configuration, click **HTTP** in the left-hand navigation frame, and select the check box.

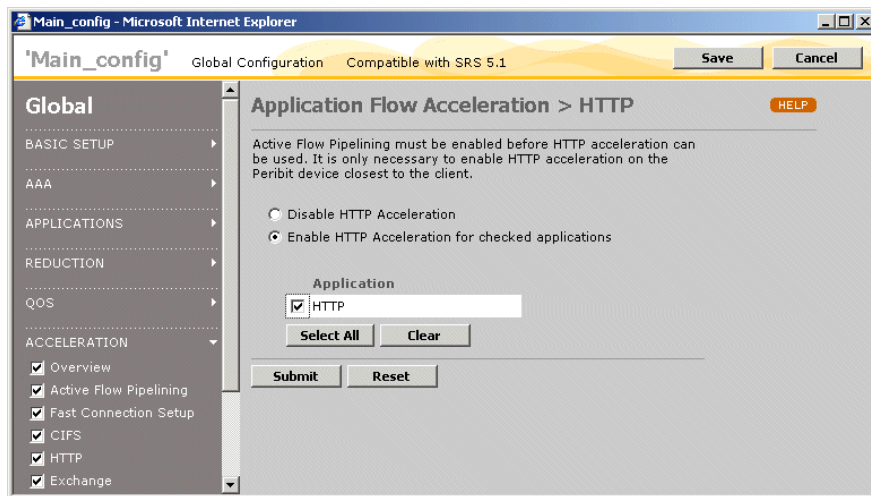


Figure 4-90 Enabling HTTP Acceleration

3. Select **Enable HTTP Acceleration for checked applications** and click the check box next to the appropriate applications, or click **Select All**. You can select only applications that have an application type of HTTP and are enabled for Active Flow Pipelining.
4. Click **Submit** to enter the changes, or click **Reset** to discard them.
5. To enable caching of static objects on a Sequence Mirror device:

On Sequence Mirror devices, static objects are cached by default (header-and-body mode). To change the cache setting, refer to the “configure acceleration CLI command in the operator’s guide.

Configuring Advanced Setup Parameters

The following sections describe the global advanced setup parameters:

- “Configuring the Feature/Topology Settings” on page 212
- “Configuring Source/Destination Filters” on page 216
- “Defining the Prime Time” on page 218
- “Configuring Packet Interception” on page 220
- “Configuring WAN Performance Monitoring” on page 231
- “Adding CLI Commands to Configurations” on page 234

Configuring the Feature/Topology Settings

The features/topology settings specify whether Active Flow Pipelining (AFP) and Application Flow Acceleration can be used, and the topology setting that describes the device’s relationship to the other devices in the same community. These settings ensure that the device’s resources are used efficiently. The topology setting can be Mesh or Hub and Spoke.

NOTE: If you do not specify the Features/Topology settings, they default to all features except Application Flow Acceleration, and a mesh topology with the lowest range of devices.

In a Mesh topology, each device can reduce and accelerate traffic for all other devices in the community. In a Hub and Spoke topology, a central device (hub) processes traffic for all other devices, but spoke devices, by default, process traffic only for the hub (to enable reduction tunnels between spokes, refer to “Configuring Endpoints for Reduction Tunnels” on page 139).

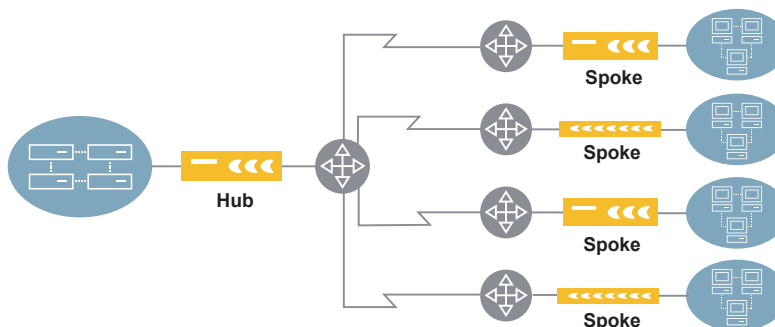


Figure 4-91 Deploying Devices in a Hub and Spoke Topology

For Hub and Mesh topologies, you must specify the range of devices in the community. Table 4-11 and Table 4-12 show the numbered ranges of devices supported by each type of device in Hub and Mesh topologies, based on whether Application Flow Acceleration (APP) and Active Flow Pipelining (AFP) are enabled. The **max-mem** option allocates all available memory for a limited number of tunnels, but all devices must be the same model and have the same topology setting (Hub or Mesh).

Table 4-11 Mesh Device Ranges

Device	Mesh Ranges		
	No APP/AFP	No APP	All Features
SR-15	0=Up to 2	0=Up to 2	0=Up to 1
SR-20	0=Up to 5 1=Up to 10 2=Up to 15 max-mem=3	0=Up to 3 1=Up to 8 2=Up to 10 max-mem=2	0=Up to 2 1=Up to 5 2=Up to 7 max-mem=1
SM-250	0=Up to 15 1=Up to 15 2=Up to 15 max-mem=15	0=Up to 5 1=Up to 10 2=Up to 15 max-mem=5	0=Up to 2 1=Up to 5 2=Up to 10 max-mem=1
SM-500 (use higher value for SR-100 clients)	0=Up to 60 1=Up to 60 2=Up to 60 3=Up to 60 4=Up to 60 5=Up to 60 max-mem=20	0=Up to 15 (20) 1=Up to 20 (25) 2=Up to 25 (30) 3=Up to 30 (40) 4=Up to 40 (50) 5=Up to 50 (60) max-mem=3 (5)	0=Up to 10 (20) 1=Up to 15 (25) 2=Up to 20 (30) 3=Up to 25 (40) 4=Up to 30 (50) 5=Up to 40 (60) max-mem=1 (5)
SR-50 SR-55	0=Up to 20 1=Up to 35 2=Up to 50 3=Up to 60 4=Up to 70 5=Up to 80 max-mem=5	0=Up to 20 1=Up to 30 2=Up to 40 3=Up to 50 4=Up to 60 5=Up to 70 max-mem=4	0=Up to 15 1=Up to 25 2=Up to 35 3=Up to 45 4=Up to 50 5=Up to 55 max-mem=3
SR-80 SR-100	0=Up to 60 1=Up to 100 2=Up to 140 3=Up to 170 4=Up to 200 5=Up to 220 max-mem=20	0=Up to 60 1=Up to 90 2=Up to 130 3=Up to 150 4=Up to 170 5=Up to 190 max-mem=15	0=Up to 60 1=Up to 90 2=Up to 120 3=Up to 140 4=Up to 160 5=Up to 180 max-mem=15
SR-100 with clients	The maximum number of devices is the sum of the ranges. If all features are used, and you select range 4 for an SR-100 that has two SM-500 clients, the top value is 260 (160 + 50 + 50).		

Table 4-12 Hub Device Ranges

Device	Hub Ranges		
	No APP/AFP	No APP	All Features
SR-15	0=Up to 2	0=Up to 2	0=Up to 1
SR-20	0=Up to 5 1=Up to 10 2=Up to 15 max-mem=3	0=Up to 3 1=Up to 8 2=Up to 10 max-mem=2	0=Up to 2 1=Up to 5 2=Up to 7 max-mem=1
SM-250	0=Up to 15 1=Up to 15 2=Up to 15 max-mem=15	0=Up to 5 1=Up to 10 2=Up to 15 max-mem=5	0=Up to 2 1=Up to 5 2=Up to 10 max-mem=1
SM-500 (use higher value for SR-100 clients)	0=Up to 60 1=Up to 60 (85) 2=Up to 60 (85) 3=Up to 60 (85) 4=Up to 60 (85) 5=Up to 60 (85) max-mem=20	0=Up to 15 (20) 1=Up to 20 (25) 2=Up to 25 (30) 3=Up to 30 (40) 4=Up to 40 (55) 5=Up to 60 (85) max-mem=3 (5)	0=Up to 10 (20) 1=Up to 15 (25) 2=Up to 20 (30) 3=Up to 25 (40) 4=Up to 35 (55) 5=Up to 50 (85) max-mem=1 (5)
SR-50 SR-55	0=Up to 20 1=Up to 40 2=Up to 60 3=Up to 80 4=Up to 100 5=Up to 120 max-mem=5	0=Up to 20 1=Up to 35 2=Up to 50 3=Up to 65 4=Up to 80 5=Up to 100 max-mem=4	0=Up to 15 1=Up to 30 2=Up to 40 3=Up to 55 4=Up to 70 5=Up to 85 max-mem=3
SR-80 SR-100	0=Up to 60 1=Up to 120 2=Up to 170 3=Up to 220 4=Up to 270 5=Up to 320 max-mem=20	0=Up to 60 1=Up to 110 2=Up to 150 3=Up to 190 4=Up to 230 5=Up to 280 max-mem=15	0=Up to 60 1=Up to 110 2=Up to 140 3=Up to 180 4=Up to 220 5=Up to 260 max-mem=15
SR-100 with clients	The maximum range of devices is the sum of the ranges. For example, if all features are used, and you select range 4 for an SR-100 that has two SM-500 clients, the top value is 330 (220 + 55 + 55).		

NOTE: If SM and SR devices are in the same community, on SM devices you should select one higher device range than needed to accommodate all combinations of device types. Also, on SR devices, when the maximum range of devices is selected for a hub (range 5), the hub conserves memory by not assembling data from the spokes—only data sent from the hub to the spokes is reduced. In this case, tunnel switching cannot be enabled on the hub.

To change the community topology settings:

1. In the Configuration window, click **ADVANCED SETUP** in the left-hand navigation frame, click **Features/Topology**, and select the check box.

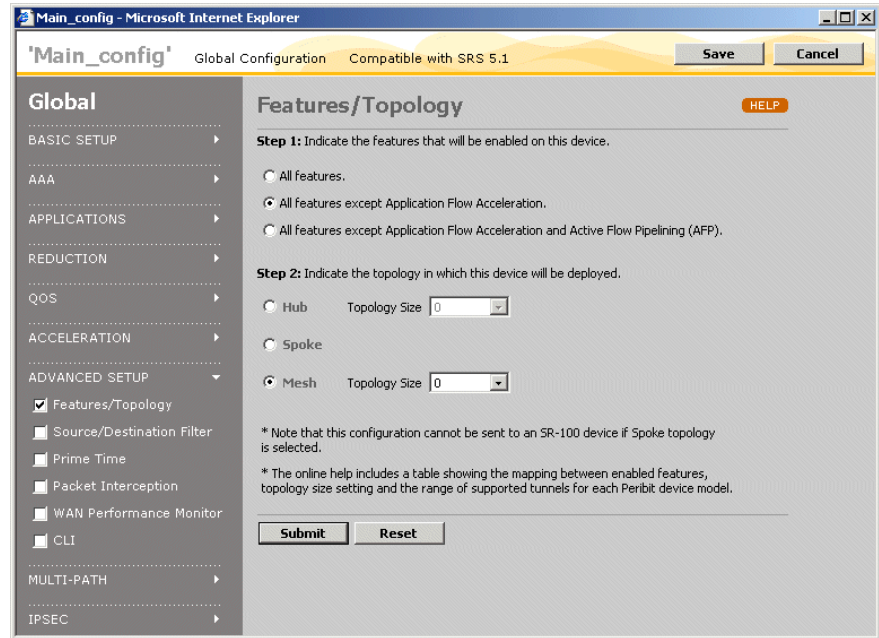


Figure 4-92 Changing the Topology Settings

2. Select the set of features to be used. If this setting is changed on a device, all reduction tunnels will be reset.

All Features	Allows all licensed features to be used.
All features except Application Flow Acceleration	Allows the basic features and Active Flow Pipelining to be used, but not CIFS, Exchange, and HTTP acceleration. Note that Sequence Mirror devices require AFP for optimum performance.
All features except Application Flow Acceleration and AFP	Allows all licensed features to be used, except Active Flow Pipelining and CIFS, Exchange, and HTTP acceleration. Excluded features cannot be accessed from the Web console.

3. Select one of the following topology settings:

Hub	In a Hub and Spoke topology, a hub can reduce and accelerate traffic for all devices in the community. Select the range of devices in the community. If a community has multiple hubs, each hub must specify the same range of devices.
Spoke	By default, a spoke can reduce and accelerate traffic only for devices that are designated as hubs. To enable reduction between spoke devices, refer to “Configuring Endpoints for Reduction Tunnels” on page 139. Note that an SR-100 cannot be a spoke.
Mesh	In a Mesh topology, each device can reduce and accelerate traffic for all other devices in the community. Select the range of devices in the community.

NOTE: Selecting an accurate device range allows each device to allocate its resources efficiently. Mixing hub and spoke with mesh designations in the same community is not recommended.

4. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Source/Destination Filters

You can create a list of source and destination addresses or subnet pairs that are either included or excluded from data reduction. This source/destination filter applies to all application traffic sent from the LAN to the WAN. To enable or disable data reduction by application, refer to “Reducing Applications” on page 144. The source/destination filter is applied before the application filter, and is more efficient.

NOTE: If you disable data reduction between a source and destination, Packet Flow Acceleration between those points is also disabled.

For example, to disable data reduction for all traffic from a local subnet, create a “Do not reduce” entry and specify the subnet as the source and enter an asterisk (*) as the destination. To disable data reduction for all traffic sent to the subnet from other devices, you must disable the advertisement of the subnet (refer to “Advertising Reduction Subnets” on page 94).

Note the following:

- If you disable data reduction between a source and destination, traffic between those points cannot be accelerated. Also, the traffic is managed by the Outbound QoS policies defined for the Default traffic class under the “Other traffic” endpoint.
- Source/destination filters are disallowed on off-path devices that use RIP for packet interception. Also, they should not be used with the External packet interception mode.

To define source and destination subnets:

1. In the Configuration window, click **ADVANCED SETUP** in the left-hand navigation frame, click **Source/Destination Filter**, and select the check box.

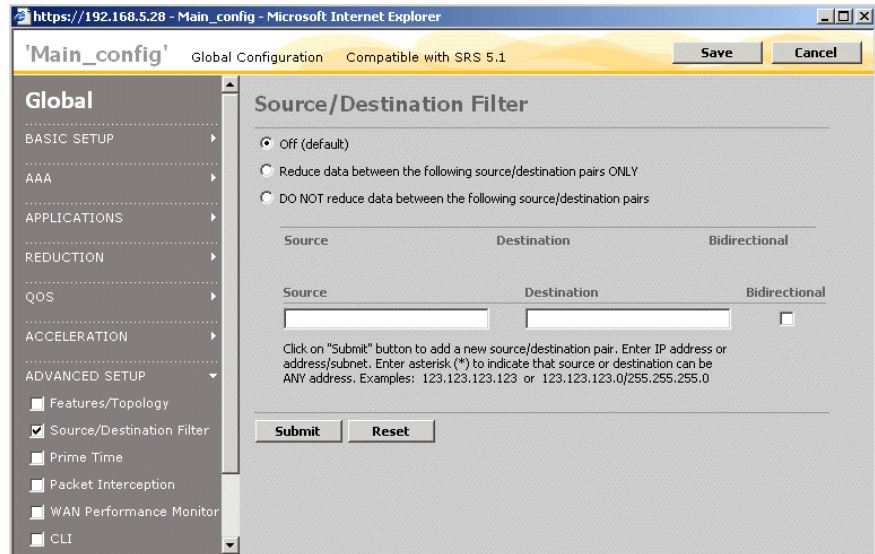


Figure 4-93 Filtering Data Reduction by Source and Destination

2. Select the type of source/destination filter you want to create.
 - **Off (default).** Data is reduced for all eligible application traffic from all local routes to all remote routes advertised by the other WX devices.
 - **Reduce data between the following source/destination pairs ONLY.** Data is reduced only for the specified source and destination pairs. Specify at least one address pair.

- **DO NOT reduce data between the following source/destination pairs.**
All data is reduced, except for traffic between the specified source and destination pairs (the traffic cannot be accelerated, and is managed by the outbound QoS policies defined for the Default traffic class under the “Other traffic” endpoint).

3. Specify the following information:

Source	Enter a source IP address or subnet. The general format is: <code>address/subnetmask</code> The default subnet mask is “255.255.255.255”. An asterisk (*) with no subnet mask indicates any source IP address.
Destination	Enter a destination IP address or subnet (same format as the source address). An asterisk (*) indicates any destination IP address.
Bidirectional	Select the check box to include traffic sent from the destination to the source. This option is particularly useful for creating “do not reduce” lists in Profile Mode. In Profile Mode, you should exclude all traffic sent to the subnet where the device is installed. For more information about Profile Mode, refer to the <i>Sequence Reducer/Sequence Mirror/Sequence Mirror Operator’s Guide</i> .

4. Click **Submit to enter the changes, or click **Reset** to discard them.**

Defining the Prime Time

The prime time setting lets you specify the days of the week and hours of the day when network performance is most important. The prime time can be used to filter performance statistics, and to specify bandwidth management policies for prime-time and non prime-time hours.

NOTE: The prime time can be used to filter reports in the device Web console, but not in CMS.

For example, to view reduction and acceleration statistics during business hours, you can set the prime time to 9:00 AM to 5:00 PM on Monday through Friday. Prime time is disabled by default, which means the effective “prime time” is 24-hours a day, seven days a week.

To define the prime time:

1. In the Configuration window, click **ADVANCED SETUP** in the left-hand navigation frame, click **Prime Time**, and select the check box.

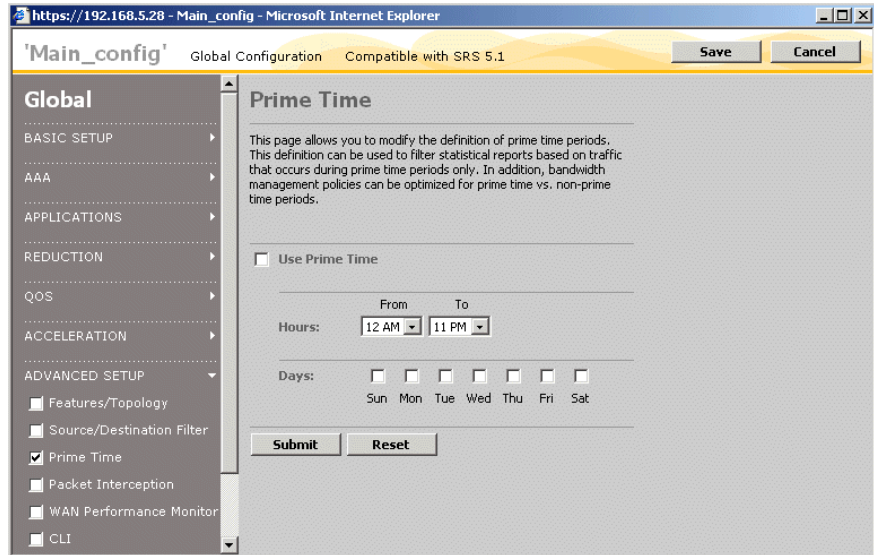


Figure 4-94 Defining the Prime Time

2. To set the prime time, select the **Use Prime Time** check box, select a time range, and select the days of the week.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Packet Interception

For “off-path” devices, you can configure one of three methods of packet interception. Devices are usually deployed in the data path between a LAN switch and a WAN edge router. When interrupting the data path is not practical, such as in collapsed backbone environments, you can deploy devices off path (Figure 4-95).

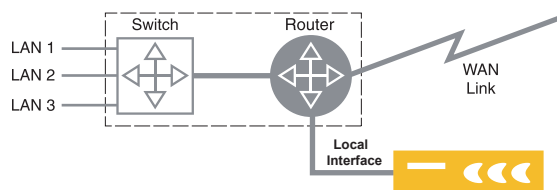


Figure 4-95 Off-Path Deployment

In an off-path deployment, the Local interface is connected to the switch or the router, and the Remote interface is not used (connecting the Local interface directly to the router is recommended).

The following topics describe how to configure packet interception. A few alternatives to packet interception are also described.

- “Configuring Packet Interception for Off-Path Devices” in the next section
- “RIP Router/Switch Configuration Commands” on page 224
- “WCCP Router Configuration Commands” on page 228
- “External Policy-Based Router Commands” on page 229
- “Alternatives to Packet Interception” on page 230

Configuring Packet Interception for Off-Path Devices

In an off-path deployment, the traffic to be reduced must be routed to the device using packet interception. Both the router and the SR/SM device must be configured using one of the following methods of packet interception:

- **Route injection.** The Routing Information Protocol (RIPv2) is used to advertise the off-path device as the lowest cost “router” for all the reduction subnets advertised by the other devices in the community. Note the following:
 - If a remote device advertises its own subnet for reduction, the off-path device generates several new subnets to exclude (carve out) the IP address of the remote device. This prevents the router from returning the traffic sent to the remote device.
 - If a remote device goes down, or carves out a reduction subnet or host, RIP updates are sent immediately to the adjacent router to ensure fast convergence.
 - The off-path device has no passthrough data. Both reduced and unreduced traffic is sent through the reduction tunnel.

To configure a router to use RIP routes, refer to the sample router commands in “RIP Router/Switch Configuration Commands” on page 224.

- **WCCP.** The Web Cache Communication Protocol is used to redirect traffic by protocol from the router to the off-path device. The router must support WCCP version 2. To configure a router to use WCCP, refer to the sample router commands in “WCCP Router Configuration Commands” on page 228.
- **External.** The WAN edge router is configured to route traffic to the off-path device. The off-path device should be connected directly to the router, and must be the only device on the port. You can also connect the off-path device to a dedicated VLAN on a Layer 3 switch. Refer to the sample router commands in “External Policy-Based Router Commands” on page 229.

In each case, the redirected traffic is reduced (if eligible) and returned to the router or switch over the Local interface. Note that for off-path deployments, outbound bandwidth management is limited to the WAN traffic that is routed through the WX device. Also, off-path devices do not support multi-node configurations, but an SR-100 with up to six client devices can be installed off path.

To configure packet interception for an off-path device:

1. In the Configuration window, click **ADVANCED SETUP** in the left-hand navigation frame, click **Packet Interception**, and select the check box.

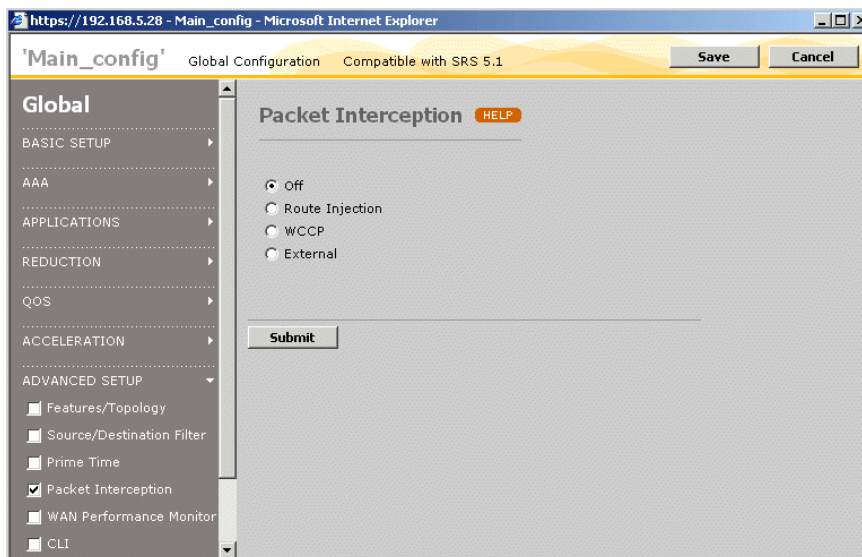


Figure 4-96 Configuring Packet Interception

2. Select one of the following methods of packet interception:

CAUTION: Enabling packet interception disables the Remote interface. If the device is installed in the data path, all data transmission through the device will stop.

- To use RIPv2 for packet interception, click **Route Injection**, and specify the following:

Authentication Type	If the WAN edge router uses RIP authentication, click Password and enter the RIP password. This is the same password used to discover dynamic routes.
Inter-packet delay	To reduce the load on slower routers, enter the number of milliseconds between each packet when multiple packets are generated for a single RIP update (0 through 50). The default is 0.

You can lower the RIP update timers to reduce the failover time (not recommended if RIP is used for network-wide routing). To change the frequency of RIP updates or the cost assigned to each advertised route, refer to the “configure packet-interception” CLI command.

- To use WCCP for packet interception, click **WCCP**, and specify the following:

Router IP Address	Enter the IP address of the WAN edge router (the router must support WCCP version 2).
WCCP Priority	Enter a number (0 through 255) that indicates the order in which packets are compared against the selected services (protocols), relative to the other services redirected by the router. Higher values have a higher priority. The default is 230. For example, if the router is redirecting HTTP traffic to a Web cache using priority 240, and you want to redirect all TCP traffic to the off-path device, specify a lower value to avoid “stealing” traffic from the Web cache.
WCCP Auth. Password	If the WAN edge router uses WCCP authentication, enter the WCCP password specified on the router.

Specify the following for each service (up to five):

IP Protocol	Select a protocol whose traffic you want redirected to the off-path device. You can also type in a protocol number (0 through 255). The standard protocol numbers are defined at: http://www.iana.org/assignments/protocol-numbers
WCCP Service ID	Enter a service ID number for the protocol (51 through 99). The ID must be unique among all the WCCP services defined on the router. In high-availability environments, where two off-path devices use the same router, they must use different IDs for the same protocol. Heartbeat packets are sent to the router every 10 seconds for each service. If the off-path device fails, the router stops redirecting traffic in 30 seconds.

- To configure packet interception by defining routing policies on the router, click **External**. Refer to the sample router commands in “External Policy-Based Router Commands” on page 229.

3. Click **Submit** to enter the changes.

4. Review the reduction subnets and be sure to advertise only the subnets on the LAN side of the off-path device (refer to “Advertising Reduction Subnets” on page 94). Since only the Local interface is connected to the network, the device cannot distinguish between LAN- and WAN-side subnets.

CAUTION: If you use RIP for packet interception, and you have multiple remote WX devices installed on the same subnet, you must disable advertisement of the local subnet on all (or all but one) of the remote devices. Otherwise, the off-path device cannot carve out the remote device addresses, and all traffic sent to them is returned by the router.

The following sections provide sample router configuration commands to support each method of packet interception.

RIP Router/Switch Configuration Commands

In general, an off-path device should be connected to a dedicated port on a router or Layer 3 switch. RIP is then configured on the router or switch where the device is connected. If the off-path device is connected to a Layer 2 switch, RIP is configured on the router. In each case, the RIP configuration is essentially the same.

Single Layer 3 Switch

The following commands provide an example of how to configure RIP on a Layer 3 Cisco switch (Figure 4-97). Installing the off-path device on a dedicated VLAN is recommended to reduce the routing failover time if the device fails. The port where the off-path device is connected should be the only port in the VLAN. Note that the load balancing done by the switch across the two routers is not affected.

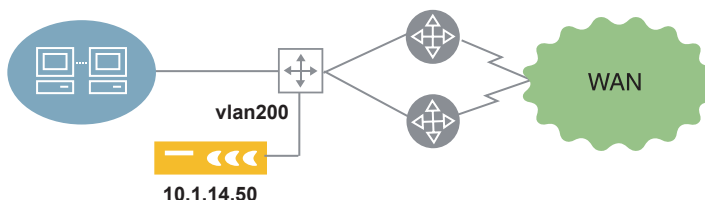


Figure 4-97 Off-Path Device Connected to a Layer 3 Switch

1. Enable RIP version 2:

```
router rip
version 2
```

2. If RIP is used only for packet interception, you can lower the RIP timers to reduce the failover time (may cause instability if RIP is used for network-wide routing):

```
timers basic 5 15 15 30
```

3. Enable RIP to listen passively on all interfaces:

```
passive-interface default
```

4. Specify the subnet where the off-path device is installed:

```
network 10.0.0.0
```

5. Specify the RIP administrative distance to be lower than all other methods used by the router or switch to discover routes (such as OSPF):

```
distance 30
```

6. Disable auto-summarization of routes:

```
no auto-summary
```

Do not redistribute the RIP routes to any other routing protocol, such as OSPF. The advertised RIP routes apply only to the configured router or switch and the off-path WX device. If RIP is used only for packet interception, no other routers should be affected.

NOTE: If you change the number of seconds between RIP updates (the default is 30), you must specify the same value on the off-path device. To match this example, enter the following CLI command on the off-path device:

```
config packet-interception rip set update-timer 5
```

To view the RIP routes advertised by the off-path device, enter the following command:

```
show ip route rip
```

If packet interception is working correctly, you should see routes like the following. In this example, 10.1.14.50 is the off-path device, and the IP address of the remote WX device (10.1.203.50) has been carved out.

```
10.1.0.0/16 is variably subnetted, 24 subnets, 9 masks
```

```
R 10.1.203.128/25 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.51/32 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.48/31 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.52/30 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.56/29 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.32/28 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.0/27 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
R 10.1.203.64/26 [30/2] via 10.1.14.50, 00:00:23,
Ethernet0/1
```

To view debugging information for RIP events on a Cisco router:

```
debug ip rip events
```

Sample debugging information:

```
1w1d: RIP: received v2 update from 10.1.14.50 on Ethernet0/1
1w1d: RIP: Update contains 8 routes
```

You can also enter "debug ip rip database" or "debug ip rip trigger" for more details.

Dual Off-Path Devices on Two Layer 3 Switches

In Figure 4-98, two off-path devices are connected to dedicated VLANs on two Layer 3 switches. To use D1 as the preferred device, SW2 is configured to add an offset to the RIP routes advertised by D2. The two switches exchange RIP routes so that if D1 fails, the “higher cost” routes from D2 are used automatically by both switches. Also, D3 specifies D1 as the preferred assembler.

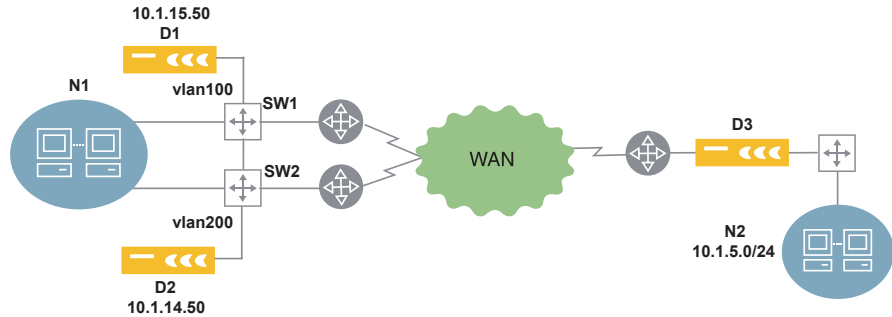


Figure 4-98 Dual Off-Path Devices on Two Layer 3 Switches

1. Enable RIP on SW1. Note that RIP is not passive because SW1 and SW2 exchange routes.

```
router rip
  version 2
  timers basic 5 15 15 30
  network 10.0.0.0
  distance 30
  no auto-summary
```

2. Enable RIP on SW2 so that a five-hop offset is added to the RIP routes received from D2 (which are the routes advertised by D3):

```
access-list 10 permit host any

router rip
  version 2
  timers basic 5 15 15 30
  offset-list 10 in 5 interface vlan200
  network 10.0.0.0
  distance 30
  no auto-summary
```

Thus, the routes from D2 have six hops on SW2, and seven hops on SW1, while the same routes from D1 have one hop on SW1 and two hops on SW2. The routes from D2 are used only if D1 fails.

If the D1 and D2 are on the same subnet, you can specify the offset on D2:

```
config packet-interception rip set metric 7
```

NOTE: If you change the number of seconds between RIP updates (the default is 30), you must specify the same value on the off-path devices. To match this example, enter the following CLI command on the off-path device:

```
config packet-interception rip set update-timer 5
```

WCCP Router Configuration Commands

The following commands provide an example of how to configure WCCP on a Cisco router for the deployment shown in Figure 4-99. The actual commands used will vary, depending on the network's topology and the type of traffic to be redirected. For more information about WCCP, go to the Cisco documentation page at <http://www.cisco.com/univercd/home/home.htm> and search for “wccp”.

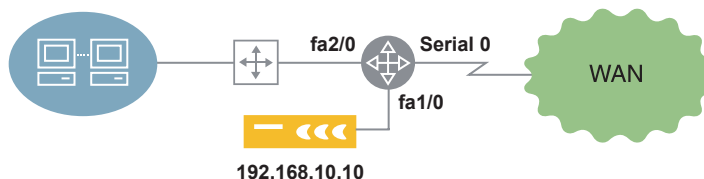


Figure 4-99 Off-Path Device Connected to a Router

1. Define an access list that specifies the traffic that is eligible for redirection to the off-path device:

```
access-list 120 permit ip any any
```

2. If the off-path device assigns WCCP service IDs 85 and 87 to TCP and UDP, respectively, create the two service IDs on the router. Include the password if authentication is enabled.

```
ip wccp 85 redirect-list 120 password <password>
ip wccp 87 redirect-list 120 password <password>
```

3. To redirect traffic from the outbound WAN interface, specify the WCCP service IDs to be redirected:

```
interface Serial 0
ip address 192.168.5.103 255.255.255.0
ip wccp 85 redirect out
ip wccp 87 redirect out
```

Alternatively, to redirect traffic from the inbound interface from the switch:

```
interface FastEthernet 2/0
ip address 192.168.5.103 255.255.255.0
ip wccp 85 redirect in
ip wccp 87 redirect in
```

NOTE: If you define a service ID on the router, but omit the redirect commands, no traffic is redirected to the off-path. However, entering the “show packet-interception” command on the off-path device will indicate the service is connected.

External Policy-Based Router Commands

The following commands provide examples of how to configure policy-based routing on Cisco routers and Layer 3 switches.

If the off-path device is connected to a dedicated port on a router, the policy is applied to the inbound interface from the LAN switch. In the following example, any incoming packet on interface FastEthernet 0/0 that matches access-list 120 is routed to the off-path device at IP address 192.168.10.10. The access list shown here redirects all packets, but it can be as specific as necessary.

```
interface FastEthernet 0/0
ip address 192.168.9.1 255.255.255.0
ip policy route-map juniper

access-list 120 permit ip any any

route-map juniper permit 50
match ip address 120
set ip next-hop 192.168.10.10
```

If the off-path device is connected to a dedicated VLAN on a Layer 3 switch, the commands are almost the same, except that the policy is applied to the switch on the inbound interface from the LAN:

```
interface Vlan200
ip address 192.168.9.1 255.255.255.0
ip policy route-map juniper
```

NOTE: Use the “set ip next-hop” command to redirect packets to the IP address of the off-path device. Do not use the “set interface” command to redirect traffic to the interface where the off-path device is connected.

Alternatives to Packet Interception

In some environments, you can install an off-path device by connecting the Local and Remote interfaces to different VLANs on the same switch. Packet interception is not used.

Layer 2 Switch Sandwich

In the high-availability environment in Figure 4-100, D1 and D2 are connected in “two-legged” VLANs on two Layer 2 switches. All traffic is switched through the devices as it passes to and from the WAN routers.

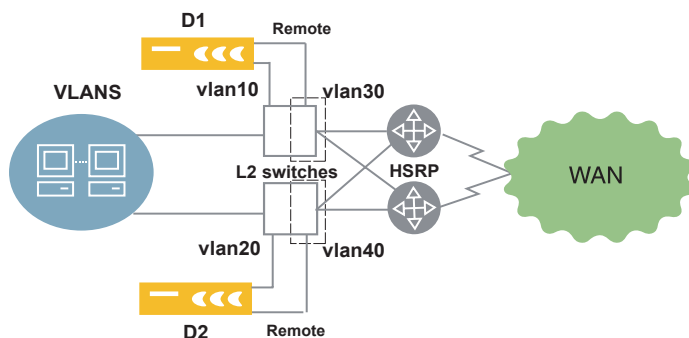


Figure 4-100 Layer 2 Switch Sandwich

Note the following:

- The Local interface is placed in the original VLAN that previously connected the switch port to the WAN router.
- The Remote interface is placed in a new VLAN along with the switch ports that feed the WAN routers.
- The default gateway of D1 and D2 is the HSRP address of the WAN routers. If one router fails, traffic is directed to the other router.
- Use a cross-over cable to connect the Local interface to the switch so that traffic is blocked if one device fails. The Layer 3 switches can then route the traffic through the other device.

Layer 3 Switch Sandwich

Figure 4-101 shows a single device connected across Layer 2 and Layer 3 VLANs on an L2/L3 switch. All traffic is switched through D1 as it passes to and from the WAN routers.

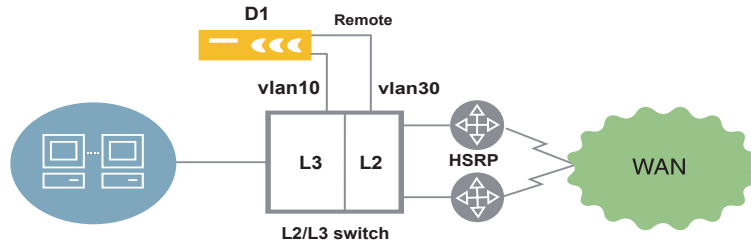


Figure 4-101 Layer 3 Switch Sandwich

Note the following:

- Hosts on the local LAN must point to the HSRP default gateway on same subnet.
- The Local interface is placed in the original VLAN that previously connected the switch port to the WAN router.
- The Remote interface is placed in a new Layer 2 VLAN along with the switch ports that feed the WAN routers.
- The default gateway of D1 is the HSRP address of the WAN routers. If one router fails, traffic is directed to the other router.

Configuring WAN Performance Monitoring

WAN performance monitoring lets each WX device measure the latency and loss to one or more remote WX devices. Probes are sent at an adjustable rate to each selected endpoint, and the loss and latency calculated for each WAN path is shown on the WAN Performance report. If the loss or latency exceeds the specified thresholds, an informational SNMP trap and Syslog entry are generated, and an event icon is shown on the report.

Data reduction is not required for WAN performance monitoring.

NOTE: If both Multi-Path and WAN performance monitoring are enabled for the same remote endpoint, the Multi-Path loss and latency settings take precedence. However, the WAN performance settings take effect if Multi-Path is disabled (refer to “Configuring Multi-Path Addresses” on page 101).

To enable WAN performance monitoring:

1. In the Configuration window, click **ADVANCED SETUP** in the left-hand navigation frame, click **WAN Performance Monitor**, and select the check box.

Main_config Global Configuration Compatible with SRS 5.1 Save Cancel

Global

- BASIC SETUP
- AAA
- APPLICATIONS
- REDUCTION
- QOS
- ACCELERATION
- ADVANCED SETUP
 - Features/Topology
 - Source/Destination Filter
 - Prime Time
 - Packet Interception
 - ☒ WAN Performance Monitor
 - CLI
- MULTI-PATH
- IPSEC

WAN Performance Monitoring HELP

☒ Enable WAN Performance Monitoring for listed endpoints

Device Name	IP Address	Latency Threshold (msec)
10.10.20.20	10.10.20.20	100
SR-192.168.0.195	192.168.0.195	100
SR-192.168.5.50	192.168.5.50	100
SM-192.168.16.60	192.168.16.60	100

Add/Remove Endpoints...

Submit Reset

Figure 4-102 Configuring WAN Performance Monitoring

2. Select the **Enable WAN Performance Monitoring for listed endpoints** check box.
3. To add or remove the devices to be monitored:
 - a. Click **Add/Remove Endpoints**.
 - b. Select a community or device group from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.
 - c. Select the devices you want to monitor, and click **Add**. To remove devices from the list, select the devices and click **Remove**.

- d. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - e. If one or more devices you want to add are not listed for the community/device group, you can add the devices manually. Click **Manual Entry**, enter the device IP addresses (one per line), and click **Submit**.
 - f. Click **Submit** to enter the changes.
4. Specify the following for each monitored endpoint:
- | | |
|-------------------|---|
| Latency Threshold | <p>Enter the round-trip time (RTT) threshold in milliseconds (20 to 5000). Traps, Syslog entries, and report events are generated when the threshold is exceeded, and again when latency drops below the threshold.</p> <p>By default, a probe tests the path 12 times per minute. Traps are generated when the median latency exceeds the threshold for four consecutive minutes or if two or more probes are lost per minute for four consecutive minutes. To change these settings, refer to the “configure wan-performance-monitor” CLI command).</p> <p>Note that availability on the WAN Performance report is measured as the percentage of minutes for which at least one probe was acknowledged.</p> |
|-------------------|---|
5. Click **Submit** to enter the changes, or click **Reset** to discard them.

Adding CLI Commands to Configurations

You can append CLI commands to a global configuration or an Advanced Setup partial configuration. This is intended primarily for use by support representatives to troubleshoot problems.

NOTE: For SRS 5.0 configurations, CLI commands can be defined only in the last section of a global configuration.

To append CLI commands to a global configuration:

1. In the Configuration window, click **ADVANCED SETUP** in the left-hand navigation frame, click **CLI**, and select the check box.

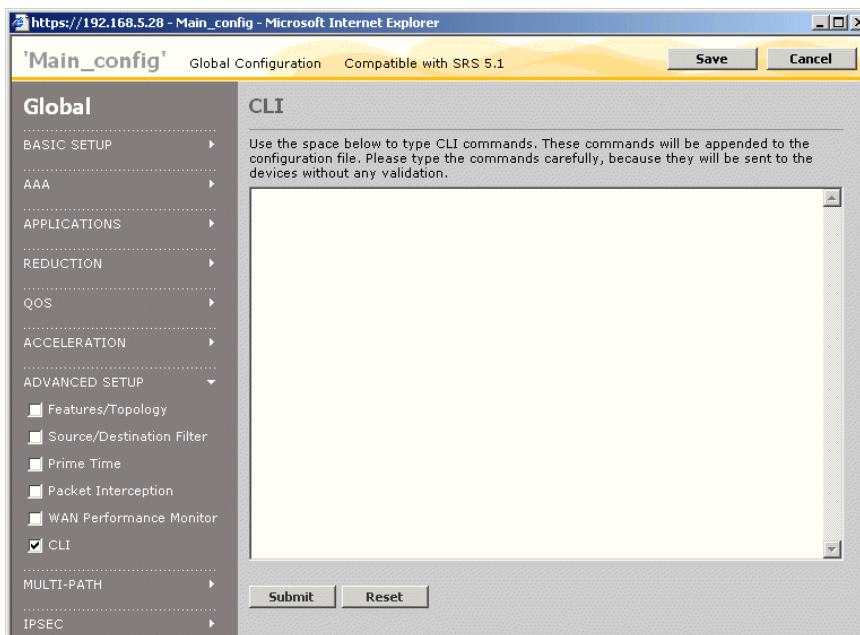


Figure 4-103 Appending CLI Commands to a Global Configuration

2. Enter CLI commands provided by the support representative.
3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Configuring Policy-Based Multi-Path

If a pair of WX devices has two possible WAN paths between them, you can designate one path as the primary and the other as the secondary. You can then route application traffic to the primary or secondary path based on the performance requirements of the application and the actual performance of the path.

To use Multi-Path, you configure both devices so that outgoing packets intended for the secondary path are marked with a secondary source IP address and, optionally, with a specific gateway address or ToS/DSCP value. In most cases, you must configure the WAN routers to route the marked packets to the appropriate path. The traffic for the preferred path (primary or secondary) is specified by traffic class, where each class contains one or more applications.

NOTE: The secondary IP address on each device must be specified using the device Web console or a Device Settings partial configuration.

For example, in Figure 4-104, most traffic might normally be sent over the private WAN, while email traffic is sent over the Internet. D1 and D2 mark email traffic with a secondary IP address, and R1 and R2 are configured to route the marked traffic to the Internet. If the private WAN fails, selected application traffic can be diverted automatically to the Internet; if the Internet latency exceeds a specified threshold, email traffic can be diverted to the private WAN. Traffic is switched back to the preferred path when conditions return to normal.

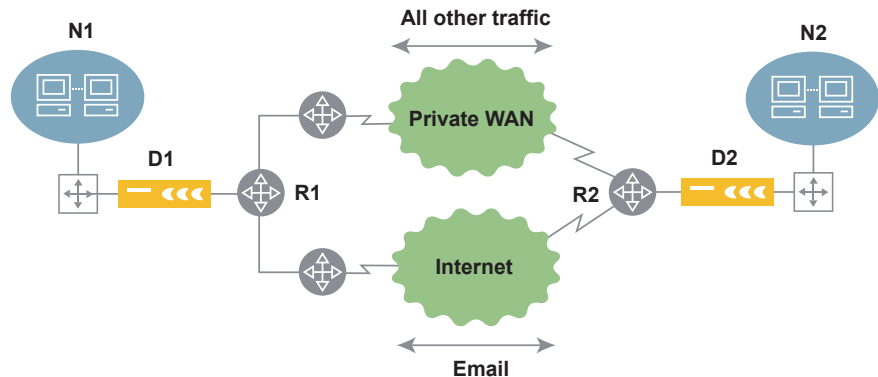


Figure 4-104 Multi-Path Deployment

The following topics describe how to configure Multi-Path:

- “Procedure for Configuring Multi-Path” in the next section
- “Enabling Policy-Based Multi-Path” on page 237
- “Defining Multi-Path Templates” on page 238
- “Defining Multi-Path Endpoints” on page 240
- “Configuring Routers to Support Multi-Path” on page 243

Procedure for Configuring Multi-Path

To configure Multi-Path for a pair of devices, do the following on BOTH devices:

1. Verify that data reduction is enabled between the two devices (refer to “Configuring Endpoints for Reduction Tunnels” on page 139).
2. Verify that the appropriate traffic classes are defined (refer to “Assigning Applications to Traffic Classes” on page 136). Outbound QoS can be on or off.
3. For each device, specify a secondary IP address and primary and secondary gateway addresses (if applicable) using the device Web console or a Device Settings partial configuration (refer to “Configuring Multi-Path Addresses” on page 101). The Device Settings configuration must be loaded on each device before you can configure the Multi-Path endpoints.
4. Enable Multi-Path and, optionally, specify primary and secondary ToS/DSCP values (refer to “Enabling Policy-Based Multi-Path” on page 237).
5. Define templates that specify the preferred path (primary or secondary) for each traffic class and the conditions when the traffic for each class can be switched to the alternate path (refer to “Defining Multi-Path Templates” on page 238).
6. Apply a template to each remote device that supports Multi-Path, and specify the congestion and latency thresholds for each path (refer to “Defining Multi-Path Endpoints” on page 240).
7. If necessary, configure the WAN router to route traffic to the appropriate path (refer to “Configuring Routers to Support Multi-Path” on page 243).

Enabling Policy-Based Multi-Path

To enable Multi-Path on a device from CMS, you must first specify a secondary IP address and primary and secondary gateway addresses (if applicable) using the device Web console or a Device Settings partial configuration (refer to “Configuring Multi-Path Addresses” on page 101). The Device Settings configuration must be loaded on each device before you can configure the Multi-Path endpoints.

To enable Multi-Path:

1. In the Configuration window, click **MULTI-PATH** in the left-hand navigation frame, click **Start/Stop**, and select the check box.

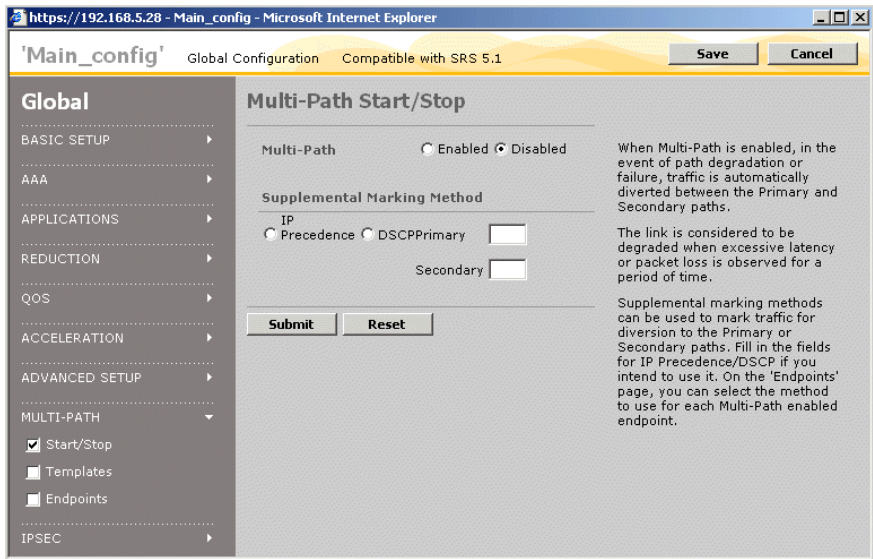


Figure 4-105 Multi-Path Start/Stop Page

2. Specify the following information:

Multi-Path	Select Enabled to activate the Multi-Path feature.
IP Precedence/DSCP	Optionally, you can mark packets sent on the primary and secondary paths with different ToS/DSCP values. Select IP Precedence or DSCP and enter a ToS IP precedence value (0 to 7) or DSCP value (0 to 63) for packets sent on the primary and/or secondary paths.

NOTE: These values override the IP precedence or DSCP settings for:

- Outbound QoS (refer to “Changing Outbound ToS/DSCP Values” on page 184)
- Control packets (refer to the “configure reduction” CLI command)

The multi-path DSCP values also override ToS marking for router-based balancing (refer to the “configure route” CLI command).

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Defining Multi-Path Templates

At least one Multi-Path template is required to specify the preferred path for each traffic class, and the conditions under which the traffic for each class can be switched to the alternate path. To assign a template to each remote device that supports Multi-Path, refer to “Defining Multi-Path Endpoints” on page 240.

To define Multi-Path templates:

1. In the Configuration window, click **MULTI-PATH** in the left-hand navigation frame, click **Templates**, and select the check box.

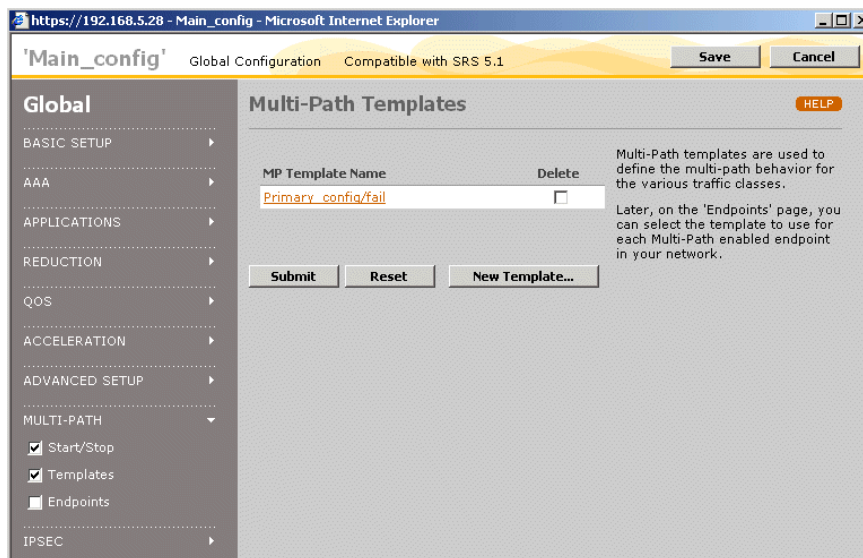


Figure 4-106 Defining Multi-Path Templates

2. To add a new template, click **New Template**, specify the following information, and click **Submit**:

Template Name	Enter the template name (up to 20 characters).
For each traffic class, select the following (to add new traffic classes, refer to “Assigning Applications to Traffic Classes” on page 136).	
Preferred Path	Select Primary or Secondary to indicate the path used for each traffic class under normal conditions.
Divert	<p>Select the conditions under which each traffic class can be switched to the alternate path:</p> <ul style="list-style-type: none">• Never. The traffic class is never diverted from the preferred path.• Failure Only. The traffic class is diverted to the alternate path only if the reduction tunnel for the preferred path goes down and the reduction tunnel for the alternate path is active.• Congestion/Failure. The traffic class is diverted to the alternate path if the loss or latency threshold is exceeded on the preferred path or the reduction tunnel goes down. A diversion for loss or latency occurs only if the alternate path’s loss and latency are not exceeded. <p>If Congestion/Failure is selected for any traffic class, probe packets are sent to the remote devices to measure the loss and latency of each path. To specify a latency threshold for each remote device, refer to “Defining Multi-Path Endpoints” on page 240. By default, the loss threshold is exceeded if two or more probes are lost per minute for four consecutive minutes.</p> <p>All of the threshold settings can be changed using the CLI (refer to the “configure multi-path” command).</p>

NOTE: Outbound QoS settings do not affect how traffic is diverted between alternate paths.

3. To change a template name or settings, click the template name, change the template name and/or the settings for each traffic class, and click **Submit**.
4. To delete a template, click the check box next to the template name, and click **Submit**. If a template is applied to an endpoint, it cannot be deleted.

Defining Multi-Path Endpoints

For each device that has a secondary IP address, you can select a multi-path template and supplemental marking method (if any), and specify a latency threshold for the primary and secondary paths to the device.

To specify a secondary IP address for a device, use the device Web console or load a Device Settings partial configuration on the device (refer to “Configuring Multi-Path Addresses” on page 101).

To define Multi-Path endpoints:

1. In the Configuration window, click **MULTI-PATH** in the left-hand navigation frame, click **Endpoints**, and select the check box.

The screenshot shows the 'Main_config' web interface in Microsoft Internet Explorer. The left navigation pane has 'MULTI-PATH' expanded, and 'Endpoints' is checked. The main content area is titled 'Multi-Path Endpoints' and includes a 'HELP' button. Below the title, there is explanatory text and a note about endpoints selected for Reduction. A table lists the configured endpoints:

Device Name	Latency Threshold (msec)		Multi-Path Template	Supplemental Marking Method
	Primary	Second.		
SR-192.168.0.195	5000	5000	Primary_config/fail	None (Sec. IP Only)

Below the table are buttons for 'Add/Remove Endpoints', 'Submit', and 'Reset'.

Figure 4-107 Defining Multi-Path Endpoints

2. To add or remove remote endpoints for Multi-Path:
 - a. Click **Add/Remove Endpoints**.
 - b. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have their primary and secondary addresses enclosed in parentheses, separated by a slash.

- c. Select the Multi-Path devices you want to configure, and click **Add**. To remove devices from the Multi-Path Endpoints list, select the devices and click **Remove**.
- d. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
- e. If one or more devices are not listed, click **Manual Entry** and enter the primary and secondary IP addresses for each device, separated by a slash (one address pair per line), and click **Submit**.
- f. When you are done, click **Submit**.

NOTE: Reduction is required for Multi-Path. When you save a global configuration, an error occurs if reduction is disabled for an endpoint using Multi-Path. If you add an endpoint to a Multi-Path partial configuration, an error occurs if you load the configuration on a device where reduction is disabled for that endpoint.

3. For each Multi-Path endpoint, specify the following:

Latency Threshold	<p>Enter the latency threshold in milliseconds (20 to 5000) for the primary and secondary paths. Traffic is switched to the alternate path when the threshold is exceeded, and is switched back when latency drops below the threshold. This setting is ignored for traffic classes where the selected template disallows switching between paths.</p> <p>NOTE: If you set the threshold too low, minor fluctuations in latency may cause constant switching between paths.</p> <p>By default, a probe tests the path 12 times per minute, and traffic is switched when the median latency exceeds the threshold for four consecutive minutes. Traffic is also switched if two or more probes are lost per minute for four consecutive minutes. To change these settings, refer to the “configure multi-path” command</p>
Multi-Path Template	<p>Select a template for this endpoint that specifies the preferred path and the conditions under which traffic can be switched to the alternate path. To add a new template, refer to “Defining Multi-Path Templates” on page 238.</p>

Latency Threshold	<p>Enter the latency threshold in milliseconds (20 to 5000) for the primary and secondary paths. Traffic is switched to the alternate path when the threshold is exceeded, and is switched back when latency drops below the threshold. This setting is ignored for traffic classes where the selected template disallows switching between paths.</p> <p>NOTE: If you set the threshold too low, minor fluctuations in latency may cause constant switching between paths.</p> <p>By default, a probe tests the path 12 times per minute, and traffic is switched when the median latency exceeds the threshold for four consecutive minutes. Traffic is also switched if two or more probes are lost per minute for four consecutive minutes. To change these settings, refer to the “configure multi-path” command</p>
Supplemental Marking Method	<p>Optionally, select one of the additional marking methods for the packets sent on each path (refer to “Configuring Multi-Path Addresses” on page 101 and “Enabling Policy-Based Multi-Path” on page 237).</p> <p>By default, all packets to be sent on the secondary path have the source address set to the secondary IP address.</p>

4. Click **Submit** to enter the changes, or click **Reset** to discard them.

To view the status of the primary and secondary paths from a specific device, access the device Web console and open the Multi-Path Endpoints page and the Multi-Path monitoring report.

Configuring Routers to Support Multi-Path

You can configure a WAN router to select a gateway for multi-path traffic based on the source IP address, or based on the source address and a ToS or DSCP value. The following configuration examples apply to router R1 in Figure 4-108. A similar configuration is needed for R2.

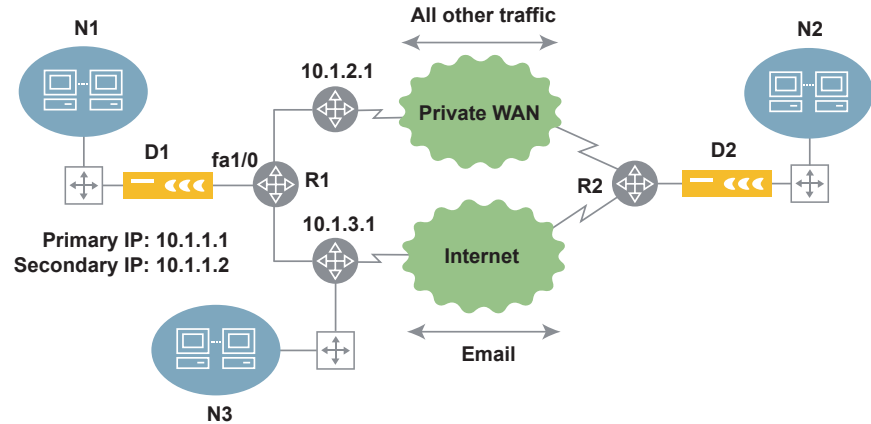


Figure 4-108 Multi-Path Router Configuration Example

To configure the WAN router R1 to use only the source IP address:

1. On the inbound interface from the WX device, define a route map for Multi-Path. For example:

```
interface FastEthernet 1/0
 ip address 10.1.1.5 255.255.255.0
 ip policy route-map mpath
```

2. Define access lists for the primary and secondary source IP addresses. For example:

```
access-list 50 permit 10.1.1.1
access-list 51 permit 10.1.1.2
```

3. Match the primary and secondary source IP addresses with the appropriate primary and secondary gateways. For example:

```
route-map mpath permit 10
 match ip address 50
 set ip next-hop 10.1.2.1
```

```
route-map mpath permit 20
  match ip address 51
  set ip next-hop 10.1.3.1
```

To configure R2, use the commands above, but change the interface address and use the primary and secondary address for device D2.

To configure the WAN router R1 to use both the source address and the ToS IP precedence or DSCP values:

1. Define a route map for Multi-Path (see the previous example).
2. Define extended access lists for the primary and secondary source IP addresses and their associated IP precedence or DSCP values. For example, for IP precedence values:

```
access-list 100 permit ip host 10.1.1.1 any precedence 10
access-list 101 permit ip host 10.1.1.2 any precedence 11
```

For DSCP values:

```
access-list 100 permit ip host 10.1.1.1 any dscp 1
access-list 101 permit ip host 10.1.1.2 any dscp 2
```

3. Match the primary and secondary source IP addresses with the appropriate primary and secondary gateways. For example:

```
route-map mpath permit 10
  match ip address 100
  set ip next-hop 10.1.2.1

route-map mpath permit 20
  match ip address 101
  set ip next-hop 10.1.3.1
```

NOTE: Unless you use a console server to manage devices, you may need to change the access lists to allow management access from some locations using SSH or Web/SSL. For example, in Figure 4-108, you may not be able to access P1 from N3 because management responses have the primary IP address, and are routed to the private WAN.

Configuring IPSec

IPSec can be used to authenticate and encrypt traffic between any pair of WX devices in the same community. Enabling IPSec allows you to:

- Compress traffic before it is encrypted (encrypted traffic cannot be compressed).
- Encrypt traffic over unprotected networks, such as the Internet.

To configure IPSec, you define templates that specify the security algorithms and key lifetimes for outgoing traffic, and then apply a template to each of the remote endpoints that act as IPSec peers. For a pair of devices to use IPSec, IPSec must be enabled on both devices, and both devices must be configured with the same pass phrase (preshared key) and security algorithms. Each device can encrypt traffic for up to 100 remote peers (the SR-20 is limited to five devices).

The following sections describe how to configure IP security (IPSec) to authenticate and encrypt traffic between any pair of devices:

- “Default IPSec Policy” in the next section
- “IPSec Implementation Details” on page 246
- “Procedure for Configuring IPSec Policies” on page 247
- “Defining IPSec Settings by Endpoint” on page 247
- “Defining IPSec Templates” on page 249
- “Defining the Default IPSec Policy” on page 251

Default IPSec Policy

When two devices are configured as IPSec peers, all compressed and passthrough traffic sent between them is encrypted. For passthrough traffic destined for subnets that are not served by a WX device, a “default IPSec policy” is provided that lets you specify, by subnet, whether the traffic is dropped and logged or sent unencrypted. Initially, the default IPSec policy allows all traffic to be sent unencrypted.

The default IPSec policy also applies to traffic between peer devices where IPSec is enabled, but the key negotiation has failed. Note that an IPSec-enabled device never encrypts traffic destined for a remote device where IPSec is disabled.

After you verify that IPSec is working correctly, all subnets advertised by IPSec-enabled peers should be added to the encryption-required list to avoid sending unencrypted traffic to those subnets if a remote device fails.

NOTE: If an inline device fails, all traffic is passed through without encryption. To block all traffic during a hardware failure, use a cross-over cable (rather than a straight-through cable) to connect the device to the WAN router. This works only if Ethernet auto-MDI negotiation is disabled on the router.

IPSec Implementation Details

IPSec is implemented in compliance with RFCs 2401-2409, and includes the following:

- Encryption algorithms—Advanced Encryption Standard (AES) encryption algorithm, with 128, 192, and 256 bit keys, and Triple DES (3DES)
- Authentication algorithms—HMAC/SHA-1 and HMAC/MD5
- Internet Key Exchange (IKE) protocol for dynamic key exchange
- Encapsulated Security Protocol (ESP) in transport mode used for all encrypted packets

AES with a 256 bit key and HMAC/SHA-1 authentication provides the highest security, while AES with a 128 bit key and HMAC/MD5 authentication provides the highest throughput (primarily because SHA-1 is two to three times slower than MD5). 3DES is supported for environments where AES is not approved, but 3DES is slower and less secure than AES, and is not recommended.

Although the IPSec protocols allow two peers to communicate using different policies, such as having Peer 1 use AES to encrypt for Peer 2, while Peer 2 uses DES to encrypt for Peer 1, both WX devices must use the same encryption and authentication algorithms.

Using IPSec allows you to compress traffic before encrypting it (encrypted traffic cannot be compressed because it contains few recognizable patterns). Since outgoing traffic is both compressed and encrypted, 3rd party IPSec devices cannot support our implementation because they cannot decompress the traffic. However, uncompressed IPSec traffic has been validated against Cisco and Microsoft IPSec implementations to ensure IPSec compliance.

NOTE: The IPSec Authentication Header (AH) is not used, and only Diffie-Hellman Group 5 is supported.

Procedure for Configuring IPSec Policies

To encrypt the traffic sent between two or more devices:

1. Select the devices that you want to encrypt traffic and specify the pass phrase(s) needed to establish a secure connection (refer to “Defining IPSec Settings by Endpoint” on page 247).
2. To change the default Wizard template or define new templates, refer to “Defining IPSec Templates” on page 249.
3. To change the default IPSec policy, refer to “Defining the Default IPSec Policy” on page 251.

Alternatively, you can run the IPSec Setup Wizard on each device from the device Web console.

Defining IPSec Settings by Endpoint

On the IPSec Overview page, you can enable or disable IPSec for all endpoints or specific endpoints, change the IPSec template or pass phrase for an endpoint, or enable encryption for management traffic. To add or change IPSec templates, refer to “Defining IPSec Templates” on page 249.

To view or change the IPSec settings by endpoint:

1. In the Configuration window, click **IPSEC** in the left-hand navigation frame, click **Overview**, and select the check box.

Global Configuration Compatible with SRS 5.1

Global

- BASIC SETUP
- AAA
- APPLICATIONS
- REDUCTION
- QOS
- ACCELERATION
- ADVANCED SETUP
- MULTI-PATH
- IPSEC**
 - ☒ Overview
 - ☒ Templates
 - ☐ Default Policy

IPSec Overview

☒ Enable IPSec Encryption for the endpoints selected below

Enter Pass Phrase Verify Pass Phrase

☒ Use a common pass phrase

☐ Use individual pass phrases for each endpoint

Endpoint	Name	Template	Mgmt. Traffic*	Enter Pass Phrase	Verify Pass Phrase
192.168.0.195	SR-192.168.0.195	Wizard	<input type="checkbox"/>		

Add/Remove Endpoints

* When checked, management traffic (SSH/SSL) is included in the encryption tunnel.

Submit Refresh Reset

Figure 4-109 IPSec Overview

2. To enable IPSec, click the check box next to **Enable IPSec Encryption for the endpoints selected below**.
3. To add or remove remote endpoints for IPSec:
 - a. Click **Add/Remove Endpoints**.
 - b. Select a community from the **Community/Device Group** list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

 Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable IPSec for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 101.
 - c. Select the devices you want to configure, and click **Add**. To remove devices from the IPSec Endpoints list, select the devices and click **Remove**. If you remove an endpoint, all subsequent traffic to that endpoint is sent unencrypted.
 - d. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - e. If one or more devices are not listed, click **Manual Entry** and enter the primary or secondary IP addresses for each device (one per line), and click **Submit**.
 - f. When you are done, click **Submit**.
4. Enter and verify a pass phrase for each endpoint, or select **Use a common pass phrase** and enter one pass phrase for all endpoints (eight or more characters is recommended). The pass phrase is used to generate a preshared key of the appropriate length.
5. To change the template for an endpoint, select a template from the **Template** drop-down list. To create new templates, refer to “Defining IPSec Templates” on page 249. Two endpoints can establish a secure connection only if their IPSec templates specify the same authentication and encryption algorithms. The default Wizard template uses AES-128 and HMAC/SHA-1.

- To encrypt all management traffic sent to a remote endpoint, including SNMP, Syslog, and registration server traffic, click the **Mgmt. Traffic** check box for the endpoint. Encrypting management traffic is recommended after you verify that the IPsec connection is operating normally.

To view the status of the IPsec connections from a specific device, access the device Web console and open the IPsec Overview page.

- Click **Submit** to enter the changes, or click **Reset** to discard them.

Defining IPsec Templates

IPsec templates specify the algorithms used to protect traffic between endpoints, and the lifetime of each generated key. You can change the default Wizard template or create new templates. The default Wizard template uses the AES-128 and HMAC/SHA-1 encryption and authentication algorithms, and the generated keys do not expire.

To apply a template to an endpoint, refer to “Defining IPsec Settings by Endpoint” on page 247.

To define IPsec templates:

- In the Configuration window, click **IPSEC** in the left-hand navigation frame, click **Templates**, and select the check box.

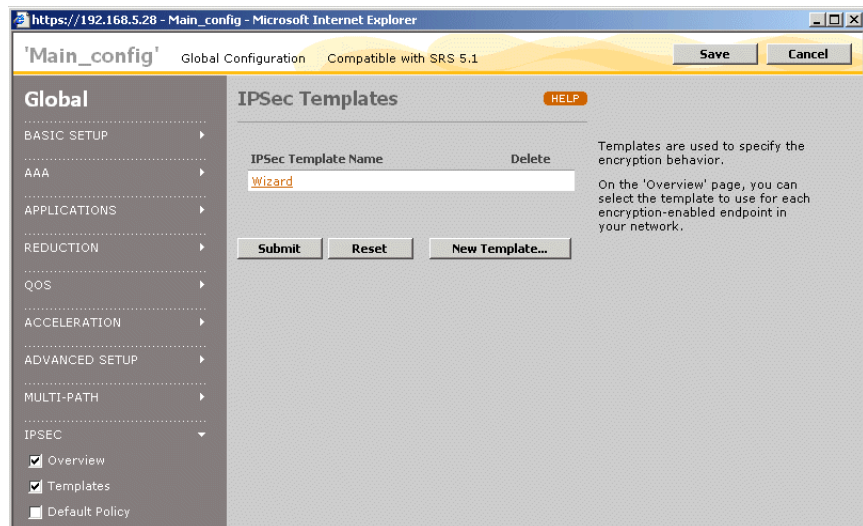


Figure 4-110 Defining IPsec Templates

2. To add a new template, click **New Template**, specify the following information, and click **Submit**:

Template Name	Enter the name of the template (up to 20 characters).
Encryption Algorithm	<p>Select the algorithm used to encrypt outbound traffic:</p> <ul style="list-style-type: none"> • Any. The algorithm selected for the peer endpoint is used. If both endpoints specify Any, AES-128 is used. • AES-128. Advanced Encryption Standard with a 128-bit key. • AES-192. AES with a 192-bit key. • AES-256. AES with a 256-bit key. • 3DES. Triple Digital Encryption Standard with a 168-bit key.
Authentication Algorithm	<p>Select the algorithm used to authenticate outbound traffic:</p> <ul style="list-style-type: none"> • Any. The algorithm selected for the peer endpoint is used. If both endpoints specify Any, HMAC/SHA-1 is used. • HMAC/SHA-1. Secure Hash Algorithm. • HMAC/MD5. Message Digest 5.
Key Lifetime	<p>Specify the time and data limits for generated keys:</p> <ul style="list-style-type: none"> • Time. Enter the number of hours before a generated key expires (up to 2160), or select Never expires. • Data. Enter the number of megabytes of traffic allowed before a generated key expires (up to 4000), or select Never expires. <p>Key negotiation begins when the key lifetime reaches 80% of the time limit or 50% of the data limit. Keys should be negotiated periodically for security purposes.</p>

3. To change a template name or settings, click the template name, change the template, and click **Submit**.
4. To delete a template, click the check box next to the template name, and click **Submit**. If you load the configuration on a device where the deleted template is applied to an endpoint, the endpoint reverts to the Wizard template.

Defining the Default IPSec Policy

The default IPSec policy is applied to the following types of traffic:

- Passthrough traffic sent to unadvertised subnets
- Traffic between devices where IPSec is enabled, but the key negotiation has failed

By default, all such traffic is unencrypted. However, you can change the default policy so that traffic to specific destinations is dropped and logged, rather than sent unencrypted. The number of packets dropped for each destination is written to the system log every five minutes. Use the device Web console to view the system log.

After you verify that IPSec is working correctly, all subnets advertised by IPSec-enabled peers should be added to the encryption-required list to avoid sending unencrypted traffic to those subnets if a remote device fails.

NOTE: All passthrough traffic between IPSec-enabled devices is encrypted. For example, traffic is encrypted even when reduction is disabled.

To change the default IPSec policy:

1. In the Configuration window, click **IPSEC** in the left-hand navigation frame, click **Default Policy**, and select the check box.



Figure 4-111 Defining the IPSec Default Policy

2. In the two text boxes, specify the destination addresses and subnets where encryption is required or optional, as follows:

Encryption Required	Enter destination addresses or subnets (one per line) for which traffic must be dropped and logged. The subnet format is: <IP address>/<subnet mask>
Encryption Optional	Enter destination addresses or subnets (one per line) for which traffic can be sent unencrypted. For example, if subnet 10.10.0.0/255.255.0.0 is specified as encryption required, you can specify one or more smaller subnets in that range where encryption is optional, such as 10.10.20.0/255.255.255.0. If an address or subnet is in both lists, or in neither list, the traffic is not encrypted.

3. Click **Submit** to enter the changes, or click **Reset** to discard them.

Chapter 5 Automatic Deployment of Devices

This chapter describes how to use CMS to automatically download configurations and SRS software to new WX devices, and to generate permanent licenses for devices that need them. It covers the following topics:

- “About Automatic Deployment” in the next section
- “Configuring Auto-Deployment” on page 254
- “Configuring License Management” on page 261

About Automatic Deployment

When a WX device running SRS 5.x (or later) is powered on for the first time, it attempts to contact the CMS server. If you know the subnet where the device is installed, you can configure CMS to download a configuration and an SRS software image to the new device. On-site personnel simply connect the cables and apply power, and the device becomes operational.

After a successful auto-deployment, you can generate a permanent license for the device (refer to “Configuring License Management” on page 261).

Auto-deployment has two requirements:

- A DHCP server must be reachable from the WX device. When first powered on, the device sends DHCP requests over its Local and Remote interfaces. The DHCP server must reply with an IP address and the address of one or more DNS servers. Up to three DNS servers will be queried.
- One of the three DNS servers must have an entry for “peribit-cms” in the domain hierarchy. For example, if the domain name in the DHCP reply is “sales.company.com”, the WX device issues DNS requests in the following order to locate the CMS server:
 - peribit-cms.sales.company.com
 - peribit-cms.company.com
 - peribit-cms.com
 - peribit-cms

If DHCP does not specify a domain, and a reverse lookup on the DNS server’s address does not obtain one, then “peribit-cms” is the only request.

After obtaining the IP address of the CMS server, the device contacts the server over HTTPS. CMS can then download the prepared configuration and software image based on the device's subnet.

NOTE: Only one device can be auto-deployed per subnet. For example, a multi-node configuration, where two devices are connected together, cannot be auto-deployed. Also, an SR-100 can be auto-deployed, but its client devices must be configured locally.

Configuring Auto-Deployment

The following topics describe how to configure auto-deployment:

- “Auto-Deployment Procedure” in the next section
- “Defining Deployment Groups” on page 255
- “Defining Deployment Records” on page 257
- “Viewing the Auto-Deployment Status” on page 259

Auto-Deployment Procedure

Use the following procedure to configure auto-deployment:

1. Prepare full configurations to be downloaded to the auto-deployed devices. You can load a global configuration, a full set of partial configurations, or a combination of both (refer to “Defining Configuration Settings” on page 83).

For example, in a hub and spoke environment, you might create a global configuration for the spokes, and partial configurations that override the topology and other settings for the hubs (refer to “Configuring the Feature/Topology Settings” on page 212).

2. Define deployment groups that specify the configuration and software image to be downloaded (refer to “Defining Deployment Groups” on page 255).
3. Define a deployment record for each auto-deployed device that specifies the device's subnet, deployment group, and other settings (refer to “Defining Deployment Records” on page 257).
4. Monitor the status of the auto-deployed devices (refer to “Viewing the Auto-Deployment Status” on page 259).

- When the deployment is complete, verify that the communities for the auto-deployed devices have been imported from the registration server(s) specified in the device configurations. Initially, all devices are in the Default community (refer to “Managing Communities” on page 328).
- Configure licenses for the auto-deployed devices (refer to “Configuring License Management” on page 261).

Defining Deployment Groups

After you prepare configurations for the WX devices to be auto-deployed, you must define at least one deployment group. A deployment group specifies the global and/or partial configurations that you want to download to one or more auto-deployed devices. Optionally, you can also specify an SRS software image to be loaded at the same time.

To define deployment groups:

- Click **MANAGEMENT** in the menu frame, and then click **Auto-Deployment** in the left-hand navigation frame.



Figure 5-1 Defining Deployment Groups

The Deployment Groups page lists the device compatibility (SRS 5.1 or 5.0), configuration(s), and software image specified by each deployment group.

- To change a deployment group, click the name, make any needed changes and click **Submit**.
- To define a new deployment group, from the Task menu select **New (5.1 compatible)** for SRS 5.1 devices, or **New (5.0 compatible)** for SRS 5.0 devices, and click **Go**.

The screenshot shows the Periscope 1S Central Management System interface. The top navigation bar includes 'MY PERIBIT', 'MONITOR', 'MANAGEMENT', 'SETUP', and 'ABOUT'. The user is logged in as 'root'. The left sidebar shows the 'Management' section with a 'View' sub-section containing 'Deployment Groups', 'Setup', and 'Status'. The main content area is titled 'Auto-Deployment > Deployment Groups > New'. It contains the following fields:

- Deployment Group Name:** A text input field.
- Global Configuration:** Two radio buttons: 'Do not load global configuration' and 'Main_config' (selected). A 'History' link is next to 'Main_config'.
- Partial Configurations:** A list of configuration types with dropdown menus and 'History' links:
 - AAA: -- History
 - Acceleration: -- History
 - Advanced Setup: -- History
 - Applications: -- History
 - Basic Setup: -- History
 - IPSEC: -- History
 - Multi-Path: -- History
 - QoS: -- History
 - Reduction: -- History
- Software Image:** Two radio buttons: 'Do not load image' and 'srs.zip' (selected).

At the bottom are three buttons: 'Submit', 'Preview...', and 'Cancel'.

Figure 5-2 Adding a Deployment Group

4. Specify the following information:

- | | |
|-----------------------|--|
| Deployment Group Name | Enter a name for the deployment group. |
| Global Configuration | Select a global configuration or, to load only partial configurations, select Do not load global configuration . Click History to view the selected configuration and its past changes. To create global configurations, refer to “Managing Configurations” on page 75. |
| Partials | <p>If you are not loading a global configuration, you must select one of each type of partial configuration. The settings in each partial configuration replace the corresponding settings in the selected global configuration (if any) or are combined into one configuration.</p> <p>Click History to view each selected configuration and its past changes.</p> |
| Software Image | Select an SRS software image to be loaded, or select Do not load image . The image must first be loaded on the CMS server (refer to “Uploading a Boot Image” on page 327). |

5. Click **Preview** to view the resulting configuration. Any settings that are not defined in the global and partial configurations will remain in the factory default state on the device.

All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *Sequence Reducer/Sequence Mirror Operator's Guide*.

6. Click **Submit** to enter the changes, or click **Cancel** to discard them.

Defining Deployment Records

After you define the appropriate deployment groups, you must create a deployment record for each device to be auto-deployed. Each deployment record specifies the subnet where the device is installed, various network settings for the device, and a deployment group.

To define deployment records:

1. Click **MANAGEMENT** in the menu frame, click **Auto-Deployment** in the left-hand navigation frame, and then click **Setup**.

The screenshot shows the Juniper Periscope MS web interface. The top navigation bar includes 'MY PERISBIT', 'MONITOR', 'MANAGEMENT', 'SETUP', and 'ABOUT'. The left-hand navigation frame shows 'Management' with sub-items: 'Devices', 'Configurations', 'Auto-Deployment', 'License Management', 'Schedules', and 'Schedule Log'. The 'Auto-Deployment' section is expanded, showing 'View' with radio buttons for 'Deployment Groups', 'Setup', and 'Status'. The 'Setup' radio button is selected. The main content area is titled 'Auto-Deployment > Setup' and contains a table of deployment records. The table has the following columns: 'Originating Subnet', 'Static IP Addr.', 'Subnet Mask', 'Gateway', 'Device Settings Partial Configuration', and 'Deployment Group'. A single record is displayed with the following values: 'Originating Subnet: 1.1.1.0/24', 'Static IP Addr.: 1.1.1.5', 'Subnet Mask: /24', 'Gateway: 1.1.1.1', 'Device Settings Partial Configuration: AD1_1_1_5', and 'Deployment Group: Deployment_group1'. Below the table, there is a 'Time Zone' dropdown menu set to '(GMT -08:00) Pacific Time (US and Canada), Tijuana' and a 'Daylight Saving' checkbox. At the bottom, there are 'Add', 'Update', and 'Delete' buttons, and a 'Submit' button. A note at the bottom right says 'When you are done, click "Submit" to save any changes.'

Originating Subnet	Static IP Addr.	Subnet Mask	Gateway	Device Settings Partial Configuration	Deployment Group
<input type="checkbox"/> 1.1.1.0/24	1.1.1.5	/24	1.1.1.1	AD1_1_1_5	Deployment_group1

Time Zone: (GMT -08:00) Pacific Time (US and Canada), Tijuana Daylight: No Ready to deploy: ☐

View
☐ Deployment Groups
☒ Setup
☐ Status

Originating Subnet Static IP Addr. Subnet Mask Gateway Device Settings Partial Configuration Deployment Group

1.1.1.0/24 1.1.1.5 /24 1.1.1.1 --Create one-- Deployment_group1

Time Zone: (GMT -08:00) Pacific Time (US and Canada), Tijuana Daylight Saving: ☐

Add Update Delete When you are done, click "Submit" to save any changes.

Submit

Figure 5-3 Defining Auto-Deployment Records

The Deployment Setup page lists the properties of each deployment record.

2. To add a new deployment record, specify the following information:

Originating subnet	<p>Enter the subnet where a new device is (or will be) installed. The format is:</p> <p><i>subnet/mask</i></p> <p>where <i>mask</i> is the number of binary digits used for the network portion of the address.</p>
Static IP Addr.	Enter a static IP address for the new device. It need not be in the originating subnet.
Subnet Mask	<p>Enter the number of binary digits used for the network portion of the static address. The format is:</p> <p><i>/mask</i></p>
Gateway	Enter the IP address of the default gateway for the device. It must be in the same subnet as the static IP address.
Device Settings Partial Configuration	<p>If you created a Device Settings partial configuration for the device, you can select it here. The default setting (--Create one--) generates a Device Settings partial configuration named:</p> <p>AD<IP address></p> <p>where the dots in the static IP address are replaced by underscores. This partial configuration specifies only the settings in the deployment record (static address, subnet mask, gateway address, time zone, and daylight savings indicator).</p>
Deployment Group	Select a deployment group that specifies the configuration(s) to be loaded on the device. To add deployment groups, refer to "Defining Deployment Groups" on page 255.
Time Zone	Select the time zone of the device.
Daylight Saving	Select the check box to enable Daylight Savings Time on the device (if applicable).

3. When you are done, click **Add**, and click **Submit**.

To add a new record by copying and modifying an existing record, select the check box next to the subnet for the record you want to copy, and then clear the check box. You can then change the copied record, click **Add**, and click **Submit**.

4. To change a deployment record, click the check box next to the subnet, make any needed changes, click **Update**, and then click **Submit**.

- After you click **Submit**, you can leave the page and return later to add or edit deployment records. You can complete the deployment records over time as you establish the required network information for each device to be auto-deployed.
- When a deployment record is complete, click the **Ready to Deploy** check box, and click **Submit**. The configuration and software image (if any) are now ready to be downloaded when the device checks in. To view the status of the deployment, refer to “Viewing the Auto-Deployment Status” on page 259.

Note that the check box next to the subnet is greyed out. To make any additional changes to the record, you must first clear the **Ready to Deploy** check box, and click **Submit**.

Viewing the Auto-Deployment Status

After a deployment record is defined and marked “Ready to Deploy,” you can monitor the status of the auto-deployment to see when the device checks in and whether the deployment is successful.

To view the auto-deployment status:

- Click **MANAGEMENT** in the menu frame, click **Auto-Deployment** in the left-hand navigation frame, and then click **Status**.

The screenshot shows the PeriScope CMS Setup interface. The top navigation bar includes links for MY PERIBIT, MONITOR, MANAGEMENT, CMS SETUP, and ABOUT. The left-hand navigation frame shows the Management menu with options for Devices, Configurations, Auto-Deployment, License Management, Schedules, and Schedule log. The main content area displays the Auto-Deployment > Status page. A table lists deployment records with columns for IP Address, Originating Subnet, Deployment Attempts, Last Attempt, and Status. The first two records are marked 'Successful' and the third is 'In progress'. A 'Remove' button is present below the table, and a note indicates that IP addresses are unknown until the first deployment attempt, with a footnote for DHCP-assigned IP addresses.

IP Address	Originating Subnet	Deployment Attempts	Last Attempt	Status
192.168.52.200	192.168.52.192/28	1	2004-08-20 19:18:38.0	Successful
192.168.53.5	192.168.53.0/24	1	2004-08-19 18:42:20.0	Successful
192.168.53.140	192.168.53.128/26	1	2004-08-24 16:55:37.0	In progress
Unknown	10.10.2.0/24	0		

Remove Click this button to remove successfully deployed devices from the table

IP addresses are unknown until the first deployment attempt.
* DHCP-assigned IP address

Figure 5-4 Viewing Auto-Deployment Status

- The following information is provided for each deployment record marked “Ready to Deploy”:

IP Address	<p>The IP address shown for the device depends on the status of the deployment:</p> <ul style="list-style-type: none"> • Unknown. Device has not checked in. • <DHCP address>. Deployment is in progress. • <Static address>. Deployment successful.
Originating Subnet	The subnet where the device is installed (specified by the deployment record).
Deployment Attempts	<p>Number of times the deployment has been attempted. After five failed attempts, subsequent requests from the device are rejected. To allow another five attempts, you must the reset the “Ready to Deploy” flag on the deployment record (refer to “Defining Deployment Records” on page 257).</p>
Last Attempt	Date and time of the last deployment attempt.
Status	<p>Indicates the status of the auto-deployment:</p> <ul style="list-style-type: none"> • Blank. Device has not checked in. • In Progress. Deployment is in progress. • Successful. The configuration and software image (if any) were successfully downloaded to the device. • Failed. The last deployment attempt has failed. <p>Click the status for more details, such as the device type and MAC addresses.</p>

- Click **Remove** to remove the status entries for successful deployments (the corresponding deployment records are deleted automatically).

NOTE: You can auto-deploy a device only once. If an auto-deployed device is reset to the factory defaults, its attempts to contact the CMS server will be rejected.

Configuring License Management

The following topics describe how to configure the bulk deployment of new SRS licenses. For an upgrade license, contact your sales representative.

- “Licensing Procedure” in the next section
- “Importing and Validating Authorization Codes” on page 262
- “Generating and Applying Licenses” on page 264
- “Viewing the License Status” on page 266

Licensing Procedure

Use the following procedure to apply permanent SRS licenses to devices that have evaluation licenses:

NOTE: You must have an account on the Juniper Customer Support Center (<http://www.juniper.net/customers/support>), and the CMS server must be able to establish an HTTP connection with the license server at https://www.juniper.net/generate_license.

1. Obtain a file of Authorization Codes, and save the file in a location accessible from the browser.
2. Import and validate the Authorization Codes (refer to “Importing and Validating Authorization Codes” on page 262).
3. Match the Authorization Codes with the devices that have evaluation licenses, generate licenses for the matching devices, and then apply the licenses (refer to “Generating and Applying Licenses” on page 264).
4. Monitor the status of deployed licenses (refer to “Viewing the License Status” on page 266).

Importing and Validating Authorization Codes

After you obtain a file of Authorization Codes, you must import and validate the Authorization Codes in CMS. You can import Authorization Codes as often as needed.

To import and validate Authorization Codes:

1. When you purchase a “Right To Use” license from Juniper Networks, you receive an RTU Certificate in PDF format that lists the Authorization Codes at the end of the document. Use the Acrobat Text tool to copy and paste the Authorization Codes into a “.txt” file (one Authorization Code per line). Store the Authorization Codes file in a location accessible from the browser.
2. Click **MANAGEMENT** in the menu frame, and then click **License Management** in the left-hand navigation frame, and then click **Authorization Codes**.
3. Click **Import**, enter the Authorization Codes file location or click **Browse** to locate the file, and then click **Import**. Authorization Codes that have already been imported are marked as duplicates and excluded automatically. If format errors are displayed, contact the Juniper Technical Assistance Center.

Click **Back** to import another file, or click **Authorization Codes** in the navigation frame to view the Authorization Codes that were added to the database.

The screenshot shows the PeriScope CMS interface. The top navigation bar includes 'MY PERISBIT', 'MONITOR', 'MANAGEMENT', 'SETUP', and 'ABOUT'. The left sidebar shows 'Management' with sub-items: 'Devices', 'Configurations', 'Auto-Deployment', 'License Management' (selected), 'Schedules', and 'Schedule Log'. The main content area is titled 'License Management > Authorization Codes'. It contains two tables: 'Speed Authorization Codes' and 'IPSec Authorization Codes'. Both tables have columns for 'Model', 'Description', 'Authorization Code', and 'Status'. The 'Speed Authorization Codes' table lists codes for models WX 100, WX 15, WX 20, WX 5X, and WX 60. The 'IPSec Authorization Codes' table lists codes for models WX 100, WX 20, WX 5x, WX 60, and WX 60. At the bottom, there are buttons for 'Import...', 'Validate', and 'Delete Checked'.

Figure 5-5 Importing and Validating Authorization Codes

The following information is shown for each speed and IPSec Authorization Code:

Model	Device type, such as SR-80. An “SR-5x” indicates the Authorization Code can be applied to an SR-50 or SR-55.
Description	Indicates the base and maximum device speed (speed Authorization Codes only).
Authorization Code	Text identifying the Authorization Code (internal use only).
Status	Indicates the Authorization Code status: <ul style="list-style-type: none"> • New. Initial status of all imported Authorization Codes that have not been validated. • Valid. Validated by the License Server, but not yet assigned to a device. • Invalid. Not recognized by the License Server (contact Technical Support). • Canceled. No longer valid (contact Technical Support). • Temp-Assigned. Matched with a device, but not yet used to generate a license (refer to “Generating and Applying Licenses” on page 264). • Perm-Assigned. License generation has started or is complete, so the Authorization Code cannot be assigned to another device. • Used. Has been used to generate a license. This status may be set by the License Server when you validate the Authorization Codes. If a New Authorization Code is set to Used, and you have not used it to generate a license, contact Technical Support.

4. Select the check box next to the speed and IPSec Authorization Codes that you want to validate, or click **Select All**, and then click **Validate**.

The status for all **New** Authorization Codes should be changed to **Valid**. If any new Authorization Codes are set to invalid or canceled, contact Technical Support.

5. To delete the **Used** Authorization Codes, select the check box next to the appropriate codes, and click **Delete**.

You can now use the valid Authorization Codes to generate and apply licenses to your devices, as described in the next section.

Generating and Applying Licenses

After you import and validate your Authorization Codes, you can match them with the deployed devices that have evaluation licenses, generate permanent licenses, and then apply the licenses to each device.

NOTE: You must have an account on the Juniper Customer Support Center to generate the licenses.

To generate and apply licenses to your devices:

1. Click **MANAGEMENT** in the menu frame, click **License Management** in the left-hand navigation frame, and then click **License Generation**.
2. To view the devices in a community or device group that have an evaluation license (active or expired), select a community or device group and click **Submit**. It may take a few minutes to poll a large community. The page displays the progress of the poll (polling continues if you leave the page).

To stop the polling and select another community/device group, click **Stop**.

PeriScope CMS
CENTRAL MANAGEMENT SYSTEM

Juniper NETWORKS

MY PERIBIT MONITOR MANAGEMENT SETUP ABOUT

Logged in as: root LOGOUT

Management

Devices
Configurations
Auto-Deployment
License Management
Schedules
Schedule Log

View
● License Status
● License Generation
● Authorization Codes

License Management > License Generation Community: Inline [Change](#)

PeriScope CMS automatically matches deployed Peribit devices with available authorization codes. To exclude a device from the matching process, uncheck it. Then, click "Match Authorization Codes" to view the updated match results. If you are satisfied with the results, click "Generate Licenses" to create licenses based on the matched authorization codes. Click "Apply Licenses" to apply the generated licenses to the corresponding Peribit devices.

IP Address	Serial No.	Model	Speed	IPsec	Permanent License Status		
					Authorization Code Match	Generate	Apply
<input checked="" type="checkbox"/> 10.87.0.27	0500002056	WXC 500	---	---	Not Ready	N/A	N/A
<input checked="" type="checkbox"/> 10.87.41.2	0050003173	WX 5x	---	---	Not Ready	N/A	N/A
<input checked="" type="checkbox"/> 10.87.42.2	0500001995	WXC 500	---	---	Not Ready	N/A	N/A
<input checked="" type="checkbox"/> 10.87.45.2	0050003072	WX 5x	---	---	Not Ready	N/A	N/A

[Match Authorization Codes](#) [Apply Licenses](#)

Available Valid Speed Authorization Codes

WX 15 (64 K base)	Quantity
Valid authorization codes are not available	

Available Valid IPsec Authorization Codes

Device Model No.	Quantity
Valid authorization codes are not available	

WX 20 (64 K base) Quantity

64 K --> 512 K	1
----------------	---

WX 5x (256 K base) Quantity

256 K --> 4 M	1
---------------	---

Figure 5-6 Generating and Applying Licenses

The devices with evaluation licenses are listed (if any), followed by the imported speed and IPSec Authorization Codes that are available to be matched with a device. If a device could not be reached, its serial number and model are displayed as “Unknown.”

NOTE: The speed Authorization Codes show the new device names (an SR is a “WX” and an SM is a “WXC”), but they can still be applied to SRs and SMs.

To poll another community or device group, click **Change** in the upper-right corner of the page.

3. To match the available Authorization Codes with the listed devices, you must change the default selection (---) for both the **Speed** and **IPSec** fields for each device, and click **Match Authorization Codes**.

Speed Authorization Codes	Select the speed Authorization Code that you want assigned to the device. Select None to generate a license for the base speed (no Authorization Code required). The base speeds for each device type are shown in parentheses below the device list.
IPSec Authorization Codes	Select whether you want an IPSec Authorization Code assigned to the device (Yes or No).

If you do not have enough Authorization Codes for all devices, clear the check box next to the less-critical devices, and click **Match Authorization Codes** again. This runs the match for just the selected devices. You can run the match as often as needed.

4. The following licensing information is shown for each device:

Authorization Code Match	Indicates whether an imported Authorization Code matched the device: <ul style="list-style-type: none">• Excluded. The device is excluded from the matching process (the check box is not selected).• Not Ready. The speed and/or IPSec Authorization Code have not been selected.• Successful. Imported Authorization Codes matched the device. The list of available Authorization Code is adjusted accordingly.• Unavailable. No match for one or both of the selected Authorization Code (unmatched Authorization Code are highlighted in yellow).
--------------------------	---

Generate	Indicates whether a license has been generated: <ul style="list-style-type: none">• N/A. No attempt made.• Successful. License has been generated.• Failed. License generation failed (contact Technical Support).
Apply	Indicates whether a license has been applied: <ul style="list-style-type: none">• N/A. No attempt made.• Successful. License has been applied.• Failed. License could not be applied. Verify that the device is reachable and try again. If the problem persists, contact Technical Support.

5. To generate licenses for devices that have a successful Authorization Code match:
 - a. Click **Generate Licenses**.
 - b. Enter the user name and password for your customer account on the License Server, and click **Submit**. Enter the requested information in all of the fields, and click **Submit**.
 - c. License generation begins. The **Generate** column indicates the success or failure of the license generation for each device.
6. When license generation is complete, click **Apply Licenses** to download the successfully generated licenses to each device. If the last attempt to apply a license failed, the CMS server tries to apply the license again.

Applying the licenses may take some time. You can view the status for each device on the License Status page, as described in the next section.

Viewing the License Status

For each device for which you have successfully generated a license, the License Status page shows the number of attempts to apply the license to the device (if any), and the results of the last attempt.

To view the license status:

1. Click **MANAGEMENT** in the menu frame, click **License Management** in the left-hand navigation frame, and then click **License Status**.

The screenshot shows the PeriScope 1S Central Management System interface. The top navigation bar includes links for MY PERIBIT, MONITOR, MANAGEMENT, CMS SETUP, and ABOUT. The user is logged in as 'root'. The left sidebar shows the 'Management' menu with options like Devices, Configurations, Auto-Deployment, License Management (highlighted), Schedules, and Schedule log. The main content area is titled 'License Management > License Status' and contains a table with the following data:

IP Address	Download Attempts	Last Attempt	Status	Description
192.168.52.199	1	2004-08-18 16:49:44.0	License Application Successful	
192.168.52.200	1	2004-08-24 14:13:11.0	License Application Successful	
192.168.53.5	1	2004-08-20 12:36:25.0	License Application Successful	
192.168.53.180	1	2004-08-23 19:40:29.0	License Application Successful	
192.168.53.181	0	2004-08-24 15:19:19.0	License Generation Successful	
192.168.55.200	1	2004-08-18 17:34:20.0	License Application Successful	

Below the table is a 'Remove' button with the text: 'Click this button to remove devices that have licenses applied successfully'.

Figure 5-7 Viewing the License Status

2. The following information is provided for each device:

IP Address	The IP address of the device.
Download Attempts	Number of attempts to apply the license to the device.
Last Attempt	Date and time of the last attempt to apply the license.
Status	Indicates one of the following: <ul style="list-style-type: none"> • License generated. No attempt to apply the license. • License applied. License applied successfully. • License application failed. Last attempt to apply the license failed.
Description	Provides additional information if the license application fails. The most common problems are: <ul style="list-style-type: none"> • AUTH_FAILURE. The device belongs to a community that has not been imported. To import the community, refer to “Managing Communities” on page 328. • CONNECT_TIMEOUT or CONNECT_FAILURE. Network problem or the device may be down. For other types of errors, contact Technical Support.

3. Click **Remove** to remove the status entries for the licenses that were applied successfully.

Chapter 6 Monitoring Performance

This chapter describes how to use CMS to monitor the device performance. It covers the following topics:

- “Viewing and Printing Reports” in the next section
- “Configuring the My Peribit Page” on page 270
- “Viewing Reports on the Monitor Page” on page 274

Viewing and Printing Reports

Note the following about viewing and printing reports:

- Most reports are generated from a local database populated by periodic polling of the SRS 5.x devices, and report times are based on the local server time, not the device time. On device-specific reports, you can select **Show in Device Time** to view the report in the device’s time (the reported times will be accurate only if the device’s time zone is set correctly).

For example, if the device time is 8:30 AM and the server time is 11:30 AM, a report for “Today” displays 11 hours of data (12:00 AM through 11:00 AM) in the server’s time, and 8 hours of data in the device’s time.

- Hourly aggregation may take up to 28 minutes. For example, performance data from 2:00 PM to 3:00 PM may not be shown on reports until 3:28 PM.

NOTE: The CMS report for a device may not match the SRS report if the device does not respond to a poll, or if the device is rebooted after a poll. When you reboot a device, the data for the current and past hour are purged from the device.

- If you enter an arbitrary date and time range for a report, 24 hours of performance data are shown for each day, regardless of the specified start and end times.
- To print a report on the Monitor page, select **Printer Friendly Format** in the left-hand navigation frame and click **Submit**. The report opens in a new browser window. Use the browser’s Print function to print the report.

Configuring the My Peribit Page

The My Peribit page lets each user create a customized mix of charts that depict the overall performance of the devices in one or all communities or device groups. The available charts include:

- The ten applications or endpoints with the highest or lowest reduction or acceleration
- The total traffic and dropped traffic for the top ten outbound QoS traffic classes, and the four inbound traffic classes
- The ten monitored applications with the highest percentage of traffic
- The ten monitored endpoints with the worst latency, loss, or availability as measured by WAN performance monitoring

To configure the My Peribit page:

1. Click **MY PERIBIT** in the menu frame to view the My Peribit page for the current user account.

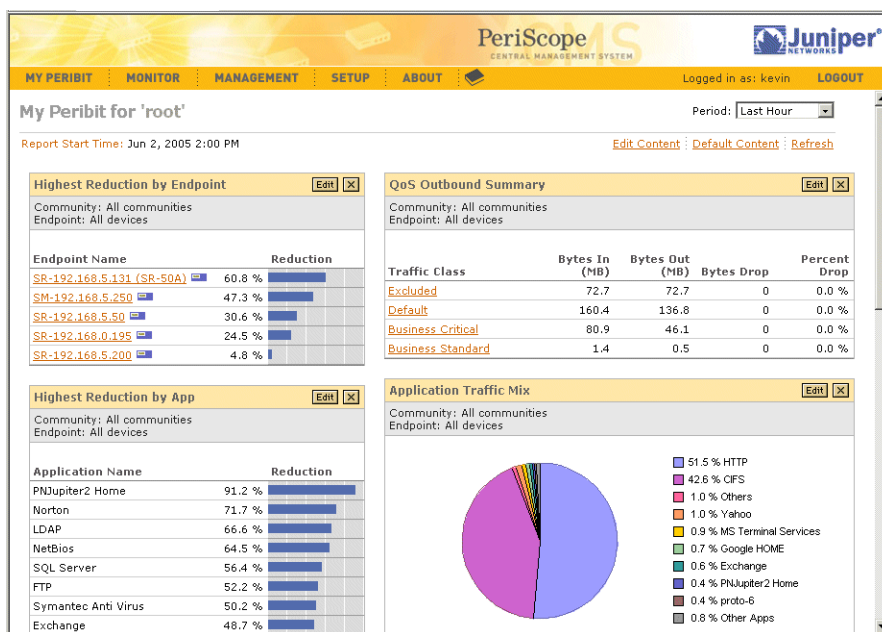


Figure 6-1 My Peribit Page

2. To display the default charts, click **Default Content**.

3. To change the time period for the displayed charts, select a time period from the **Period** menu in the upper-right corner of the page. You can view the My Peribit charts for up to the last month.
4. To change or delete a specific chart on the My Peribit page, click the Edit or delete buttons in the title bar of the chart.
5. To change the content or layout of the page, click **Edit Content**.

Figure 6-2 My Peribit Page

6. To add a chart, specify the following and click **Add**:
 - a. Select the chart from the **Item** menu.
 - b. Select a specific community (or device group) and device, as needed, from the **Community/Device Group** and **Endpoint** menus (the default is all device groups and communities). You cannot select a specific device for the “by Endpoint” charts.
 - c. Optionally, change the default title of the chart in the **Title** field.


Table 6-1 describes the available charts. If you add the same chart multiple times, such as for different communities or device groups, a number is appended to the title automatically. Note that narrow charts are displayed in the left column; wide charts are displayed in the right column.

Table 6-1 My Peribit Charts

Chart	Size	Description
Highest Acceleration by App	Narrow	The ten applications that have the highest acceleration gains from Active Flow Pipelining, Fast Connection Setup, or CIFS, Exchange, and HTTP acceleration.
Lowest Acceleration by App	Narrow	The ten applications that have the lowest acceleration gains.
Highest Acceleration by Endpoint	Narrow	The ten devices that have the highest acceleration gains from Active Flow Pipelining or Fast Connection Setup. Click a device name on the chart to view acceleration results by application for the selected device (refer to Figure 6-26 on page 309 and Figure 6-27 on page 312).
Lowest Acceleration by Endpoint	Narrow	The ten devices whose applications have the lowest acceleration gains.
Highest Reduction by App	Narrow	The ten applications that have the highest percentage of data reduction.
Lowest Reduction by App	Narrow	The ten applications that have the lowest percentage of data reduction.
Highest Reduction by Endpoint	Wide	The ten devices that have the highest percentage of data reduction. Click a device name on the chart to view the reduction details from the selected device to each of the remote devices (refer to Figure 6-11 on page 287).
Lowest Reduction by Endpoint	Wide	The ten devices that have the lowest percentage of data reduction.
Application Traffic Mix	Wide	A pie chart of the nine monitored applications that have the highest percentage of the traffic into the selected device(s). The tenth "Other Apps" category indicates the percentage of the traffic for all of the other reduced applications. The "Others" category is for reduced applications that are undefined or unmonitored.

Table 6-1 My Peribit Charts

Chart	Size	Description
QoS Outbound Summary	Wide	The ten QoS traffic classes with the most outbound traffic. Includes the total number of bytes in and out of the selected devices for each class, and the number and percentage of bytes dropped (if any). Click a class name to view the traffic by device for the selected class, and then click a device to view traffic from the selected device to each QoS endpoint (refer to Figure 6-20 on page 301).
QoS Inbound Summary	Wide	The total number of bytes into and out of the selected devices for each inbound traffic class, and the number and percentage of bytes dropped (if any). Click a class name to view the traffic by device for the selected class, and then click a device to view traffic into the selected device from all other devices (refer to Figure 6-23 on page 305)
Lowest Availability	Wide	The ten devices that have the lowest percentage of availability as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the low-availability device (refer to Figure 6-11 on page 287).
Highest Time Above Latency Threshold	Wide	The ten devices where the average latency exceeded the latency threshold for the highest percentage of time, as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the high-latency device (refer to Figure 6-11 on page 287).
Highest Loss	Wide	The ten devices that have the highest percentage of probe packet loss as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the high-loss device (refer to Figure 6-11 on page 287).

7. To delete a chart or change its position on the page:
- a. To position a chart on the page, select the chart in the Narrow Column or Wide Column lists, and click the up or down arrow keys.
 - b. To delete a chart, select the chart, and click .
 - c. Click **Apply** to save the changes and stay on the page, or click **Finish** to return to the My Peribit page.

Viewing Reports on the Monitor Page

The following topics describe the reports available on the Monitor page:

- “WAN Statistics” on page 274
- “Data Reduction Statistics” on page 283
- “QoS Statistics” on page 300
- “Acceleration Statistics” on page 308
- “Top Traffic Statistics” on page 317
- “Executive Summary” on page 321

WAN Statistics

This section describes the WAN reports.

- “WAN Performance Statistics” on page 275
- “WAN Throughput Statistics” on page 280
- “WAN Application Summary” on page 282

WAN Performance Statistics

The WAN Performance, Loss, Latency, and Availability reports show the WAN performance statistics and events, as measured from each device where WAN Performance Monitoring or Policy-Based Multi-Path is enabled for one or more remote devices.

To view the WAN performance reports:

1. Click **MONITOR** in the menu frame, and **WAN** in the navigation frame.
2. Select one of the following reports from the **Statistic** menu. Each report provides the same statistics when a specific device is selected. When **All devices** is selected, each report shows a different statistic as a matrix of color-coded cells to indicate the status between each pair of devices:
 - **WAN Performance.** Best, average, or worst values measured for loss, latency, and availability (the default is average). To change the performance ranges represented by each color, refer to “Setting WAN Reporting Thresholds” on page 348.
 - **WAN Latency.** Percentage of the selected time period that the average latency exceeded the specified threshold.
 - **WAN Loss.** Percentage of probe packets that were lost.
 - **WAN Availability.** Percentage of the minutes in the selected time period for which at least one probe was acknowledged.
3. Change one or more of the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a specific device from the **Device** menu to view a table of performance statistics measured from the selected device to each remote device. The default is **All devices**, which shows a color-coded matrix for all monitored devices.
 - Select a time period from the **Period** menu. You can select the previous hour, day, week, month, or six months. The default is **Last Hour**.
4. If **WAN Performance** and **All devices** are selected, click **Submit**. to open the WAN Performance page for the selected community or device group.

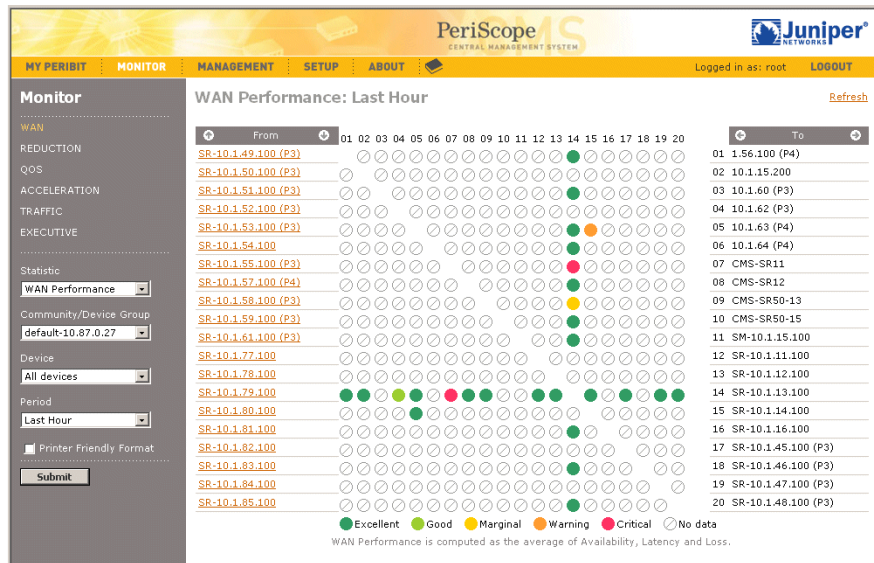


Figure 6-3 WAN Performance for All Devices in a Community/Device Group

From the report page, you can:

- View the WAN performance from each device in the **From** column to each device in the **To** column. The same devices are listed in both columns so you can see the performance in both directions. A cell is white if both WAN Performance Monitoring and Policy-Based Multi-Path are disabled between the devices.
 - Move the cursor over a cell to highlight the two devices and display the exact percentages in the browser's status bar for loss, time above the latency threshold, and availability (measured by the **From** device).
 - View the next group of devices by moving the cursor over the **From** or **To** column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.
5. To view the WAN performance statistics between a specific device and each of the other devices it is monitoring, click the device name in the **From** column, or select the device from the **Device** menu and click **Submit**.

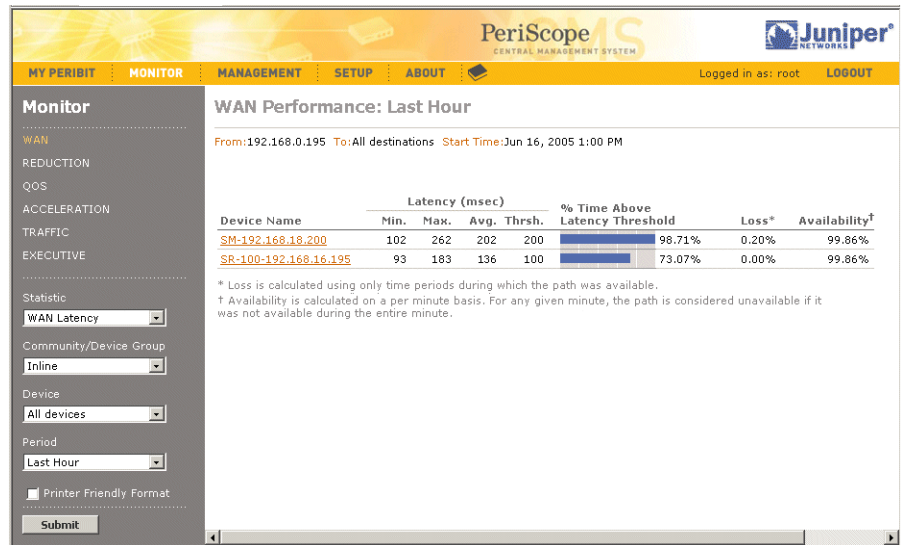


Figure 6-4 WAN Performance Statistics for All Destinations

The following information is shown for each monitored device.

- **Device Name.** Name of the remote device. Devices that support Multi-Path have a “_Pri” or “_Sec” appended to the device name to indicate the primary or secondary path.
 - **Latency (msec).** Probes are used to measure the lowest, highest, and average round-trip times between the selected device and the remote device (in milliseconds). The latency threshold is also displayed.
 - **% Time Above Latency Threshold.** Percentage of the selected time period that the average latency exceeded the specified threshold.
 - **Loss.** Percentage of the probes that were lost.
 - **Availability.** Percentage of the minutes in the selected time period for which at least one probe was acknowledged.
6. To view the WAN performance graphs and events for a specific device, click the device name or select a colored cell on the matrix shown for **All devices**. The information on the performance graphs depends on whether the device is enabled for Multi-Path (Figure 6-5) or WAN performance monitoring (Figure 6-6).

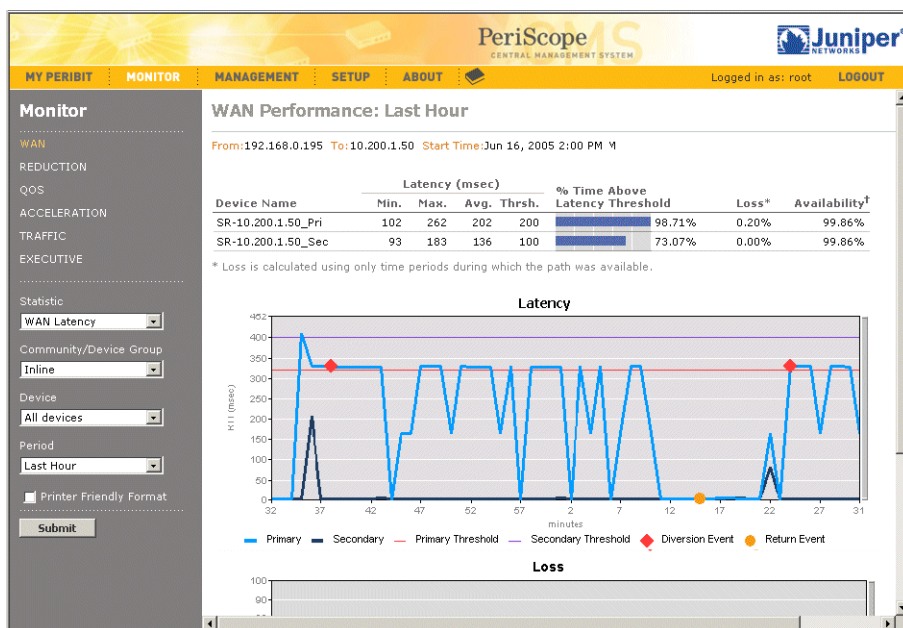




Figure 6-5 Multi-Path WAN Performance Charts

For a Multi-Path device, the following information is shown on the Loss and Latency charts (Figure 6-5):

- The Latency chart shows the average round-trip time for the primary path (blue) and secondary path (black), and indicates the configured latency threshold for each path.
- The following icons are used to indicate performance events. Move the cursor over the icon to view the number of events in the time period.

Icon	Description
	<p>Indicates that traffic was switched to the alternate path due to one of the following conditions:</p> <ul style="list-style-type: none"> • Loss or latency threshold exceeded. Eligible traffic is diverted only if the alternate path's reduction tunnel is up and the loss and latency are below the specified thresholds. • Reduction tunnel is down. Eligible traffic is diverted regardless of the alternate path's performance (if the alternate reduction tunnel is up). Traffic that cannot be switched to the alternate path is passed through without reduction (if the link is up and only the reduction tunnel is down).

Icon	Description
	Indicates that performance has returned to normal, and traffic was switched back to the preferred path (the reduction tunnel must be up).
	Indicates the loss or latency threshold was exceeded, but no traffic was diverted (such as when both paths are degraded). For time periods longer than one hour, the icon may represent multiple types of events. Move the cursor over the icon to view the number of each type of event that occurred in the time period.

- The Loss chart shows the percentage of the probes that were lost on the primary and secondary paths. If the loss threshold is exceeded, a diversion to the alternate path is indicated on the Latency chart (if the alternate path is not degraded).

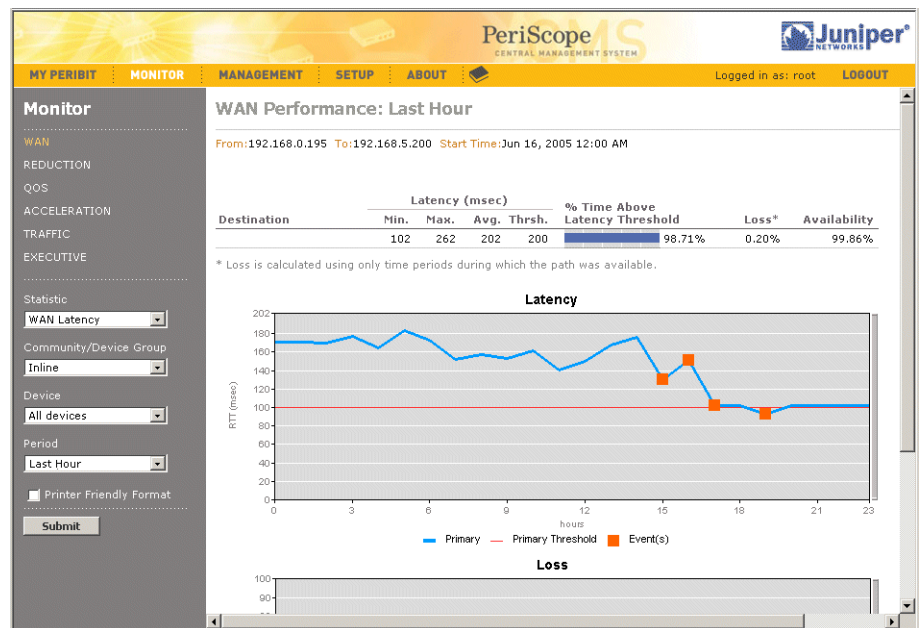



Figure 6-6 Single-Path WAN Performance Charts

For WAN performance monitoring endpoints (Figure 6-6), the loss and latency are shown for a single path, and the  icon indicates the loss or latency threshold was exceeded.

NOTE: If the remote device is unreachable, all paths will be down, the Latency chart will be blank (latency cannot be measured), and the Loss chart will show 100% probe loss on all paths.

WAN Throughput Statistics

The WAN Throughput report shows the speed of the traffic to and from the WAN for a specific device. You can view a device's WAN throughput to all remote destinations, a specific remote SR/SM device, or all non-SR/SM destinations.

To view WAN throughput:

1. Click **MONITOR** in the menu frame, and **WAN** in the navigation frame.
2. Select **WAN Throughput** from the **Statistic** menu.
3. Select a community or device group from the **Community/Device Group** menu., and select a device from the **Device** menu.
4. Change one or more of the following report parameters, and click **Submit**.
 - Select a specific monitored application from the **Application** menu. Select **Others** to view statistics for applications that are undefined or unmonitored. The default is **All applications**. To specify the monitored applications on a device, refer to “Monitoring Applications” on page 138.
 - Select a specific SR/SM device from the **Destination** menu. Select **Other Destinations** to view statistics for all non-SR/SM destinations.
 - Select a time period from the **Period** menu. You can select the previous hour, day, week, month, or six months. The default is **Last Hour**.

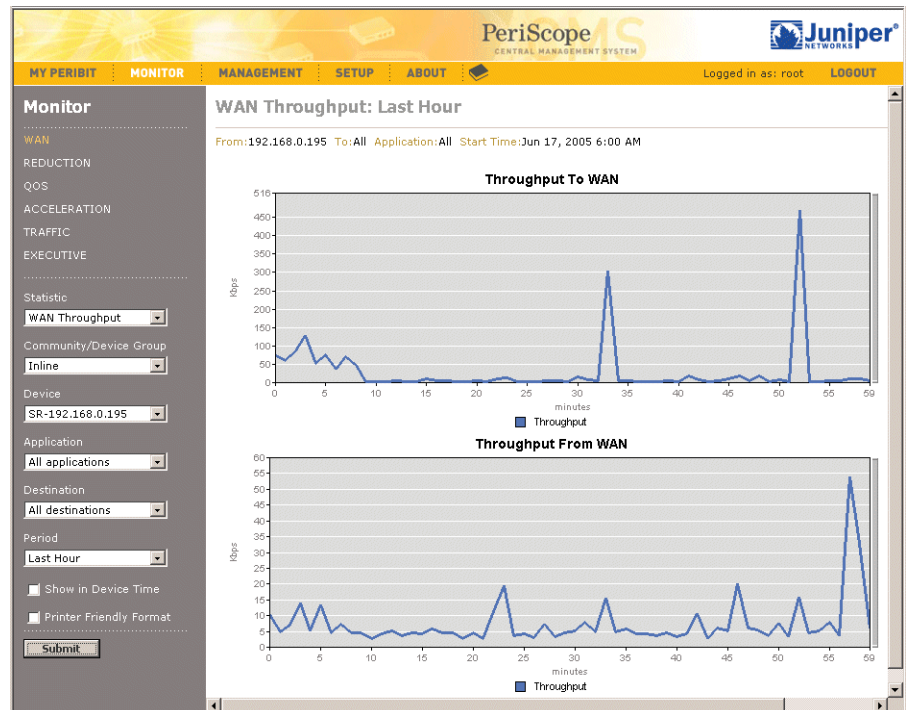


Figure 6-7 WAN Throughput Report

5. Review the following information on the two throughput graphs. Keep in mind that all values are for the selected application, destination, and time period.
 - The Throughput to WAN graph shows the average throughput of data sent to the WAN.
 - The Throughput From WAN graph shows the average throughput of data received from the WAN. This graph is blank when the device is in Profile Mode.

WAN Application Summary

The WAN Application Summary shows the application traffic to and from the WAN for a specific device. You can view a device's WAN traffic to all remote destinations, a specific remote SR/SM device, or all non-SR/SM destinations. The traffic is shown for up to 40 monitored applications.

To view the WAN Application Summary:

1. Click **MONITOR** in the menu frame, and **WAN** in the navigation frame.
2. Select **WAN App. Summary** from the **Statistic** menu.
3. Select a community or device group from the **Community/Device Group** menu., and select a device from the **Device** menu.
4. Change one or more of the following report parameters, and click **Submit**.
 - Select a specific SR/SM device from the **Destination** menu. Select **Other Destinations** to view statistics for all non-SR/SM destinations.
 - Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

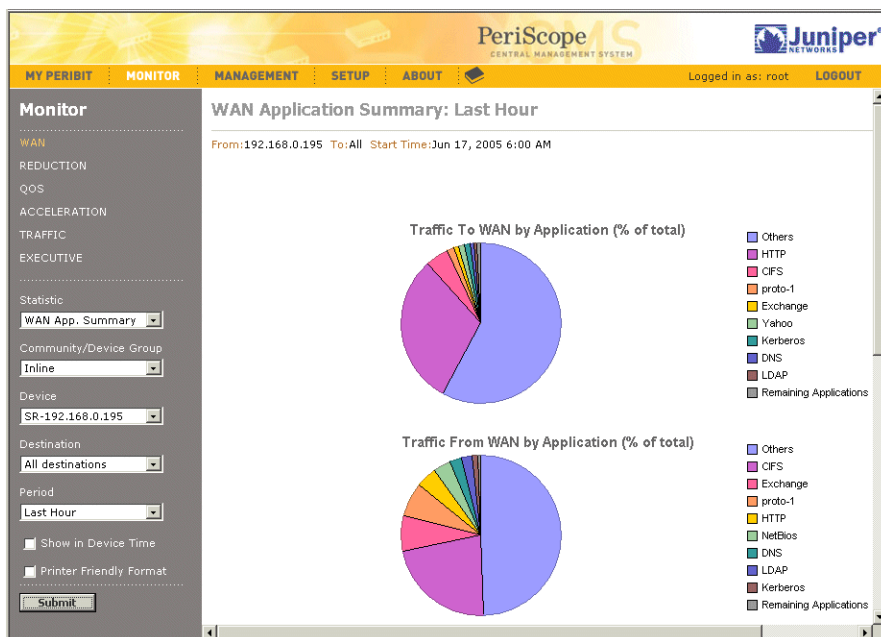


Figure 6-8 WAN Application Summary

5. Review the information on the following charts. Keep in mind that all values are for the selected destination and time period.
 - The two pie charts show the nine monitored applications that have the highest percentage of the total traffic sent to and from the WAN for the selected device. The **Remaining applications** category shows the traffic percentage for all other applications.
 - The application table shows the traffic in megabytes sent to and from the WAN for each monitored application. The applications are sorted in descending order by total traffic. The **Others** category indicates the traffic for applications that are undefined or unmonitored.

Reduction Statistics

This section describes the reduction reports:

- “Data Reduction Statistics” in the next section
- “Application Summary Statistics” on page 291
- “Passthrough Statistics” on page 293
- “Packet Size Distribution Statistics” on page 295
- “Monitoring Tunnel Status” on page 296

Data Reduction Statistics

The Reduction reports let you view the percentage of data reduction for:

- Each pair of SR/SM devices in the same community (matrix view).
- A selected device and each of the other devices in a community/device group.
- Each application for a selected pair of devices.
- A selected application from a specific device to each of the other devices in the same community.

The percentage of data reduction for the selected time period is based on the total number of bytes in and out of each device, as follows

$$\% \text{ of Reduction} = \left(\frac{\text{Bytes In} - \text{Bytes Out}}{\text{Bytes In}} \right) \times 100$$

To view the Reduction reports:

1. Click **MONITOR** in the menu frame, and **REDUCTION** in the navigation frame.
2. Select one of the following reports from the **Statistic** menu:
 - **Reduction Overview.** Percentage of reduction shown in a matrix between each pair of devices in the selected community or device group (when **All devices** is selected). Select a cell in the matrix to view graphs of reduction by time. Select a device from the **Device** menu to view reduction from the selected device to each remote device for one or all applications (in one or both directions).
 - **Reduction by Time.** Percentage of reduction by time from a selected device to one or all remote destinations, for one or all applications.
 - **Reduction by Dest.** Same as the Reduction Overview, except that a specific device must be selected (no matrix view).
 - **Reduction by App.** Percentage of reduction for all applications from a selected device to a remote destination.
3. Change one or more of the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a specific device from the **Device** menu to view reduction statistics measured from the selected device to each remote device.
 - Select a specific monitored application from the **Application** menu. Select **Others** to view statistics for applications that are undefined or unmonitored. The default is **All applications**. To specify the monitored applications on a device, refer to “Monitoring Applications” on page 138.
 - Select a specific device from the **Destination** menu to view reduction statistics measured to a specific device (on the by-time and by-application reports only).
 - Select a time period from the **Period** menu. You can select the previous hour, day, week, month, or six months. The default is **Last Hour**.
4. If **Reduction Overview** and **All devices** are selected, click **Submit**. to open the Reduction page for all devices in the selected community/device group.

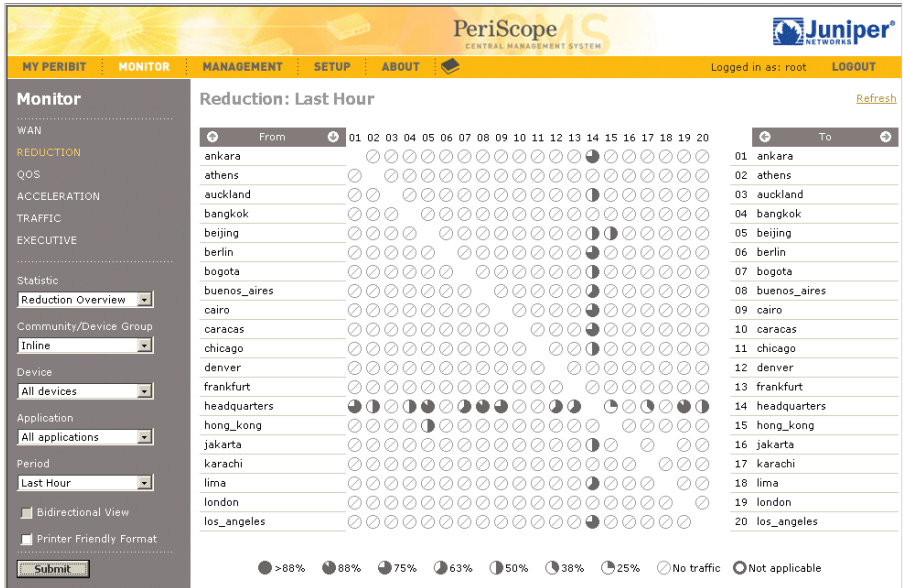



Figure 6-9 Percentage of Data Reduction for a Community/Device Group

From the report page, you can:

- View the percentage of data reduction for traffic sent from each device in the **From** column to each device in the **To** column. The same devices are listed in both columns so you can see the reduction in both directions. The icons indicate a percentage range.
- The  icon indicates that the device is down or unreachable, or there is no data reduction.
- Move the cursor over an icon to highlight the two devices and display the exact data reduction percentage in the browser's status bar, along with the number of bytes and packets in and out of the **From** device. Note that the reduction percentage indicated by the icon is approximate.
- View the next group of devices by moving the cursor over the **From** or **To** column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.

- Click the icon for a pair of devices to view graphs of reduction by time for the traffic sent from a device in the **From** column to a device in the **To** column.

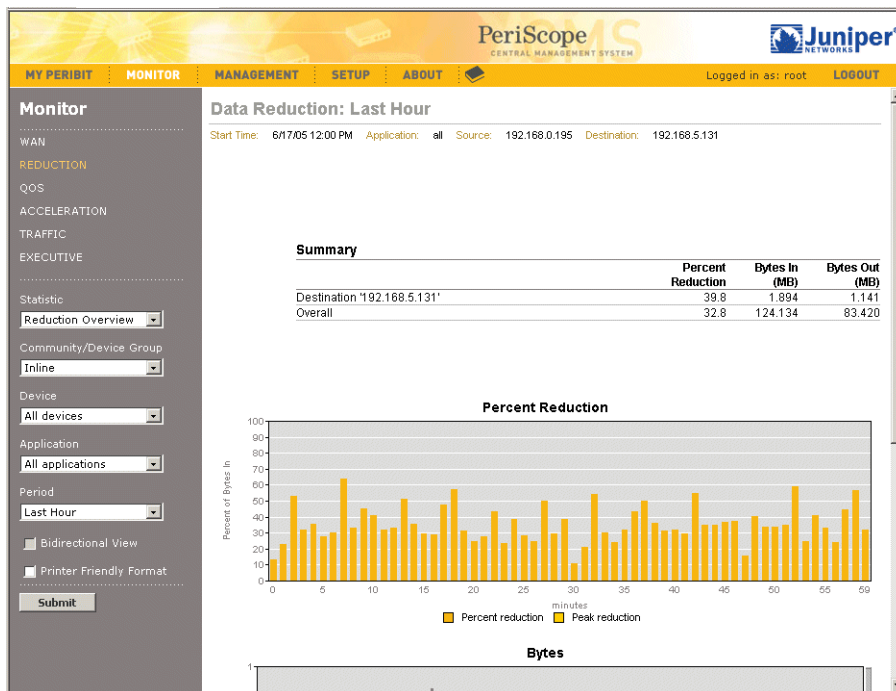


Figure 6-10 Percentage of Data Reduction By Time

The report shows the percentage data reduction for the remote device, the overall reduction for all remote devices, the number of bytes in and out of the reduction engine, and graphs of the reduction, bytes, and packets for the selected time period.

To view the same graphs from a selected device to all remote devices, select **Reduction by Time** from the **Statistic** menu and select the device from the **Device** menu.

- To view the percentage of data reduction between a specific device and each of the other devices in the same community, select the device name from the **Device** menu and click **Submit**. Alternatively, select **Reduction by Dest.** from the **Statistic** menu and select the device from the **Device** menu.

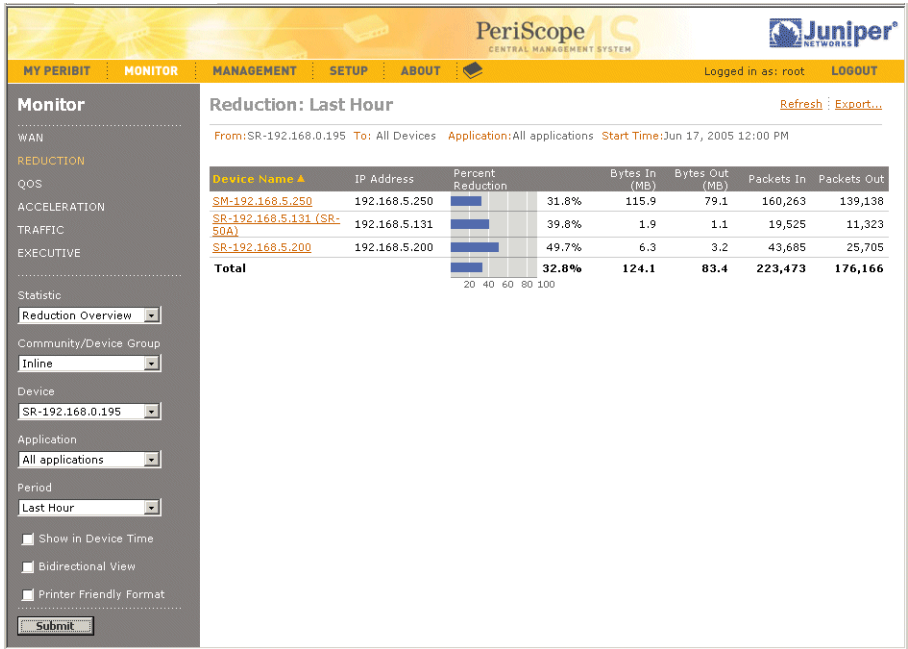


Figure 6-11 Percentage of Data Reduction for a Selected Device

From the report page, you can:

- View the outbound percentage of data reduction achieved by the selected device for each of the other devices in the same community. The number of bytes and packets in and out of the reduction engine on the selected device is shown for each destination device.
- Click **Export** to view or save the displayed data in CSV format.

NOTE: If the selected device resides in multiple communities, the report includes data reduction statistics for devices in each community.

- Click a device name to view the percentage of data reduction for each reduced application in the traffic sent to the device. Alternatively, select **Reduction by App.** from the **Statistic** menu and select the device from the **Device** menu.

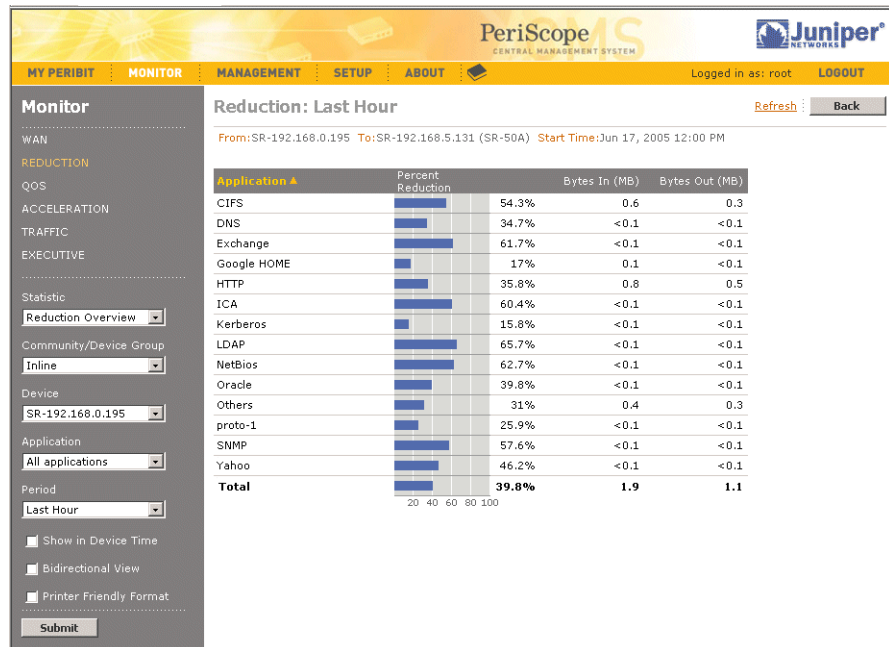


Figure 6-12 Percentage of Data Reduction By Application

The report displays the percentage data reduction and number of bytes in and out of the device for each reduced application in the selected time period.

Note that the **Others** application is for applications that are undefined or unmonitored on the **From** device.

- Click **Bidirectional View** and click **Submit** to view the inbound and outbound percentage of data reduction between the selected device and each of the other devices in the same community (Figure 6-13 on page 289).

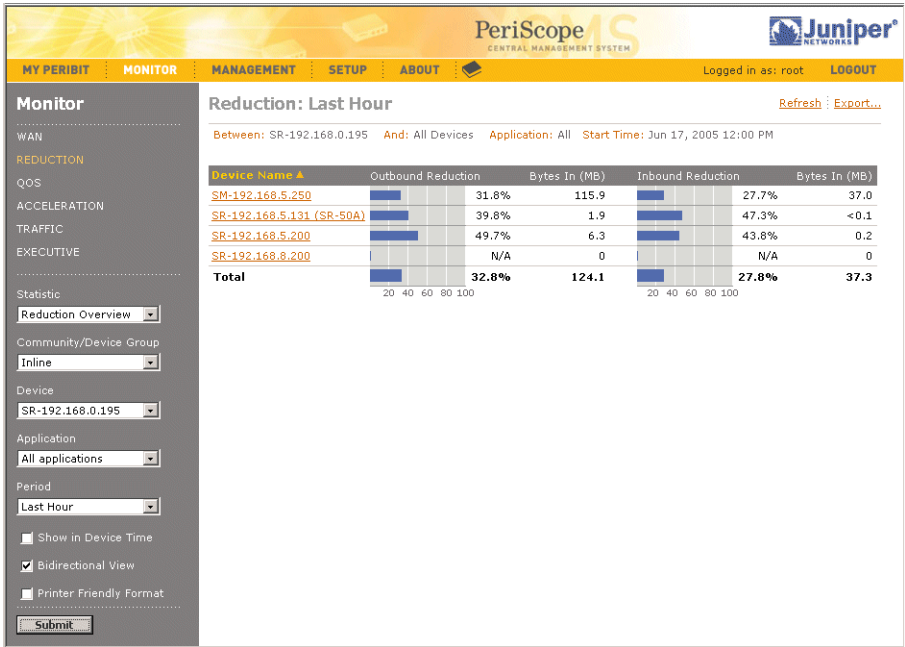


Figure 6-13 Bidirectional Data Reduction for a Selected Device

In the bidirectional view, note that the first **Bytes In** column shows the number of bytes into the reduction engine of the selected device (the **From** device); the second **Bytes In** column shows the number of bytes into the reduction engines of each of the other devices in the same community.

9. To view the percentage of data reduction for a specific application from the selected device to each of the other devices in the same community, select a monitored application from the **Application** menu, and click **Submit**.

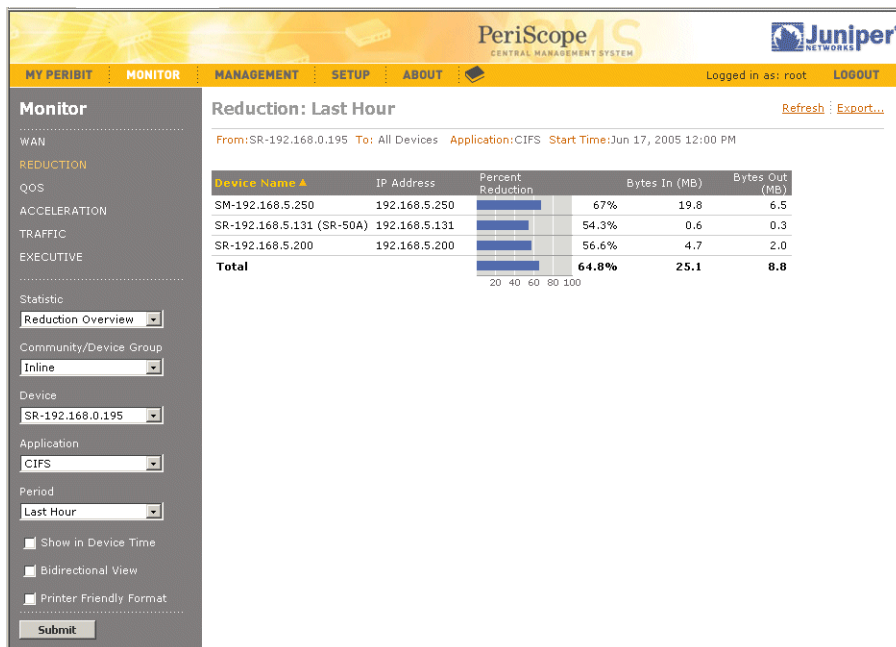


Figure 6-14 Percentage of Data Reduction for a Selected Application

From the report page, you can:

- View the percentage of data reduction achieved for each destination device by the selected device and application. For the selected application, the number of bytes and packets in and out of the reduction engine on the selected device is shown for each destination device.
- Click **Bidirectional View** and click **Submit** to view the inbound and outbound percentage of data reduction for the selected application between the selected device and each of the other devices in the same community.
- Click **Export** to view or save the displayed data in CSV format.

Application Summary Statistics

The Application Summary shows a pie chart of the nine monitored applications that have the highest percentage of the traffic into a selected device for one or all remote devices. A table is also included that shows the traffic statistics and percentage of data reduction for each monitored application (up to 40).

To view the Application Summary:

1. Click **MONITOR** in the menu frame, and **REDUCTION** in the navigation frame.
2. Select **Application Summary** from the **Statistic** menu.
3. Select a community or device group from the **Community/Device Group** menu., and select a device from the **Device** menu.
4. Change one or more of the following report parameters, and click **Submit**.
 - Select a specific device from the **Destination** menu.
 - Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

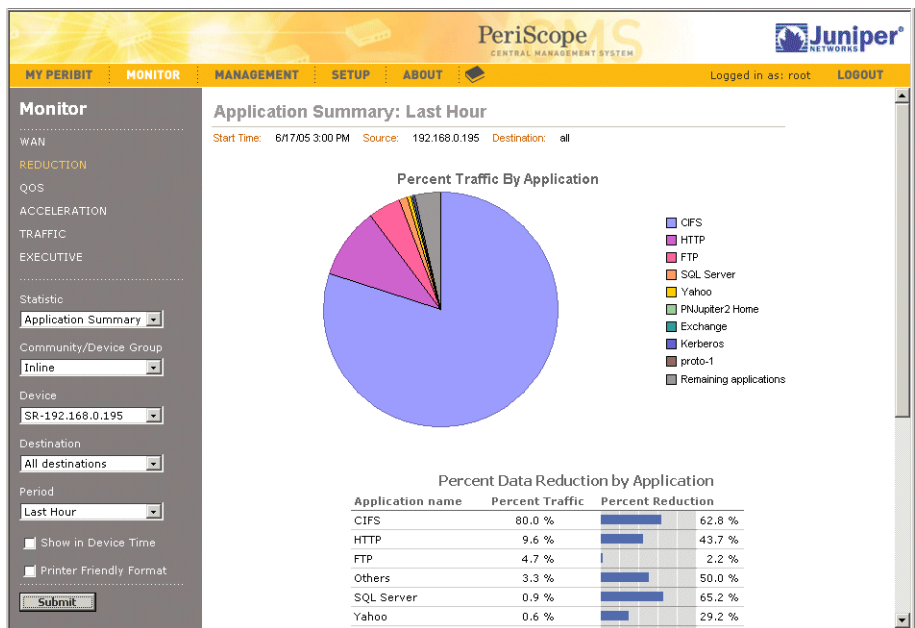


Figure 6-15 Application Summary Statistics

5. Review the following information on the Application Summary.
 - The pie chart shows the nine monitored applications with the highest percentage of the total traffic into the device for the selected destination. The **Remaining applications** category shows the traffic for all other applications (both defined and undefined).
 - The application table has the following columns.
 - **Application Name.** Names of the monitored applications, sorted in descending order by reduction percentage. The **Others** category indicates the traffic for reduced applications that are undefined or unmonitored.
 - **Percent Traffic.** Percentage of the total traffic into the device's reduction engine for each application.
 - **Percent Reduction.** Percentage of data reduction achieved for each application. A dash is shown for applications that have no traffic or cannot be reduced (such as encrypted applications). Data reduction should be disabled for applications that consistently show little or no reduction (refer to “Reducing Applications” on page 144).

Passthrough Statistics

Traffic that falls into one of several categories is passed through the devices with no attempt at data reduction. The Passthrough report shows a pie chart of the percentage of passthrough traffic in each category. A table is also included that shows the number of bytes and packets in each category.

NOTE: An off-path device that uses RIP for packet interception has no passthrough statistics. All traffic is sent through the reduction tunnel.

To view passthrough statistics:

1. Click **MONITOR** in the menu frame, and **REDUCTION** in the navigation frame.
2. Select **Passthrough** from the **Statistic** menu.
3. Select a community or device group from the **Community/Device Group** menu., and select a device from the **Device** menu.
4. fSelect a time period from the **Period** menu, and click **Submit**. You can select the previous hour, day, week, month, or six months, or enter an absolute date range.

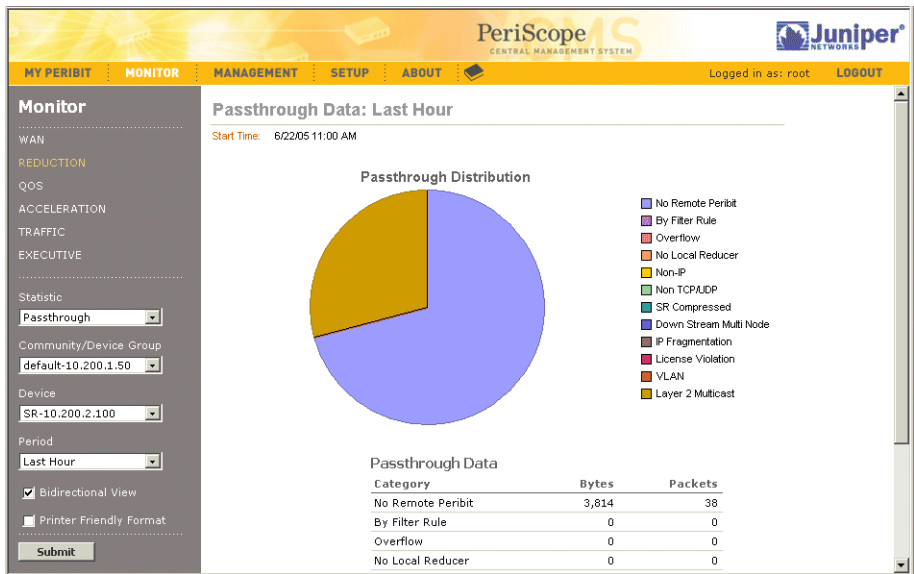


Figure 6-16 Passthrough Statistics

The following table describes the passthrough categories.

Category	Description
No Remote Peribit	No device available to assemble the data, or reduction is disabled for one or more devices.
By Filter Rule	Reduction is disabled for specific applications or source/destination addresses.
Overflow	Traffic volume exceeded the device capacity.
No Local Reducer	Reduction is disabled on this device.
Non-IP	Non-IP traffic is not reduced.
Non-TCP/UDP	By default, only TCP/UDP application traffic is reduced. This category is invalid if you define non-TCP/UDP applications.
SR Compressed	Traffic was compressed by another device.
Down Stream Multi-Node	Traffic will be reduced by the next device (refer to Multi-Node configurations in the operator's guide).
IP Fragmentation	Always zero unless reduction of IP fragments is disabled (refer to the "configure filter" CLI command in the operator's guide).
License Violation	The licensed throughput speed was exceeded.
VLAN	Total VLAN traffic that was not reduced for any reason. Includes traffic between local VLANs (non-WAN traffic) and ISL VLAN traffic.
Layer 2 Multicast	Layer 2 multicast traffic, such as for ARP, is not reduced because the intended destination is unknown.
NOTE: Jumbo Gigabit Ethernet frames are also passed through without reduction, but they are not counted in any of the above categories.	

Packet Size Distribution Statistics

For each device in a selected community or device group, the Packet Size Distribution report shows the number of packets in and out of the reduction engine for each of six packet-size ranges.

To view packet size distribution statistics:

1. Click **MONITOR** in the menu frame, and **REDUCTION** in the navigation frame.
2. Select **Packet Size Dist.** from the **Statistic** menu.
3. Select a community or device group from the **Community/Device Group** menu.
4. Select a time period from the **Period** menu and click **Submit**. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

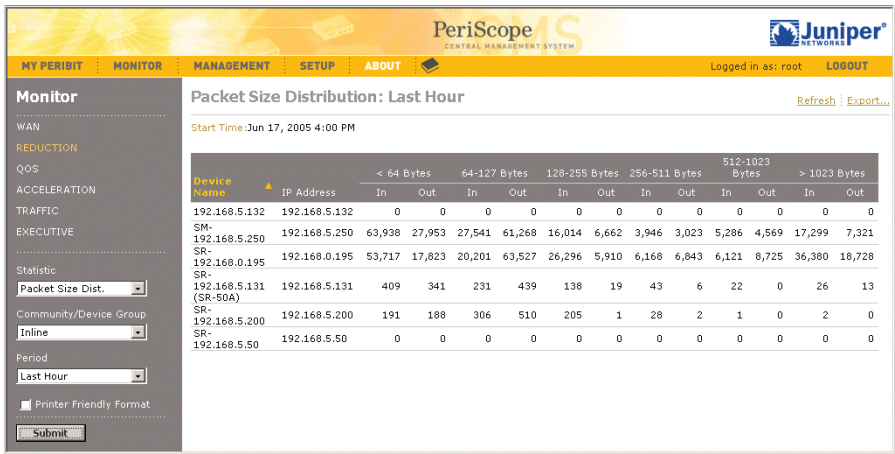


Figure 6-17 Packet Size Distribution Statistics

Monitoring Tunnel Status

By default, each SR/SM device attempts to form a pair of reduction tunnels with each of the other devices in the same community. An outbound tunnel carries reduced data to another device; an inbound tunnel carries data reduced by another device. You can configure each device to specify which tunnels are formed.

The Tunnel Status reports let you view the tunnel status for:

- The outbound tunnel for each pair of devices in the same community (matrix view).
- A selected device's outbound and inbound tunnels to and from each of the other devices in the same community (table view).

If devices in the same community are in different time zones, the CMS server time is shown in the **Last Update** field at the top of the Tunnel Status pages.

To view the Tunnel Status reports:

1. Click **MONITOR** in the menu frame, and **REDUCTION** in the navigation frame.
2. Select **Tunnel Status** from the **Statistic** menu.
3. Select a community or device group from the **Community/Device Group** menu.
4. To view a matrix showing the outbound tunnel status between each pair of devices in the same community, select **All devices** from the **Device** menu, and click **Submit**.

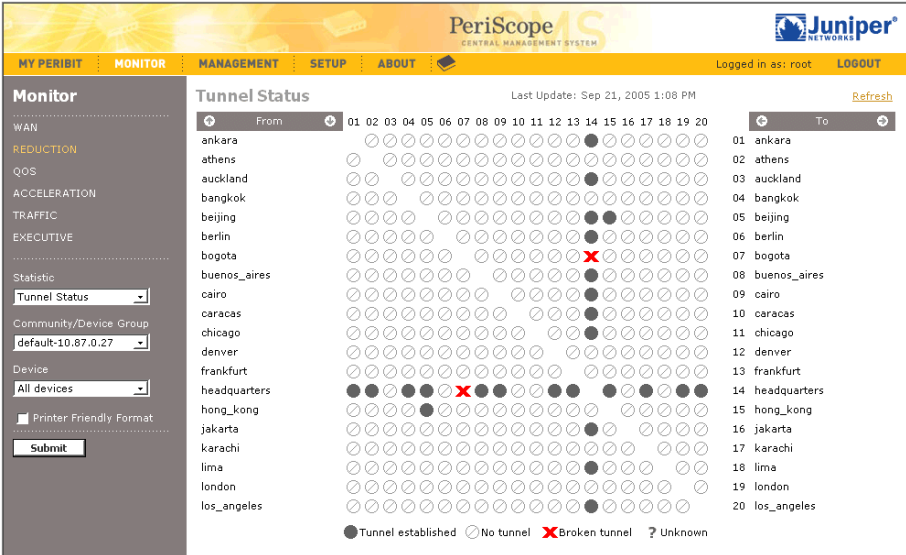






Figure 6-18 Monitoring Tunnel Status for a Community/Device Group

From the report page, you can:

- View the outbound tunnel status from each device in the **From** column to each device in the **To** column. Move the cursor over an icon to highlight both devices. The status icons are described in the following table.

Table 6-2 Tunnel Status Icons

Icon	Description
	Tunnel established — An outbound tunnel exists from a device in the From column to a device in the To column.
	No tunnel — No outbound tunnel exists due to a policy setting. For example, this icon is displayed if you manually disable data reduction from one device to another.
	Broken tunnel — No outbound tunnel exists due to an error or problem, such as low system resources. For more information, open the SRS Web console for the “from” device and click REDUCTION to view the Endpoints page. NOTE: If a device in a user-defined community is removed from the network, it will be shown with broken tunnels until it is deleted from the registration server.
	Temporarily unavailable — The tunnel is in a transitory state, or the device is down or unreachable.

- View the next group of devices by moving the cursor over the **From** or **To** column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.
 - To update the tunnel status from the devices, click **Refresh**.
5. To view a device’s outbound and inbound tunnels to and from each of the other devices in the same community, select the device name from the **Device** menu and click **Submit**.

The Tunnel Status page for the selected device opens (Figure 6-19).



Figure 6-19 Monitoring Tunnel Status for a Device

The tunnel status information shown here is the same as the status shown on the Endpoints page in the device’s SRS Web console. Note the following:

- The **OUT** column indicates the status of the outbound tunnel from the selected device to each device in the table; the **IN** column indicates the status of the inbound tunnel on the selected device from each of the listed devices.
- An **✗** icon in the **IN** column indicates that the inbound tunnel has a problem or that data reduction to the selected device is disabled on the Endpoints page of the device listed in the table.

NOTE: If the selected device resides in multiple communities, the report includes the tunnel status for devices in each community.

QoS Statistics

This section describes the QoS reports.

- “Outbound QoS Statistics” in the next section
- “Inbound QoS Statistics” on page 304

Outbound QoS Statistics

If outbound QoS is enabled on a device, the QoS Outbound reports display the following statistics for the traffic into the Local (LAN) interface and out of the Remote (WAN) interface:

- Total number of bytes and packets in and out of a selected device for each destination. Includes the number of bytes and packets dropped.
- Byte and packet counts for each traffic class on a selected device for a specific destination. Includes the throughput for each class.
- Throughput in and out of a selected device for a specific traffic class and destination. Includes the rate of dropped packets.

NOTE: Outbound QoS is not effective for an off-path device unless all outbound WAN traffic is routed through the device.

To view the QoS Outbound reports:

1. Click **MONITOR** in the menu frame, and **QOS** in the navigation frame.
2. Select **QoS Outbound** from the **Statistic** menu.
3. Select the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a device from the **Device** menu that has outbound QoS enabled. Devices using outbound QoS have a **QoS** on the Devices page (click **MANAGEMENT** in the menu frame to view the Devices page).
 - Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.



Figure 6-20 QoS Outbound Report for a Selected Device

From the QoS Outbound report page, you can:

- View the total number of bytes and packets (both reduced and unreduced) in and out of the selected device for each of the destination devices that are defined as QoS endpoints. The number of bytes and packets dropped by the device is also shown.

The **Other traffic** “device” does not have an IP address because it indicates all traffic that is not sent to a device that is designated as a QoS endpoint.

NOTE: If the selected device resides in multiple communities, the report includes the destination devices in each community.

- Click **Export** to view or save the displayed data in CSV format.
- Click a device name (destination) to view the throughput and byte and packet counts for each traffic class defined on the selected device (Figure 6-21) for the selected destination.

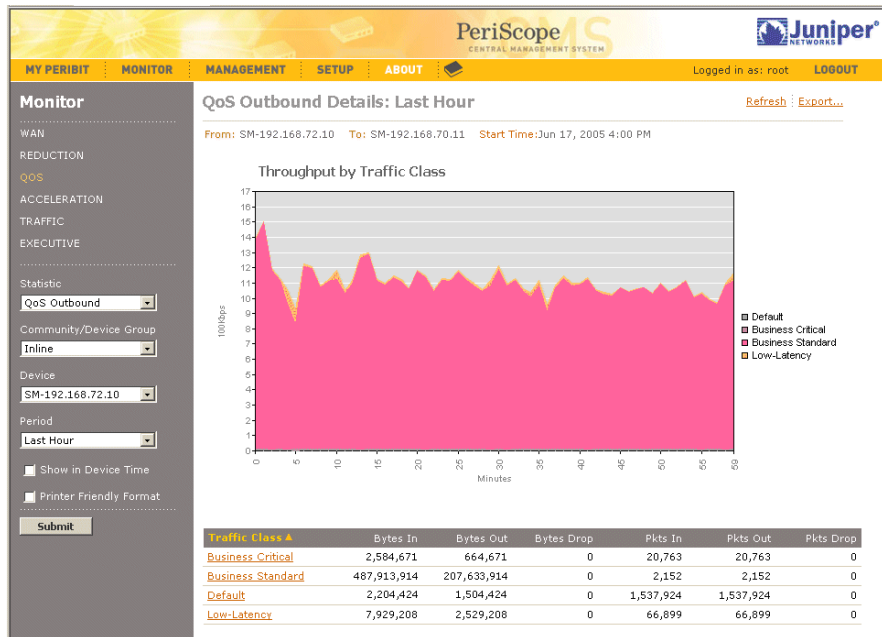


Figure 6-21 QoS Outbound Details by Traffic Class

From the QoS Outbound Details page, you can:

- View a graph of the throughput for each traffic class, and a table of the byte and packet counts for each traffic class, including the number of bytes and packets dropped by the device for this destination.
- Click **Export** to view or save the tabular data in CSV format.
- Click a traffic class name to view the throughput in and out of the device, and the rate of dropped traffic for the class (Figure 6-22).

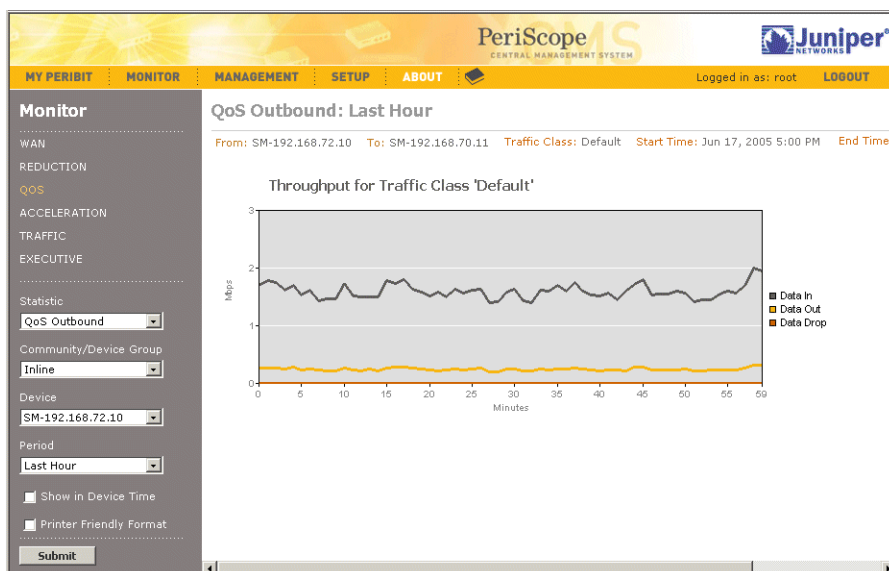


Figure 6-22 QoS Outbound Throughput for a Selected Traffic Class

The Throughput graph shows the following:

- **Data In** (grey line). Average data throughput into the Local interface from the LAN side of the device.
- **Data Out** (orange line). Average throughput to the WAN side of the device. Indicates the data reduction achieved for the selected destination.
- **Data Dropped** (red line). Average rate that outbound data was dropped. Data is dropped when the traffic for the selected class exceeds the maximum allocated bandwidth or when the guaranteed bandwidth is exceeded while the circuit is fully utilized.

Note that brief bursts of traffic can cause data to be dropped, even when the average throughput is well below the maximum bandwidth.

Inbound QoS Statistics

If inbound QoS is enabled on a device, the QoS Inbound reports display the following statistics for the traffic into the Remote (WAN) interface and out of the Local (LAN) interface:

- Total number of bytes and packets in and out of a selected device. Includes the number of bytes and packets dropped.
- Byte and packet counts for the inbound traffic classes on a selected device. Includes the throughput for each class.
- Throughput in and out of a selected device for a specific traffic class. Includes the rate of dropped packets.

NOTE: QoS Inbound reports do not apply to off-path devices.

To view the QoS Inbound reports:

1. Click **MONITOR** in the menu frame, and **QOS** in the navigation frame.
2. Select **QoS Inbound** from the **Statistic** menu.
3. Select the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a device from the **Device** menu that has inbound QoS enabled.
 - Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.



Figure 6-23 QoS Inbound Report for a Selected Device

From the QoS Inbound report page, you can:

- View the total number of bytes and packets in and out of the selected device. The number of bytes and packets dropped by the device is also shown.
- Click **Export** to view or save the displayed data in CSV format.
- Click **Inbound** to view the throughput and byte and packet counts for each of the inbound traffic classes (Figure 6-24).

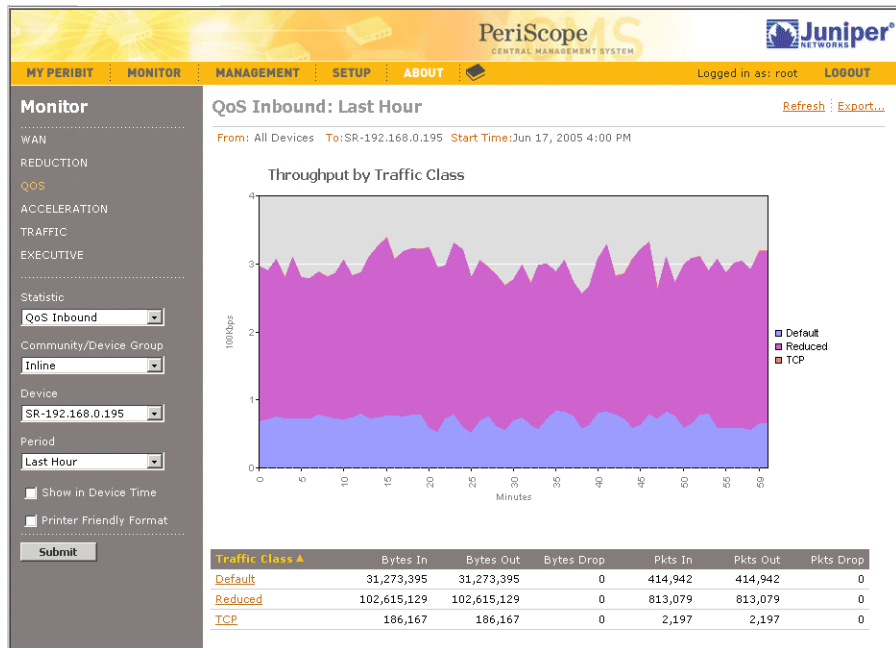


Figure 6-24 QoS Inbound Details by Traffic Class

From the QoS Inbound Details page, you can:

- View a graph of the throughput for each traffic class, and a table of the byte and packet counts for each traffic class, including the number of bytes and packets dropped by the device.
- Click **Export** to view or save the tabular data in CSV format.
- Click a traffic class name to view the throughput in and out of the device, and the rate of dropped traffic for the class (Figure 6-22).

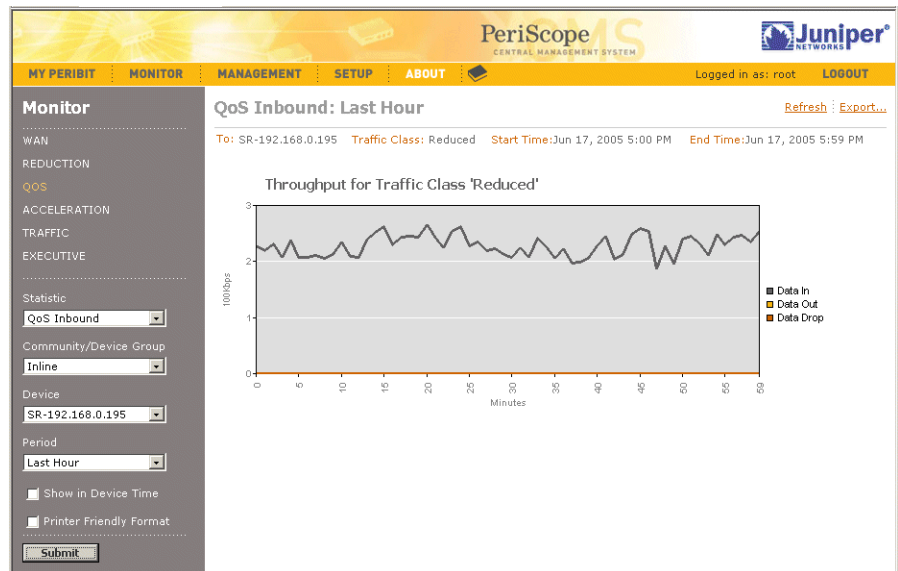


Figure 6-25 QoS Inbound Throughput for a Selected Traffic Class

The Throughput graph shows the following:

- **Data In** (grey line). Average data throughput into the Remote interface from the WAN side of the device.
- **Data Out** (orange line). Average throughput to the LAN side of the device.
- **Data Dropped** (red line). Average rate that inbound data was dropped. Data is dropped when the traffic for the selected class exceeds the maximum allocated bandwidth.

Acceleration Statistics


This section describes the acceleration reports.

- “Active Flow Pipelining Statistics” in the next section
- “Fast Connection Setup Statistics” on page 311
- “CIFS and Exchange Acceleration Statistics” on page 313
- “HTTP Acceleration Statistics” on page 315

Active Flow Pipelining Statistics

If Flow Pipelining and/or Active Flow Pipelining (AFP) is enabled for one or more endpoints and applications, the Flow Pipelining/AFP report shows the session statistics and the average throughput improvements due to Flow Pipelining and/or AFP.

To view Flow Pipelining/AFP statistics:

1. Click **MONITOR** in the menu frame, and **ACCELERATION** in the navigation frame.
2. Select **Active Flow Pipelining** from the **Statistic** menu.
3. Select the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a device from the **Device** menu that has acceleration enabled. Devices using acceleration have a  on the Devices page (click **MANAGEMENT** in the menu frame to view the Devices page).
 - Select an application from the **Application** menu to view the acceleration statistics to each remote device. Select **Others** to view statistics for applications that are undefined or unmonitored. The default is **All applications**, which shows the average acceleration for all applications to all devices.
 - Select the IP address of a specific device from the **Destination** menu to view statistics only for traffic sent to the selected device. The default is **All destinations**.
 - Select a time period from the **Period** menu. Select the previous hour, day,

week, month, or six months, or enter an absolute date range.

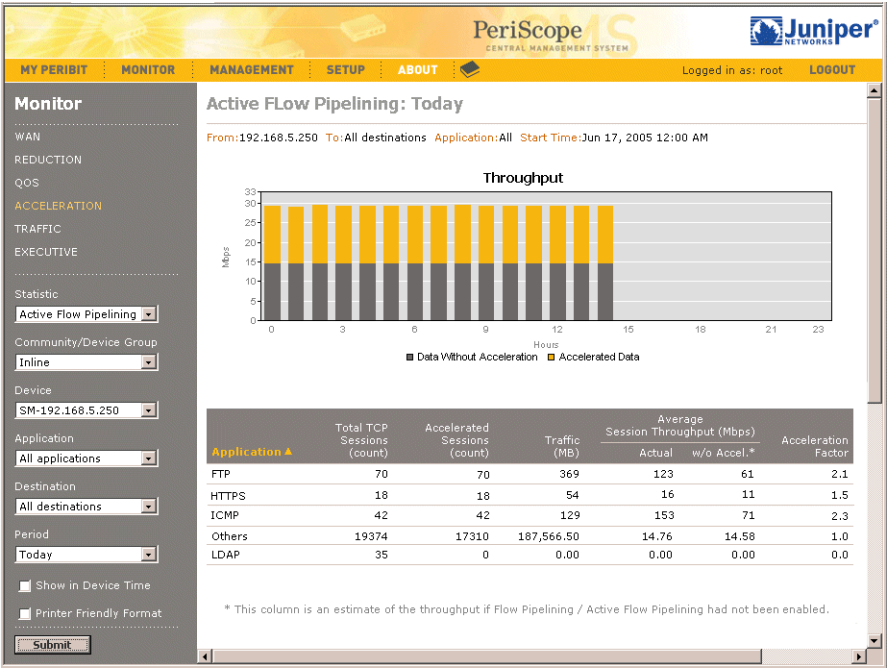


Figure 6-26 Flow Pipelining Statistics

Review the following information. Keep in mind that all values are for the selected application, destination, and time period.


- The Throughput bar graph shows the following:
 - **Data Without Acceleration** (grey bars). Average data throughput with no acceleration for applications that have Active Flow Pipelining enabled.
 - **Accelerated Data** (orange bars). Average increase in data throughput as a result of Active Flow Pipelining.
- The table has the following columns.
 - **Application** or **Destination**. Name of the accelerated application(s) or, if you select a specific application, the IP addresses of each remote device.
 - **Total TCP Sessions**. Number of sessions that ended in the selected time period.

- **Accelerated Sessions.** Number of accelerated sessions that ended in the selected time period.
 - **Traffic (MB).** Number of megabytes of traffic into the device that is accelerated.
 - **Average Session Throughput (Mbps).** Average throughput of all sessions, versus the estimated average throughput if Active Flow Pipelining was disabled.
 - **Acceleration Factor.** The performance increase for the accelerated sessions due to Active Flow Pipelining (actual throughput divided by the estimated throughput without acceleration). This value indicates the overall impact of Active Flow Pipelining.
4. Click **Export** to view or save the tabular data in CSV format.

Fast Connection Setup Statistics

If Fast Connection Setup is enabled for one or more endpoints and applications, the Fast Connection Setup report shows the session statistics and the average percentage reduction in session time due to Fast Connection Setup.

To view Fast Connection Setup statistics:

1. Click **MONITOR** in the menu frame, and **ACCELERATION** in the navigation frame.
2. Select **Fast Connection Setup** from the **Statistic** menu.
3. Select the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a device from the **Device** menu that has acceleration enabled. Devices using acceleration have a  on the Devices page (click **MANAGEMENT** in the menu frame to view the Devices page).
 - Select an application from the **Application** menu to view the acceleration statistics to each remote device. Select **Others** to view statistics for applications that are undefined or unmonitored. The default is **All applications**, which shows the average acceleration for all applications to all devices.
 - Select the IP address of a specific device from the **Destination** menu to view statistics only for traffic sent to the selected device. The default is **All destinations**.
 - Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

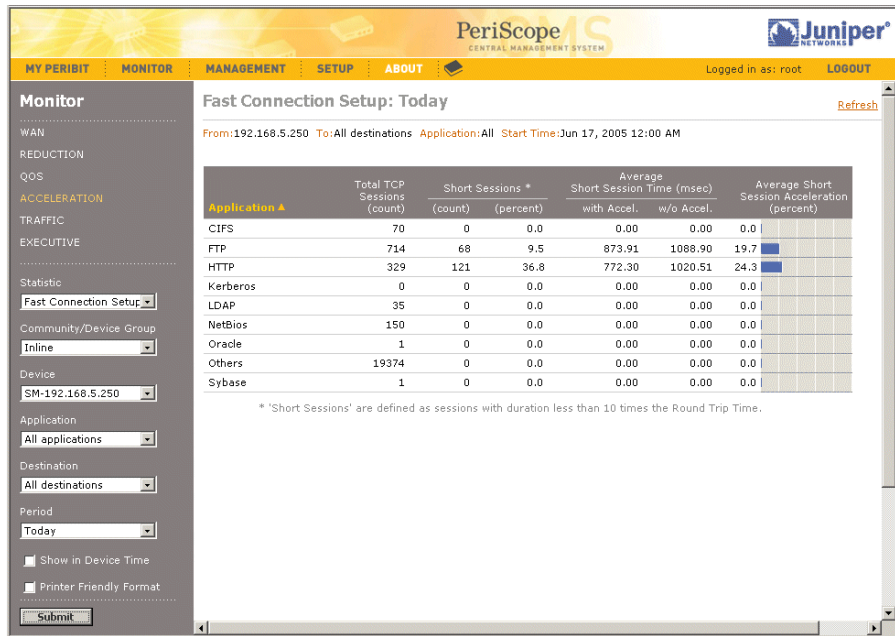


Figure 6-27 Fast Connection Setup Statistics

4. Review the following information. Keep in mind that all values are for the selected application, destination, and time period.
 - **Application or Destination.** Name of the accelerated application(s) or, if you select a specific application, the IP addresses of each remote device.
 - **Total TCP Sessions.** Number of sessions that ended in the selected time period.
 - **Short Sessions.** Number of “short” TCP sessions accelerated, and the percentage of the total sessions. These columns show the relative number of sessions that benefit from Fast Connection Setup. Short sessions are

those that last less than ten times the round-trip time (RTT). If a specific application traffic flow has five consecutive short sessions, subsequent identical traffic flows will be accelerated.

- **Average Short Session Time (msec).** Average duration of the accelerated sessions (in milliseconds), versus what the average session time would have been if Fast Connection Setup was disabled.
- **Average Short Session Acceleration (percent).** The average percentage reduction in session time, calculated as follows:

$$100 - [100 (\text{Accelerated session time}) / (\text{Session time without acceleration})]$$

This value indicates the overall impact of Fast Connection Setup on the accelerated sessions.

5. Click **Export** to view or save the data in CSV format.

CIFS and Exchange Acceleration Statistics

If CIFS or Exchange application acceleration is enabled for one or more application definitions on a device, the CIFS and Exchange acceleration reports shows the time saved due to CIFS and Exchange acceleration. Active Flow Pipelining must also be enabled.

NOTE: View CIFS and Exchange acceleration reports on the client-side device, not the server-side device. The acceleration statistics apply to the traffic in both directions. However, reduction statistics should probably be viewed on the server-side device.

To view CIFS or Exchange acceleration statistics:

1. Click **MONITOR** in the menu frame, and then click **ACCELERATION** in the left-hand navigation frame.
2. Select **CIFS Acceleration** or **Exchange Acceleration** from the **Statistic** menu.

3. Select the following report parameters, and click **Submit**.

- Select a community or device group from the **Community/Device Group** menu.
- Select a specific device from the **Device** menu.
- Select an application from the **Application** menu to view the acceleration statistics for a specific CIFS or Exchange application definition. The default is **All applications**.
- Select a specific device from the **Destination** menu to view statistics only for traffic sent to the selected device. The default is **All destinations**.
- Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

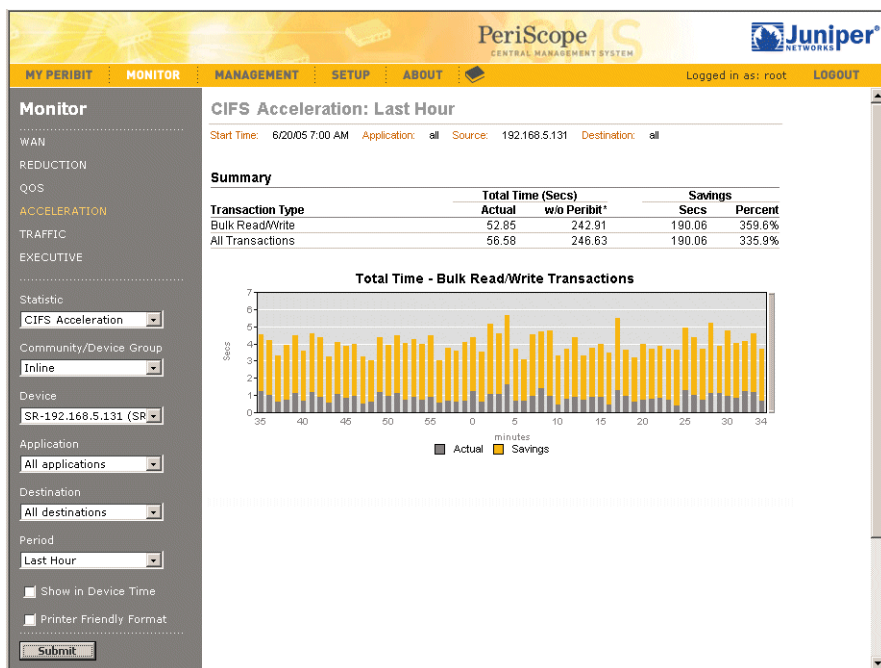


Figure 6-28 CIFS Acceleration Statistics

Review the following information. Keep in mind that all values are for the selected application, destination, and time period.

- The Summary table shows the following statistics for bulk read/write transactions and for all transactions.
 - **Total Time.** Number of seconds required to complete the transactions that ended in the selected time period for all clients, and the number of seconds that would have been required if acceleration was disabled.
 - **Savings.** Amount of time saved by acceleration, shown in seconds and as a percentage of the time required if acceleration was disabled.
- The two graphs show the following for bulk read/write transactions and for all transactions:
 - **Actual** (grey bars). Number of seconds required to complete the transactions that ended in the time period for all clients.
 - **Savings** (orange bars). Number of seconds saved by acceleration during the time period.

HTTP Acceleration Statistics

If HTTP acceleration is enabled for one or more application definitions, the HTTP Acceleration report shows the amount of time saved by HTTP acceleration. Active Flow Pipelining must also be enabled.

NOTE: View HTTP acceleration reports on the client-side device. The acceleration statistics apply to the traffic in both directions. However, reduction statistics should probably be viewed on the server-side device.

To view HTTP acceleration statistics:

1. Click **MONITOR** in the menu frame, and then click **ACCELERATION** in the left-hand navigation frame.
2. Select **HTTP Acceleration** from the **Statistic** menu.

3. Select the following report parameters, and click **Submit**.

- Select a community or device group from the **Community/Device Group** menu.
- Select a specific device from the **Device** menu.
- Select an application from the **Application** menu to view the acceleration statistics for a specific HTTP application definition. The default is **All applications**.
- Select a specific device from the **Destination** menu to view statistics only for traffic sent to the selected device. The default is **All destinations**.
- Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

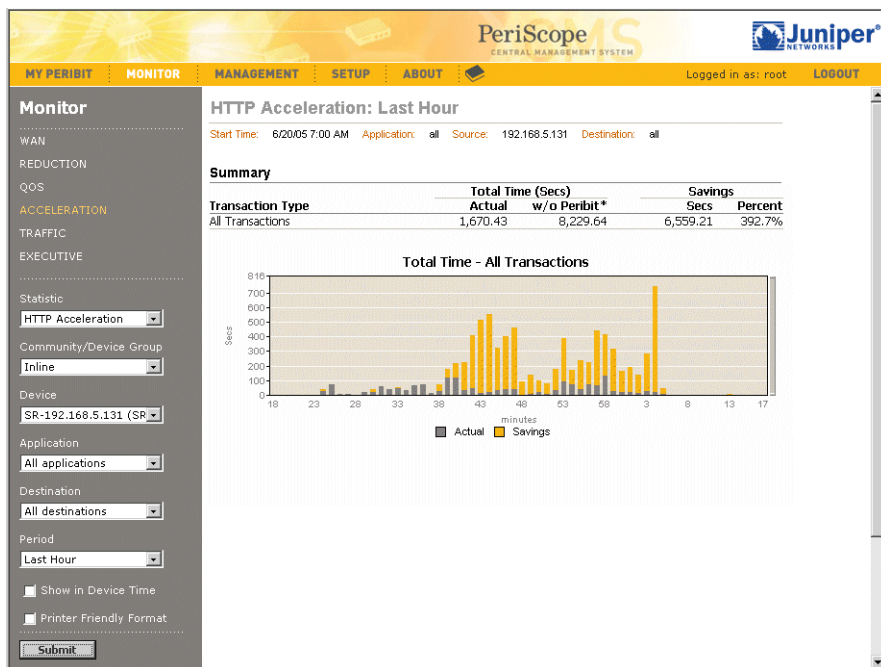


Figure 6-29 HTTP Acceleration Statistics

Review the following information. Keep in mind that all values are for the selected application, destination, and time period.

- The Summary table shows the following statistics for all transactions.
 - **Total Time.** Number of seconds required to complete the transactions that ended in the selected time period for all clients, and the number of seconds required if acceleration was disabled.
 - **Savings.** Amount of time saved by acceleration, shown in seconds and as a percentage of the time required if acceleration was disabled.
- The Total Time graph shows the following for all transactions:
 - **Actual** (grey bars). Number of seconds required to complete the transactions that ended in the time period for all clients.
 - **Savings** (orange bars). Number of seconds saved by acceleration during the time period.

Top Traffic Statistics

Each device collects statistics for its most active traffic flows, including the application name and protocol, source and destination addresses and ports, and the number of packets and bytes sent and received. The collected statistics can be sent to a Cisco NetFlow server and displayed in the Traffic report.

You can view the top traffic statistics for the past hour, the past 24 hours, or all available hours (the length of time depends on the traffic volume). The 65,000 most active flows are recorded. You can view the top 50 flows in the Web console, but the complete list can be exported to a file in CSV format.

NOTE: A flow constitutes data sent and/or received from a single source IP address and port number, to a single destination IP address and port number over the same protocol. Only the traffic flows that started in the selected time period are shown.

To view the Traffic statistics for a device:

1. Click **MONITOR** in the menu frame, and click **TRAFFIC** in the left-hand navigation frame.
2. Select a community or device group from the **Community/Device Group** menu, select a device from the **Device** menu, and click **Submit** to view the top traffic flows for the past hour. If an SRS 5.0 device is generating Cisco NetFlow records, you cannot view its traffic statistics in CMS.

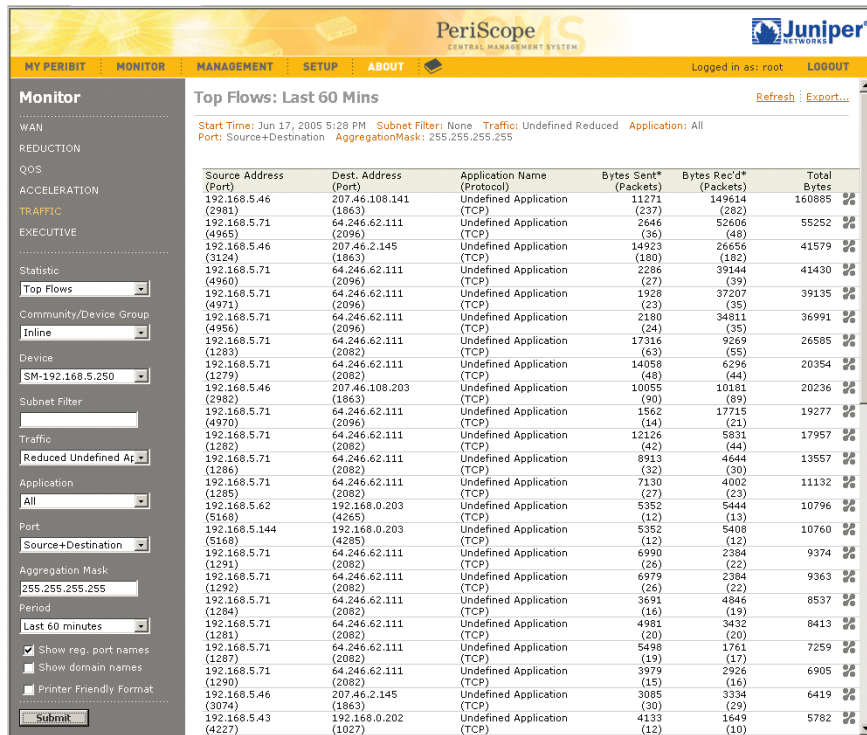



Figure 6-30 Top Traffic Statistics

Note that an  is shown next to the flows for undefined applications that are being reduced.

3. To filter the traffic statistics, specify the following information and click **Submit**.

Statistic	<p>Select a view of the traffic statistics. Each is displayed in descending order by traffic volume.</p> <ul style="list-style-type: none"> • Top Flows. The top 50 pairs of source and destination addresses and ports that have the highest total traffic (sent and received). Each traffic flow shows the number of bytes and packets sent and received by the source address. • Top Sending Addresses. Traffic sent by the top 50 addresses. • Top Sending Ports. Traffic sent by the top 50 ports. • Top Receiving Addresses. Traffic received by the top 50 addresses. • Top Receiving Ports. Traffic received by the top 50 ports.
Subnet Filter	<p>If you select the top flows, sending addresses, or receiving addresses, you can enter a subnet to view just the traffic from that subnet (SRS 5.1 devices only). The format is:</p> <p><IP address>/<subnet mask></p> <p>Where <subnet mask> is the number of bits used for the network portion of the address (such as “10.10.20.0/24”).</p>
Traffic	<p>Select a view of the traffic for the selected statistic.</p> <ul style="list-style-type: none"> • All. All traffic for the selected statistic. • All Reduced. Reduced traffic only. • Reduced Undefined Apps. Reduced traffic for undefined applications only. • Passthrough Only. Traffic sent from the WAN to the LAN that was not reduced. Does not apply to off-line devices or to in-line devices that use tunnel switching.
Application	<p>Select an application to view the traffic for a specific application.</p>

Port	<p>If you select the top flows, you can select a view of the port information.</p> <ul style="list-style-type: none"> • Ignore Port. Traffic is consolidated across all ports for each pair of source and destination addresses. • Source Only. Traffic is consolidated across the same source ports for each pair of source and destination addresses. • Destination Only. Traffic is consolidated across the same destination ports for each pair of source and destination addresses. • Source + Destination. Traffic is shown for each combination of source and destination port.
Aggregation mask	<p>If you select the top flows, sending addresses, or receiving addresses, you can enter a subnet mask to view all traffic from the same subnet as one consolidated entry. The default mask ("255.255.255.255") shows a separate flow for each host. You can also use the "/n" format. For example, enter "/24" or "255.255.255.0" to consolidate all traffic flows with the same 24-bit network address.</p>
Period	<p>Select the time period (last 60 minutes, last 24 hours, or all). Note that if you select Last 60 minutes or Last 24 hours, only the traffic flows that started in the selected time period are shown.</p>
Show reg. port names	<p>If you select the top flows, click the check box to view the registered names for all ports in the collected data. Clear the check box to view the names only for port numbers up to 1024.</p>
Show domain names	<p>If you select the top flows, click the check box to view the domain names for each IP address (SRS 5.1 devices only). The DNS server used is the one specified on the Windows machine where CMS is installed. The IP address is displayed if its domain name cannot be resolved (the DNS queries may take a few seconds).</p>

4. To export the traffic statistics to a file in CSV format, click **Export** in the upper-right corner of the page.

Executive Summary

The Executive report summarizes reduction results, traffic volume by application, and average WAN performance (latency and loss) for one or all remote devices.

To view the Executive statistics:

1. Click **MONITOR** in the menu frame, and click **EXECUTIVE** in the left-hand navigation frame.
2. Select the following report parameters, and click **Submit**.
 - Select a community or device group from the **Community/Device Group** menu.
 - Select a specific device from the **Device** menu. The default is **All devices**.
 - Select a time period from the **Period** menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

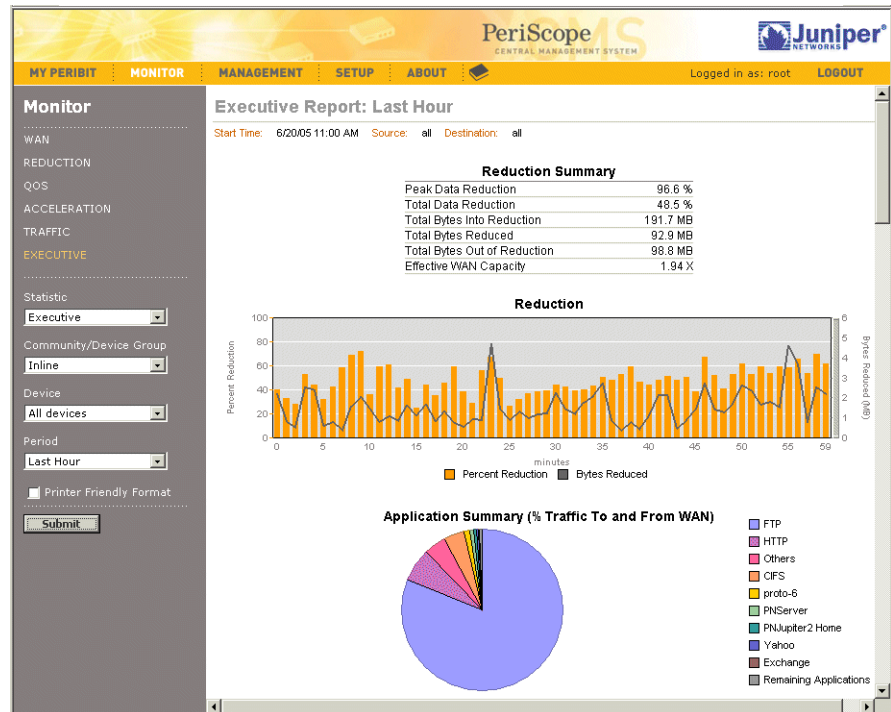


Figure 6-31 Executive Summary

3. Review the following information.

- The Reduction Summary table shows the following:
 - **Peak Data Reduction.** Highest percentage of data reduction for the selected time period. Based on five-second intervals for hourly reports, one-minute-intervals for daily reports, and one-hour intervals for weekly and monthly reports.
 - **Total Data Reduction.** Percentage of reduced data for the selected time period.
 - **Total Bytes Into Reduction.** Number of bytes into the data reduction engine.
 - **Total Bytes Reduced.** Number of bytes reduced.
 - **Total Bytes Out of Reduction.** Number of bytes of traffic output after data reduction.
 - **Effective WAN Capacity.** Factor increase in WAN capacity resulting from the total data reduction. For example, this value is 2.00 if total data reduction is 50%.
- The Reduction graph shows the average percentage of data reduction and the number of bytes reduced for the selected time period.
- The Application Summary pie chart shows the nine monitored applications with the highest percentage of the total traffic sent to and from the WAN for all destinations. The **Remaining applications** category shows the traffic for all other applications (both defined and undefined). Move the cursor over the legend to view the number of bytes for each application.
- The Volume by Application graph shows the traffic volume over the selected time period for the top nine monitored applications, plus the **Remaining applications** category.
- The Path Latency, Loss, and Availability distribution charts show the distribution of the average WAN latency, loss, and availability values measured for all destinations by devices that have WAN performance monitoring or Policy-Based Multipath enabled. To change the performance ranges represented by each color, refer to “Setting WAN Reporting Thresholds” on page 348.

Chapter 7 CMS Setup and Administration

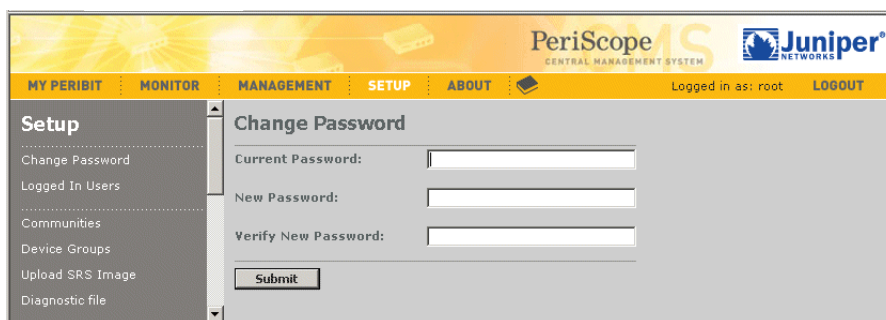
This chapter describes how to set up and administer CMS and covers the following topics:

- “Changing User Passwords” in the next section
- “Viewing Logged In Users” on page 326
- “Uploading a Boot Image” on page 327
- “Administering Devices” on page 328
- “Administering CMS” on page 334

Changing User Passwords

All CMS users can change their password, unless restricted by their user account (refer to “Defining CMS User Accounts” on page 335):

1. Click **SETUP** in the menu frame, and then click **Change Password** in the left-hand navigation frame.



The screenshot shows the PeriScope CMS web interface. The top navigation bar includes links for MY PERIBIT, MONITOR, MANAGEMENT, SETUP, and ABOUT. The user is logged in as 'root'. The left-hand navigation menu is expanded, showing options like Change Password, Logged In Users, Communities, Device Groups, Upload SRS Image, and Diagnostic file. The main content area is titled 'Change Password' and contains three input fields: 'Current Password:', 'New Password:', and 'Verify New Password:'. A 'Submit' button is located below the input fields.

Figure 7-1 Changing a User Password

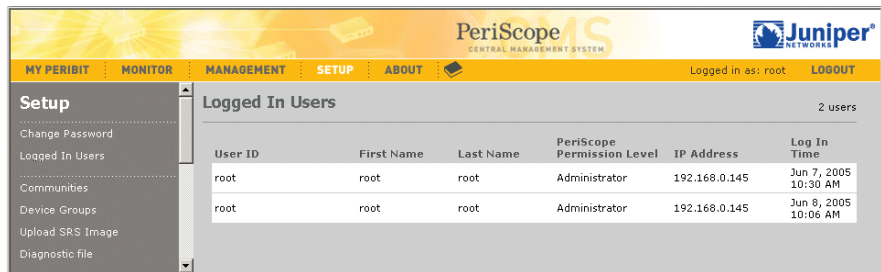
2. In the Change Password page, type the current password, and then type the new password in the **New Password** and **Verify New Password** fields.
3. Click **Submit** to activate the new password.

Viewing Logged In Users

Up to 50 users can access CMS at any given time. All users can view a list of the users who are currently logged in to CMS.

To view a list of the logged-in users:

1. Click **SETUP** in the menu frame, and then click **Logged In Users** in the left-hand navigation frame.



User ID	First Name	Last Name	PeriScope Permission Level	IP Address	Log In Time
root	root	root	Administrator	192.168.0.145	Jun 7, 2005 10:30 AM
root	root	root	Administrator	192.168.0.145	Jun 8, 2005 10:06 AM

Figure 7-2 Viewing Logged In Users

2. Review the following information for each user:
 - User ID and name.
 - Permission Level (Administrator, User, or Read-Only). For more information about user accounts and levels of access, refer to “Defining CMS User Accounts” on page 335.
 - IP address of the client that the user logged in from.
 - Date and time the user logged in.

NOTE: Users who close their Web browser without logging out are shown here until the session timeout expires.

Uploading a Boot Image

To load software upgrades on selected devices, you must first upload the boot image to CMS from a local disk or an FTP server. To distribute a boot image from the CMS server to selected devices in a community, refer to “Loading Device Boot Images” on page 38.

NOTE: You should retain the default name of the boot image to easily identify different releases and builds.

To upload a boot image to the CMS server:

1. Click **SETUP** in the menu frame, and then click **Upload SRS Image** in the left-hand navigation frame.

Figure 7-3 Uploading a Boot Image

2. Select **Local Disk** and click **Browse** to locate the boot image, or select **FTP Server** and specify the IP address of the FTP server, pathname and filename of the boot image, and the user name and password. If the FTP server accepts anonymous user access, leave the user name and password blank.

The boot image must have either a “.bin” or “.zip” extension. CMS does not recognize files with other extensions.

3. Click **Submit** to upload the boot image to the CMS server.

Administering Devices

The following sections describe the device-related administration tasks available to users with “PeriScope Administrator” privileges:

- “Managing Communities” on page 328
- “Managing Device Groups” on page 331
- “Generating a Diagnostic File” on page 333

Managing Communities

A community is a group of devices that can reduce and accelerate data for each other. Communities are defined on the devices that act as registration servers. When you install a device, you specify the IP address and password of a registration server that the device contacts periodically to identify the other devices in the same community.

To manage the devices in each community, CMS must import the communities defined on each registration server. Thereafter, the registration server is queried each day for changes to the imported communities. As new communities are added to a registration server, they must be imported into CMS.

Also, if you change the password on a registration server, you must update the password here. You can then download the new password to the devices in each community defined on the registration server (refer to “Applying a Registration Server Password” on page 58).

NOTE: Changes made here affect only CMS, not the registration server.

To import communities or change a registration server's password:

1. Click **SETUP** in the menu frame, and then click **Communities** in the left-hand navigation frame.



Figure 7-4 Community Administration

For each community, the Communities page lists the names of the primary and secondary registration servers and the number of devices in the community.

From the Communities page, you can:

- Import communities, as described in Step 2.
- Change a registration server password, as described in Step 3.
- Delete communities from CMS. Select the check box next to one or more communities, and click **Delete**. Note the following:
 - All schedule information is also deleted.
 - If the devices in a deleted community belong to no other communities, they are removed from all device groups (refer to “Managing Device Groups” on page 331).

2. To import the communities from a registration server:
 - a. Click **Import** to open the Communities > Import page.

Figure 7-5 Importing Communities from a Registration Server

- b. Specify the IP address and password of a registration server, and click **Submit**.
 - c. Select the check box next to each community you want to manage, click **Import**, and then click **OK**. Communities that have already been imported do not have a check box.
3. To update a registration server's password in CMS:
 - a. On the Communities page, click the registration server name.
 - b. Enter and verify the new password. The new password must match the password defined on the registration server. The previous password is retained and used as a security check when you apply the new password to the devices that access the registration server.

NOTE: If you make a mistake and must re-enter the password, enter and verify the previous password first, and then enter the new password. Applying the password from CMS will fail if the previous password is incorrect.

- c. Click **Submit** to activate the changes, or click **Cancel** to discard them.
 - d. Apply the new password to the devices in all communities defined on the registration server (refer to “Applying a Registration Server Password” on page 58).

Managing Device Groups

Each CMS user can be granted read or write access to one or more communities or device groups. Device groups let you create arbitrary groups of devices that are independent of communities. A device group can include a single device, or multiple devices from one or more communities. The same device can appear in multiple device groups.

NOTE: When viewing reports by device group, performance statistics between devices are shown only for devices in the same community.

To allow access to a device group:

- Add the device group to at least one user group, and specify read or write access (refer to “Defining User Groups” on page 337)
- Add the appropriate users to the user group (refer to “Defining CMS User Accounts” on page 335)

To define device groups:

1. Click **SETUP** in the menu frame, and then click **Device Groups** in the left-hand navigation frame.



Figure 7-6 Device Groups

From the Device Groups page, you can:

- Add new device groups, as described in Step 2.
- Change a device group. Select a group name, click **Add/Remove Devices**, change the devices in the group, and click **Submit**.

- Delete device groups. Select the check box next to one or more device groups, and click **Delete**. The deleted device groups are removed from all user groups, and users in those user groups may lose access to some or all of the devices in the deleted device groups.

2. To add a new device group:

- a. Click **New** to open the Device Groups > Add page.

Figure 7-7 Adding Device Groups

- b. Enter a device group name (up to 32 characters) and click **Add/Remove Devices**.
- c. Select a community from the **Community** list. The device name and IP address are shown for each device in the selected community. The IP address is enclosed in parentheses. To import more communities, refer to “Managing Communities” on page 328.
- d. Select the devices you want to add to the group, and click **Add**. To remove devices from the Members of Device Group list, select the devices and click **Remove**.
- e. Repeat Steps c and d for each community, as needed, and click **Submit**.

Generating a Diagnostic File

If you have problems with CMS, you can generate a diagnostic file to send to our support team. The diagnostic file contains current configuration, system information, and the most recent log files. By completing the form on the Diagnostic file page, your contact information is included with the file. After you generate and save the diagnostic file, email it to support@juniper.net.

NOTE: To generate a diagnostic file, the CMS Web server must be functioning.

To generate a diagnostic file:

1. Click **SETUP** in the menu frame, and then click **Diagnostic file** in the left-hand navigation frame.

Figure 7-8 Generating a Diagnostic File

2. Complete the form so that your contact information and a description of the problem is included with the diagnostic file.
3. Click **Submit** to generate the diagnostic file, and then click **Save** and specify a local file name and location.

Email the diagnostic file as an attachment to support@juniper.net. A support representative will contact you.

Administering CMS

The following sections describe the features available only to users with “Administrator” privileges:

- “Defining CMS User Accounts” in the next section
- “Defining User Groups” on page 337
- “Controlling Client Device Access to CMS” on page 339
- “Defining the Session Timeout” on page 340
- “Configuring FTP Server Parameters” on page 341
- “Enabling Syslog Reporting” on page 342
- “Entering a Permanent License Key” on page 343
- “Stopping and Starting the Scheduler” on page 344
- “Changing the Web Server Port” on page 345
- “Configuring Data Collection and Retention” on page 346
- “Setting WAN Reporting Thresholds” on page 348
- “Viewing the Polling Catch-Up and Failure Logs” on page 349
- “Viewing System Logs” on page 351
- “Backing Up and Restoring the Database” on page 352
- “Purging Temporary Java Files” on page 355

Defining CMS User Accounts

CMS provides a user account named “root” that has full CMS access. You can create up to 49 additional user accounts, each with one of the following access levels:

- **Administrator** — Users can perform all CMS tasks.
- **Read-Only** — Users can manage devices, but they cannot create or change configuration files.
- **User** — Users can manage devices and create and modify configuration files, but they cannot administer CMS.

User groups determine the devices that a user can access, and whether the access is read-only or read/write. A user cannot access any devices until the user’s account is assigned to at least one user group. To define user groups, refer to “Defining User Groups” on page 337.

To define CMS user accounts:

1. Click **SETUP** in the menu frame, and then click **Users** in the left-hand navigation frame.

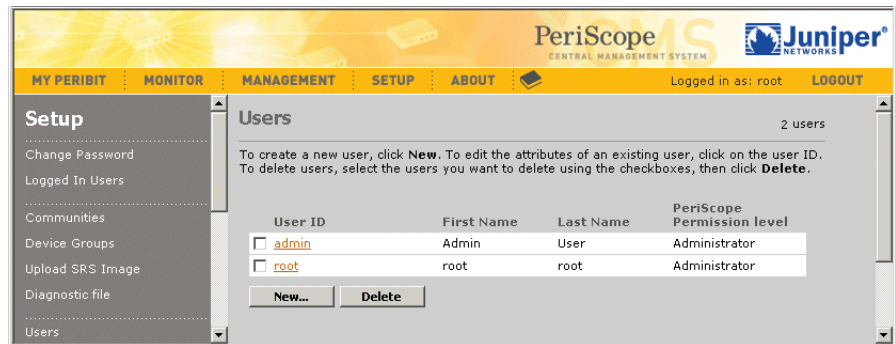


Figure 7-9 Administering CMS User Accounts

From the Users page, you can:

- Add new user accounts, as described in Step 2.
- Change a user account. Click the user ID, make any needed changes, and click **Submit**.
- Delete user accounts. Select the check box next to one or more accounts, and click **Delete**.

2. To add a new CMS user account, click **New**, specify the following information, and click **Submit**:

User ID	Enter the account login name (up to 32 characters). In general, use only letters and numbers when defining user IDs. If necessary, you can use the following special characters: : # \$ & _ - / () '
First Name	Enter the user's first and last name (up to 32 characters each). Both names are required.
Last Name	
Permission Level	Select the access level (administrator, user, or read-only).
Password	Enter the password twice (from 4 to 30 characters). Clear the check box to prevent the user from changing the password.
Member of	<p>Add the account to at least one user group by selecting the user group and clicking Add. Alternatively, you can add the account directly to the user group (refer to "Defining User Groups" on page 337).</p> <p>User groups provide read or read/write access to specific devices. Users cannot access any devices until their account is assigned to one or more user groups.</p> <p>NOTE: Since user groups allow access to devices by community or device group, access conflicts can occur (read vs. read/write) if a device belongs to multiple communities or device groups. These conflicts are resolved as follows:</p> <ul style="list-style-type: none"> • If the conflict occurs between two user groups that a user belongs to, the user receives read/write access. • If the conflict occurs within a single user group, the user receives read-only access.

Defining User Groups

User groups specify the devices that a user can access, and whether the access is read-only or read/write. A user cannot access any devices until the user's account is assigned to at least one user group. You can assign user accounts to user groups as described here, or you can select the user groups when you define the user account (refer to “Defining CMS User Accounts” on page 335).

To define user groups:

1. Click **SETUP** in the menu frame, and then click **User Groups** in the left-hand navigation frame.



Figure 7-10 Defining User Groups

From the User Groups page, you can:

- Add new user groups, as described in Step 2.
- Change a user group. Click the group name, make any needed changes, and click **Submit**.
- Delete user groups. Select the check box next to one or more user groups, and click **Delete**. The deleted user groups are removed from all user accounts, and users in those groups may lose access to some or all of the devices specified in the deleted user groups.

2. To add a new user group, click **New**, specify the following information, and click **Submit**:

User Group Name	Enter the group name (up to 32 characters).
Devices	<p>To add communities and/or device groups to the user group:</p> <ol style="list-style-type: none"> 1. Click Add/Remove Devices. 2. Select one or more communities or device groups and click Add. When the list is complete, click Submit. 3. Select the Read Only or Read/Write permission for each of the selected communities and device groups, and click Submit. <p>To import communities from a registration server, refer to “Managing Communities” on page 328. To add more device groups, refer to “Managing Device Groups” on page 331.</p>
Members	<p>To add user accounts to the user group:</p> <ol style="list-style-type: none"> 1. Click Add/Remove Users. 2. Select one or more accounts in the Users list and click Add. When the list is complete, click Submit. <p>NOTE: Since user groups allow access to devices by community or device group, access conflicts can occur (read vs. read/write) if a device belongs to multiple communities or device groups. These conflicts are resolved as follows:</p> <ul style="list-style-type: none"> • If the conflict occurs between two user groups that a user belongs to, the user receives read/write access. • If the conflict occurs within a single user group, the user receives read-only access.

Controlling Client Device Access to CMS

You can create an Include or Exclude list to allow or deny administrative access to CMS from specific IP addresses or subnets. For example, if you enter one address in the Include list, administrative users can log in only from the specified address. Alternatively, if you enter an address or subnet in the Exclude list, access from that address or subnet is denied.

By default, the Include and Exclude lists are empty, which means that administrative access is allowed from any address.

To restrict administrative access to CMS:

1. Click **SETUP** in the menu frame, and then click **PeriScope Access** in the left-hand navigation frame.

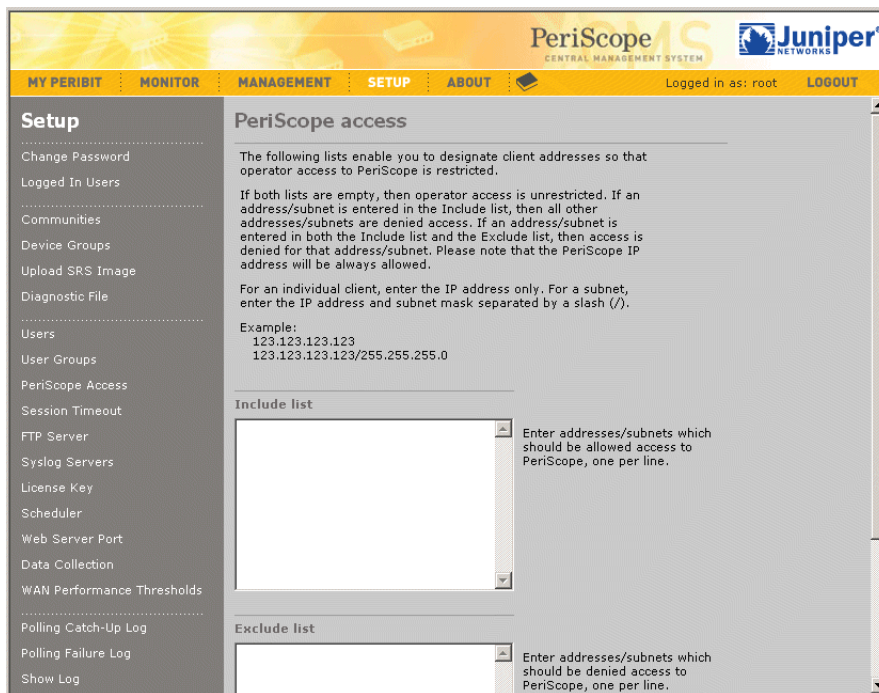


Figure 7-11 Controlling Client Device Access to CMS

2. To allow access to CMS only from specific IP addresses or subnets, enter the addresses or subnets in the **Include list** (one per line). The subnet format is:

<IP address>/<subnet mask>

All other client IP addresses are denied access to the device.

3. To deny access to CMS only from specific IP addresses or subnets, enter the addresses or subnets in the **Exclude list** (one per line).

NOTE: IP addresses that are in both the Include and Exclude lists are denied access.

4. Click **Submit** to activate the changes, or click **Reset** to discard them.

Defining the Session Timeout

The session timeout is the length of time a session can be idle before the session is closed (from 15 minutes to 24 hours). The default is 30 minutes.

To change the session timeout:

1. Click **SETUP** in the menu frame, and then click **Session Timeout** in the left-hand navigation frame.

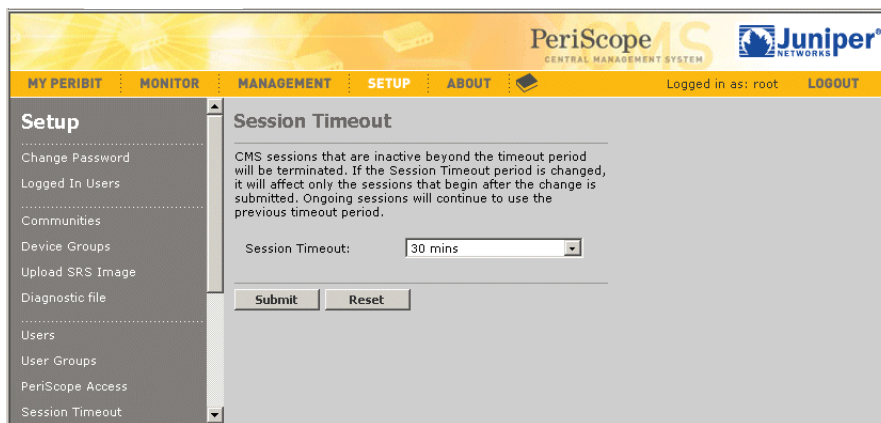


Figure 7-12 Defining the Session Timeout

2. Select a timeout, and click **Submit**.

The new timeout affects only future sessions, not current sessions.

Configuring FTP Server Parameters

The Microsoft FTP server must be installed and running on the CMS server. During Quick Setup, you specified the FTP user name, password, and root directory. FTP is used by CMS to upload boot images, and by the devices to send data to CMS. You can change the FTP parameters at any time.

To change the FTP server parameters:

1. Click **SETUP** in the menu frame, and then click **FTP Server** in the left-hand navigation frame.

The screenshot shows the PeriScope CMS interface. The top navigation bar includes 'MY PERIBIT', 'MONITOR', 'MANAGEMENT', 'SETUP', and 'ABOUT'. The 'SETUP' tab is active. On the left, a navigation menu lists various setup options, with 'FTP Server' selected at the bottom. The main content area is titled 'FTP Server' and contains three input fields: 'User name' with the value 'peribit', 'Password' with masked characters '****', and 'Root Directory' with the value 'c:\inetpub\ftproot'. Below these fields are 'Submit' and 'Reset' buttons. The top right of the interface shows the user is logged in as 'root' and provides a 'LOGOUT' link.

Figure 7-13 Configuring FTP Server Parameters

2. Specify the FTP user name and password. If the FTP server allows anonymous user access, enter “anonymous” in the **User name** field and leave the **Password** field blank.
3. Verify that the FTP root directory is correct.
4. Click **Submit** to activate the new settings. To restore the original settings, click **Reset**.

Enabling Syslog Reporting

CMS can send Syslog messages to up to five Syslog servers. A Syslog server allows you to centrally log and analyze configuration events and system error messages such as the failure of scheduled tasks. For a description of Syslog messages, refer to “Device Events” on page 359.

To enable Syslog reporting for CMS:

1. Click **SETUP** in the menu frame, and then click **Syslog Servers** in the left-hand navigation frame.

The screenshot shows the PeriScope CMS interface. The top navigation bar includes 'MY PERIBIT', 'MONITOR', 'MANAGEMENT', 'CMS SETUP', and 'ABOUT'. The left-hand navigation frame under 'CMS Setup' lists various options, with 'Syslog Servers' selected. The main content area, titled 'Syslog Servers', contains the following configuration options:

- Enable syslog reporting:** A checkbox labeled 'Yes' is checked.
- Syslog servers:** A text input field contains the IP address '192.168.70.15'. To the right, a note says 'Enter IP addresses, one per line'.
- Syslog message severity:** Three checkboxes are present: 'Critical' (checked), 'Error' (checked), and 'Information' (unchecked). To the right, a note says 'Check message severity levels you want reported to the syslog server.'

At the bottom of the configuration area are two buttons: 'Submit' and 'Reset'.

Figure 7-14 Enabling CMS Syslog Reporting

2. Select the **Enable syslog reporting** check box to enable Syslog reporting, and then enter the IP addresses of up to five Syslog servers (one per line).
3. Select the severity levels of the messages sent to the Syslog server:
 - **Critical:** Critical error messages, such as license exceeded.
 - **Error:** Error message, such as scheduled task failure.
 - **Informational:** Informational messages, such as a restart.
4. Click **Submit** to activate the changes, or click **Reset** to discard them.

Entering a Permanent License Key

CMS requires a permanent license key to operate beyond the 45-day evaluation period. The permanent license key determines the maximum number of devices that CMS manages. For more information about the permanent license, see “CMS Licenses” on page 357.

To enter a permanent license key for CMS:

1. Click **SETUP** in the menu frame, and then click **License Key** in the left-hand navigation frame.

The screenshot shows the PeriScope CMS web interface. The top navigation bar includes links for MY PERIBIT, MONITOR, MANAGEMENT, SETUP, and ABOUT. The left-hand navigation menu is expanded to the 'Setup' section, which includes options like Change Password, Logged In Users, Communities, Device Groups, Upload SRS Image, Diagnostic file, Users, User Groups, PeriScope Access, Session Timeout, FTP Server, Syslog Servers, and License Key. The main content area is titled 'License Key' and contains the following text: 'The maximum number of Peribit devices supported by CMS is determined by the license key. You can use CMS by entering an evaluation license key for the purpose of evaluation. However, the evaluation license will expire in 45 days. A valid license key must be entered prior to the expiration date in order to ensure uninterrupted service.' Below this text is a table showing the current license details:

License key	7CSA-F9SY-BJEZ-K9FY-ESAK-GA
Current license	200 devices
Expires	Jul 30, 2005 11:59 PM

Below the table, it states: 'The license can be upgraded by entering a new license key below.' There is a text input field labeled 'License Key' and two buttons, 'Submit' and 'Reset'.

Figure 7-15 Entering a License Key

2. Enter the permanent license key in the **License Key** field. Be sure to enter all characters, including hyphens (-), of the permanent license key.
3. Click **Submit** to activate the permanent license key. To restore the original license key, click **Reset**.

Stopping and Starting the Scheduler

The CMS scheduler lets you schedule device management tasks for a future date and time (refer to “Managing CMS Schedules” on page 62). If your network is having problems, and you have critical tasks scheduled for execution, you might want to stop the scheduler until the problems are resolved.

NOTE: You must reschedule any tasks that are scheduled to run while the scheduler is off.

To stop and start the scheduler:

1. Click **SETUP** in the menu frame, and then click **Scheduler** in the left-hand navigation frame.

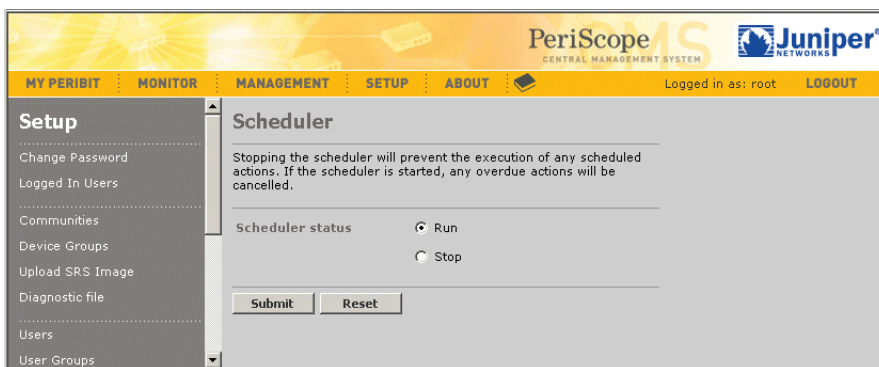


Figure 7-16 Stopping and Starting the Scheduler

2. Click **Stop** or **Run** to stop or start the scheduler., and click **Submit**.

Changing the Web Server Port

By default, the CMS Web server uses port 443 (HTTPS). You can change this port if necessary. If you change the port, CMS must be restarted for the new port to take effect.

To change the CMS Web server port:

1. Click **SETUP** in the menu frame, and then click **Web Server Port** in the left-hand navigation frame.

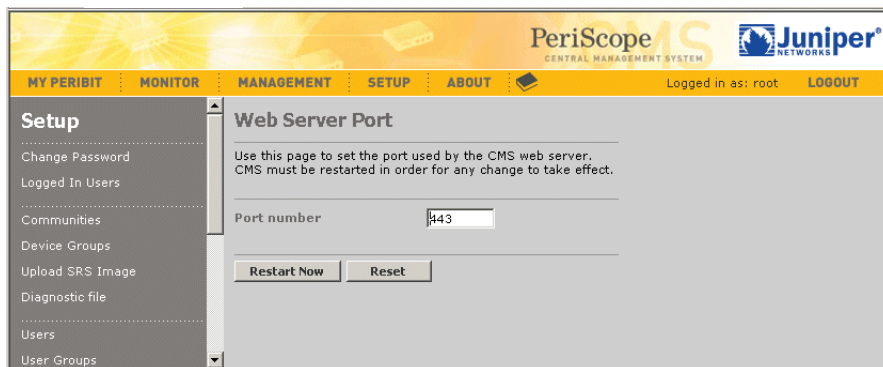


Figure 7-17 Changing the Web Server Port Number

2. Enter the port number in the **Port number** field.

You can specify any unused port, but if 443 is not used, 8443 is recommended.

3. Click **Restart Now** to restart CMS and activate the new port number. To restore the original port number, click **Reset**.

After a restart, the Web console is redirected to the new port number in about 60 seconds.

Configuring Data Collection and Retention

By default, performance data is collected every hour from SRS 5.1 devices, and every 30 minutes from SRS 5.0 devices. If you have a large number of devices, you may want to reduce the polling interval to once a day to conserve system resources. As needed, you can disable polling for one or all devices, and change the length of time that the collected data is retained.

The following table provides a rough estimate of the daily growth of the database (actual results depend on the device configurations). Polling once a day eliminates the per minute data.

Type of Data	Daily Disk Space per Device	Max Days Retention
Per minute	10 MB	10
Hourly	180 KB	180
Daily	15 KB	365

IMPORTANT: If you poll devices every hour, all devices should use an NTP server to ensure the accuracy of the hourly reports (refer to “Configuring NTP” on page 107).

To configure data collection and retention:

1. Click **SETUP** in the menu frame, and then click **Data Collection** in the left-hand navigation frame.

The screenshot shows the PeriScope CMS interface. The top navigation bar includes MY PERIBIT, MONITOR, MANAGEMENT, SETUP, and ABOUT. The left navigation pane shows a tree view with categories like Communities, Device Groups, Upload SRS Image, Diagnostic file, Users, User Groups, PeriScope Access, Session Timeout, FTP Server, Syslog Servers, License Key, Scheduler, Web Server Port, Data Collection, and WAN Performance Thresholds. The main content area is titled 'Data Collection' and contains the following configuration options:

- Polling Interval:** A dropdown menu set to '1 hour'.
- Polling Start Time:** A dropdown menu set to '1:00 AM'.
- Data Retention:** Three input fields with labels:
 - Keep minute data for: 7 days (max: 10 days)
 - Keep hourly data for: 30 days (max: 180 days)
 - Keep daily data for: 180 days (max: 365 days)
- Poll Checked Endpoints:** A table with two columns of IP addresses and checkboxes for each. The IP addresses are: 128.52.34.17, 128.54.34.17, 128.56.34.17, 129.52.34.17, 129.54.34.17, 129.56.34.17, 130.52.34.17, 130.54.34.17, and 130.56.34.17. Below the table are 'Select All' and 'Clear' buttons.
- Buttons:** 'Submit' and 'Reset' buttons at the bottom.

Figure 7-18 Configuring Data Collection and Retention

2. Specify the following information:

Polling Interval	<p>Select a polling interval:</p> <ul style="list-style-type: none"> • 1 hour. Collects data from SRS 5.1 devices for the previous hour (the default). For SRS 5.0 devices, data is collected every 30 minutes for the previous half hour. <p>If a 5.1 device does not respond to a poll, the next poll requests hourly data (rather than per-minute data) for the previous two hours. A “catch-up” poll can request up to the last 24 hours of data. To view the catch-up polls for each device, refer to “Viewing the Polling Catch-Up and Failure Logs” on page 349.</p> <ul style="list-style-type: none"> • 1 day. Collects data for the previous day. Per minute data is not retained, and the Last Hour and Today time periods will be greyed out on reports. • Never. Disables polling for all devices.
Polling Start Time	If you select a one-day polling interval, select a start time (off-peak hours are recommended)
Data Retention	<p>Enter the number of days to retain the collected data in the database. The options are:</p> <ul style="list-style-type: none"> • Per Minute. Up to 10 days (default is 7). • Hourly. Up to 180 days (default is 30). • Daily. Up to 365 days (default is 180).
Poll Checked Endpoints	Select the check box next to each device that you want to poll. To select all devices, click Select All . To deselect all devices, select Clear .

3. Click **Submit** to activate the changes, or click **Reset** to discard them.

Setting WAN Reporting Thresholds

The WAN Performance, Loss, Latency, and Availability charts use colored cells in a matrix to indicate the WAN status between each pair of devices in a community (refer to “WAN Statistics” on page 274 and “Executive Summary” on page 321). As needed, you can change the default performance ranges associated with each of the colors.

To change the WAN reporting thresholds:

1. Click **SETUP** in the menu frame, and then click **WAN Performance Thresholds** in the left-hand navigation frame.

PeriScope CENTRAL MANAGEMENT SYSTEM **Juniper** NETWORKS

MY PERIBIT MONITOR MANAGEMENT **SETUP** ABOUT

Logged in as: root LOGOUT

Setup

- Change Password
- Logged In Users
- Communities
- Device Groups
- Upload SRS Image
- Diagnostic File
- Users
- User Groups
- PeriScope Access
- Session Timeout
- FTP Server
- Syslog Servers
- License Key
- Scheduler
- Web Server Port
- Data Collection
- WAN Performance Thresholds**
- Polling Catch-Up Log
- Polling Failure Log
- Show Log

WAN Performance Thresholds

The table below allows you to modify the thresholds which determine the display of colored dots in the WAN performance reports.

	% Time Above Latency Threshold		% Loss		% Availability	
Excellent	<	0.5	<	0.25	>	99.5
Good	0.5	to 1.0	0.25	to 0.3	99.0	to 99.5
Marginal	1.0	to 5.0	0.3	to 0.5	98.0	to 99.0
Warning	5.0	to 10.0	0.5	to 2.0	95.0	to 98.0
Critical	>	10.0	>	2.0	<	95.0

The WAN Performance report provides an aggregate health metric that combines latency, loss and availability. You can select the manner in which the aggregate measure is determined from the three parameters for each path in the WAN.

☐ Display the best of the parameters
☐ Display the worst of the parameters
☒ Display the average of the parameters

Figure 7-19 Setting WAN Reporting Thresholds

2. To change the performance ranges for a color, enter the new percentage values for time above the latency threshold, loss, and availability.
3. Select whether the WAN Performance report reflects the best, average, or worst values measured for loss, latency, and availability.
4. Click **Submit** to activate the changes, or click **Reset** to discard them. To restore the default values, click **Set to Defaults**.

Viewing the Polling Catch-Up and Failure Logs

If an SRS 5.1 device does not respond to an hourly poll for performance data, the next poll requests hourly data (rather than per-minute data) for the previous two hours. These catch-up polls can request up to the last 24 hours of data. If a device does not respond for an entire day, the performance data for that day will be lost.

Note that successful catch-up polls create discrepancies between the Last Hour and daily reports for a device. For example, the Today report will have data for the last hour, but the Last Hour report will be blank (no per-minute data). You can verify the cause of these discrepancies by viewing the polling catch-up and failure logs.

To view the polling catch-up or failure log:

1. Click **SETUP** in the menu frame, and then click **Polling Catch-Up Log** or **Polling Failure Log** in the left-hand navigation frame.

The polling catch-up log (Figure 7-20) shows the time of each catch-up poll for each device.

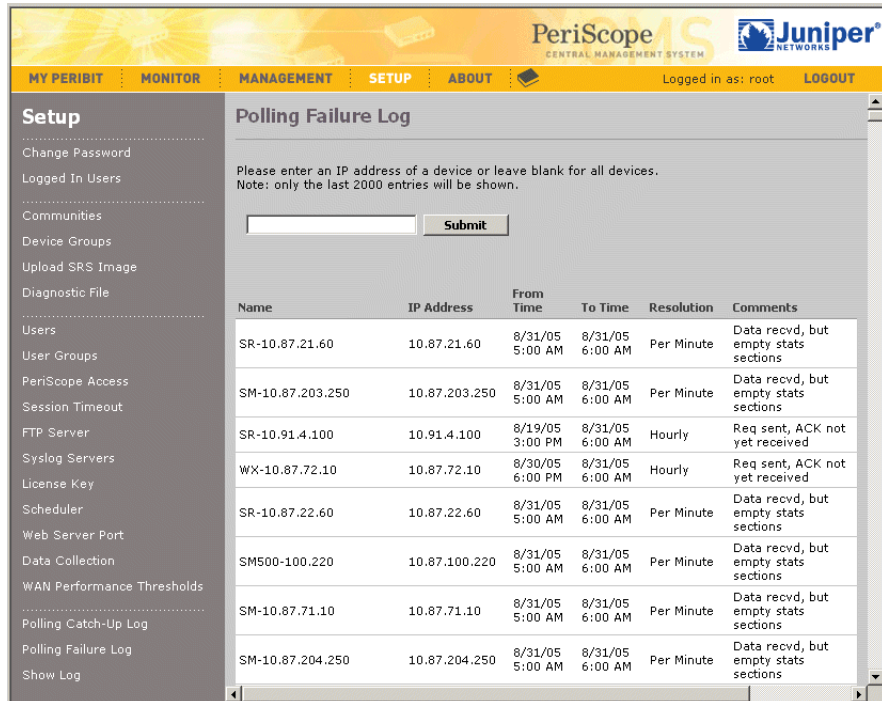
The screenshot shows the PeriScope CMS interface. The top navigation bar includes links for MY PERISBIT, MONITOR, MANAGEMENT, SETUP (selected), and ABOUT. A user is logged in as 'root'. The left-hand navigation menu under 'Setup' includes options like Change Password, Logged In Users, Communities, Device Groups, Upload SRS Image, Diagnostic File, Users, User Groups, PeriScope Access, Session Timeout, FTP Server, Syslog Servers, License Key, Scheduler, Web Server Port, Data Collection, WAN Performance Thresholds, Polling Catch-Up Log (selected), Polling Failure Log, and Show Log.

The main content area is titled 'Polling Catchup Log'. It features a search input field with the IP address '192.168.5.131' and a 'Submit' button. Below the search field is a table displaying the polling catch-up log entries.

Name	IP Address	From Time	To Time	Catchup Type
SR-192.168.5.131 (SR-50A)	192.168.5.131	6/4/05 1:00 AM	6/4/05 3:00 AM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/24/05 10:00 PM	5/25/05 12:00 AM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 5:00 PM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 4:00 PM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 3:00 PM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 2:00 PM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 1:00 PM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 12:00 PM	Hourly instead of per minute
SR-192.168.5.131 (SR-50A)	192.168.5.131	5/18/05 9:00 AM	5/18/05 11:00 AM	Hourly instead of per minute

Figure 7-20 Viewing the Polling Catch-Up Log

The polling failure log (Figure 7-21) shows the time of each failed poll for each device, whether the poll was for per minute or hourly data (a normal or catch-up poll), and whether the device failed to respond or responded with no statistics.



PeriScope CMS
CENTRAL MANAGEMENT SYSTEM

Juniper
NETWORKS

MY PERIBIT MONITOR MANAGEMENT **SETUP** ABOUT

Logged in as: root LOGOUT

Setup

- Change Password
- Logged In Users
- Communities
- Device Groups
- Upload SRS Image
- Diagnostic File
- Users
- User Groups
- PeriScope Access
- Session Timeout
- FTP Server
- Syslog Servers
- License Key
- Scheduler
- Web Server Port
- Data Collection
- WAN Performance Thresholds
- Polling Catch-Up Log
- Polling Failure Log**
- Show Log

Polling Failure Log

Please enter an IP address of a device or leave blank for all devices.
Note: only the last 2000 entries will be shown.

Submit

Name	IP Address	From Time	To Time	Resolution	Comments
SR-10.87.21.60	10.87.21.60	8/31/05 5:00 AM	8/31/05 6:00 AM	Per Minute	Data recvd, but empty stats sections
SM-10.87.203.250	10.87.203.250	8/31/05 5:00 AM	8/31/05 6:00 AM	Per Minute	Data recvd, but empty stats sections
SR-10.91.4.100	10.91.4.100	8/19/05 3:00 PM	8/31/05 6:00 AM	Hourly	Req sent, ACK not yet received
WX-10.87.72.10	10.87.72.10	8/30/05 6:00 PM	8/31/05 6:00 AM	Hourly	Req sent, ACK not yet received
SR-10.87.22.60	10.87.22.60	8/31/05 5:00 AM	8/31/05 6:00 AM	Per Minute	Data recvd, but empty stats sections
SM500-100.220	10.87.100.220	8/31/05 5:00 AM	8/31/05 6:00 AM	Per Minute	Data recvd, but empty stats sections
SM-10.87.71.10	10.87.71.10	8/31/05 5:00 AM	8/31/05 6:00 AM	Per Minute	Data recvd, but empty stats sections
SM-10.87.204.250	10.87.204.250	8/31/05 5:00 AM	8/31/05 6:00 AM	Per Minute	Data recvd, but empty stats sections

Figure 7-21 Viewing the Polling Failure Log

- To view the log entries for a specific device, enter the IP address and click **Submit**. To view all entries again, delete the IP address and click **Submit**. The logs retain the last 2000 entries.

Viewing System Logs

The CMS system log files can be displayed in the Web console. If your network has dedicated Syslog servers, you can send CMS log messages to up to five Syslog servers, as described in “Enabling Syslog Reporting” on page 342.

To view the system logs:

1. Click **SETUP** in the menu frame, and then click **Show Log** in the left-hand navigation frame.

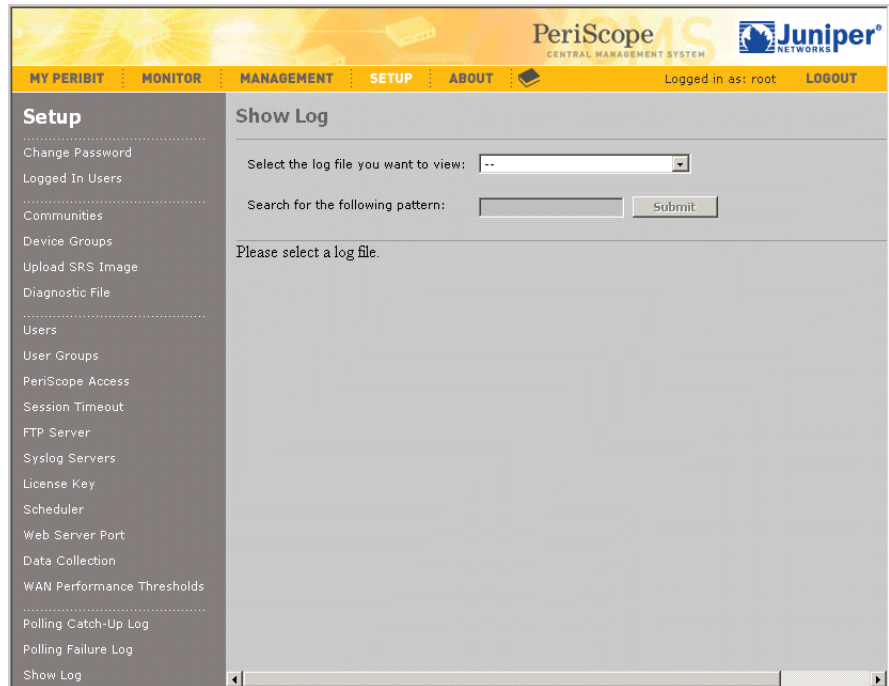


Figure 7-22 Viewing the System Logs

2. Select the log you want to view.
 - **access.log.** Lists the time, user name, and client address of each user who has logged in to CMS since the last reboot (also lists the times of failed logins).
 - **activity.log.** Lists the various types of CMS tasks performed by each user since the last reboot.

- **debug.log**. Lists the most recent informational messages, such as aggregation start and stop times, along with error messages, such as failures to communicate with a device. Each time **debug.log** reaches 5 MB, it is saved as **debug.log.1**, and existing logs are renumbered up to **debug.log.5** (older logs are discarded).
3. To view all entries that have a specific text string, enter the text in the **Pattern** field (such as a partial IP address) and click **Submit**. To view all entries again, delete the text string and click **Submit**.

Backing Up and Restoring the Database

You should periodically back up the database on the CMS server. You can either stop the server and back up the database manually, or use a script to back up the database automatically while the server is running.

Manual Database Backups

To backup the CMS database manually:

1. Stop the CMS service and the MySQL service.
 - a. Click **Start > Run**, enter “services.msc” and click **OK**.
 - b. In the Services window, right-click on **PeriScopeCMS** and click **Stop**, and then right-click on **MySQL** and click **Stop**.
2. Check the current size of the database in the following folder:


```
<Install>\MySQL\data
```

Where “<Install>” is the location where the database is installed. The default location is “C:\Program Files\Peribit\CMS”.
3. Copy the `\data` folder to a backup location that has sufficient disk space.
4. Restart the **MySQL** and **PeriScopeCMS** services.

5. To restore a database that was backed up manually:
 - a. Stop the **PeriScopeCMS** and **MySql** services.
 - b. Rename the `\MySql\data` folder.
 - c. Copy the `\data` folder from the backup location to the `\MySql` folder.
 - d. Restart the **MySql** and **PeriScopeCMS** services.
 - e. Delete the old `\data` folder that was renamed in Step b.

Automatic Database Backups

Database backup and restore scripts are provided that let you back up the database while the server is running. Some data may not be backed up, so use this process only if you cannot stop the server.

By default, the backup script (*dbbackup.bat*) copies the database to the current drive. To back up the database on another drive, change the last line of the script to specify the appropriate drive. To restore a database from another drive, you must also change the last line of the restore script (*dbrestore.bat*).

To back up the CMS database while the server is running:

1. Check the current size of the database in the following folder:

```
<Install>\MySql\data
```

Where “<Install>” is the location where the database is installed. The default location is “C:\Program Files\Peribit\CMS”.

2. Verify that the free space remaining on the current drive is at least twice the size of the `\data` folder.
3. To execute the backup script, open a Command Prompt window and enter the following command (or use the Windows “at” command to execute the script daily or weekly):

```
<Install>\MySql\pnscripsts\dbbackup.bat <Install>
```

NOTE: For the default installation location, enter “C:\Program Files\Peribit\CMS” as “C:\Program~1\Peribit\CMS”. Spaces are not supported in the path name.

The following files are created in the `\MySql\backup` folder:

- **CMSData.sql**. New database backup.

- **CMSData_1.sql**. Previous database backup.

If an error occurs during the backup, copy the **CMSData_1.sql** file to **CMSData.sql**, and execute the backup script again.

4. To restore a database that was backed up automatically:
 - a. Stop the CMS service (the MySQL service must be running).
 - b. Verify that you have a **CMSData.sql** file in the `\MySQL\backup` folder. If the last backup failed, be sure to copy the **CMSData_1.sql** file to **CMSData.sql**.
 - c. Open a Command Prompt window and enter the following command:

```
<Install>\MySQL\pnscripts\dbrestore.bat <Install>
```

NOTE: For the default installation location, enter "C:\Program Files\Peribit\CMS" as "C:\Program~1\Peribit\CMS". Spaces are not supported in the path name.

The current database is dropped, and a new database is created with the data imported from the backup database.

Moving CMS to Another Disk Drive

To move CMS to another disk drive:

1. Copy the CMS database to a temporary folder on the new drive, as described in “Manual Database Backups” on page 352, but do not restart the CMS nor MySQL services.
2. Uninstall CMS using the Microsoft Windows Add/Remove Programs function in the Control Panel. Elect NOT to remove the CMS database. This is not an option if you have a temporary CMS license.
3. Install CMS on the new drive (refer to “Installing CMS” on page 21).
4. Restore the CMS database that you saved in Step 1 (refer to “Manual Database Backups” on page 352).
5. When you are satisfied that CMS is running properly on the new drive, remove the old CMS installation directory.

Purging Temporary Java Files

In JRE version 1.5.0, Java cache files are accumulated in the Windows temporary folder on the CMS server. CMS will not start if the disk becomes full, so you should periodically delete the following files:

/WINNT/Temp/jar_cache.tmp*

Appendix A CMS Licenses

The CMS license determines the maximum number of devices that CMS can manage. CMS has two types of licenses:

- Permanent license—If you have purchased CMS, you should receive an authorization code via email. To use the authorization code to obtain a permanent license key, call 1-888-314-JTAC (or 1-408-745-9500), or go to https://www.juniper.net/generate_license.
- Evaluation license—If you are evaluating CMS, you can use the evaluation license generated during the installation of the CMS software. The evaluation license is valid for 45 days and allows CMS to manage up to 5 devices. If you do not enter a permanent license key before the 45-day evaluation period expires, on the 46th day you will lose access to all CMS Web console pages, except the License Key page.

If you use an evaluation license during installation and subsequently purchase CMS and receive a permanent license, reinstallation is not necessary. Simply enter the permanent license key, as described in “Entering a Permanent License Key” on page 343.

If you want to manage a larger number of devices than is specified in your permanent license, you can purchase an upgrade license. Again, simply enter the new license key. Reinstallation is not necessary.

If you upgrade your CMS software to a newer version of software, the permanent license key last applied to CMS is retained and honored. To download the latest version of CMS, go to <https://www.juniper.net/customers/csc/software/>.

Appendix B Device Events

Table B-1 lists the Syslog messages generated by CMS, and Table B-2 list the critical- and error-level Syslog messages generated by the devices and displayed in the Devices page as “events.” For information about how to access information about the device events, see “Viewing Device Events” on page 37.

Table B-1 CMS Events

Message	Warning: Exceeded licensed number of devices. Current licensed: <licensed#> and total devices: <actual#>.
Message Type	Critical Error
Recommended Action	Call +1-866-737-4248 to obtain a new license to support all of your devices.
Message	The schedule task <task name> for community (or device group) <dg name> failed.
Message Type	Error
Recommended Action	To reschedule or acknowledge failed tasks, see “Managing Scheduled Tasks” on page 62
Message	The CMS was started.
Message Type	Information
Recommended Action	None

Table B-2 Device Events

Message	Safe-mode suspend: case 2
Message Type	Critical Error
Recommended Action	Contact technical support.
Message	Exceeded licensed throughput
Message Type	Error
Recommended Action	This message is also sent if you explicitly reboot the system into Safe Mode from the Web console or the Command Line Interface (CLI). Call +1-866-737-4248 to obtain a new license with speed configured to a higher value.

Table B-2 Device Events (Continued)

Message	License expired, Data reduction/assembly has been disabled
Message Type	Error
Recommended Action	Call +1-866-737-4248 to obtain a new license.
Message	REG: Self registration failed. IP=<ip address>.
Message Type	Error
Recommended Action	Check the network connectivity to the primary registration server <ip address>.
Message	REG: Self registration failed for secondary registration server. IP=<ip address>.
Message Type	Error
Recommended Action	Check the network connectivity to the secondary registration server <ip address>.
Message	REG: Registration failed. Password mismatch. IP=<ip address>
Message Type	Error
Recommended Action	The device <ip address> does not have the correct registration server password. It can be corrected from CLI or Web UI.
Message	Health monitor detected anomalous system condition
Message Type	Error
Recommended Action	The health monitoring system detected an unexpected error condition. The health monitoring system will take corrective action and attempt to restore proper operating condition, including if necessary performing a system reset. Please call +1-866-737-4248 to further analyze the anomaly.
Message	SaveConfig: Cannot save <module> settings: status=<status>
Message Type	Error
Recommended Action	Call +1-866-737-4248 with this information.
Message	Login failed: access=<method> user=<name> IP=<ip-addr>
Message Type	Error
Recommended Action	The message has the access method (CONSOLE, SSH, or IWEB) and the IP address of the client (for SSH and WEB). You can check if the user is authorized to configure this system. Since CONSOLE access requires physical access to the system, any unauthorized CONSOLE access should be treated as a serious problem.
Message	Fan Error (CPU or Chassis fan not operational).
Message Type	Error

Table B-2 Device Events (Continued)

Recommended Action	CPU or Chassis fan may not be working. May have to replace the fan in the system if the message persists.
Message	Fan Speed Error (Cpu or Chassis speed variation).
Message Type	Error
Recommended Action	This message indicates that change in fan speed was noticed. May have to replace the fan in the system if the message persists.
Message	SR: Multi-Node Master Node is Down
Message Type	Error
Recommended Action	This message indicates that the master node of the multi-node configuration is down. If this node has not been taken down intentionally, please check the running configuration and the network connectivity for problems.
Message	SR: Multi-Node Last Node is Down
Message Type	Error
Recommended Action	This message indicates that the last node of the multi-node configuration is down. If this node has not been taken down intentionally, please check the running configuration and the network connectivity for problems.
Message	<p>WAN performance messages:</p> <ul style="list-style-type: none"> • WP: ***** Unacceptable Performance detected due to LOSS on Path, Path Ip=<remote IP address> ***** • WP: ***** Unacceptable Performance detected due to LATENCY on Path, Path Ip=<remote IP address> ***** • WP: ***** Acceptable Performance detected on Path, Path Ip=<remote IP address> *****
Message Type	Informational
Recommended Action	These messages indicate changes in when WAN performance between the local device and the specified remote WX device. No action is required.

Appendix C Understanding Exported Data Results

This appendix describes the top traffic data that can be exported to Cisco NetFlow servers, and the monitoring statistics that CMS can retrieve in CSV format from devices running SRS 5.1. For information about retrieving statistical data, see “Retrieving Device Files” on page 56.

This appendix covers the following sections:

- “NetFlow Version 5 Export” in the next section
- “Performance Statistics Export” on page 365
- “Top Traffic Export” on page 376

NetFlow Version 5 Export

Each device can export its top traffic data to a Cisco NetFlow server in Version 5 format (refer to “Generating NetFlow Records” on page 115).

Table C-1 describes the NetFlow packet header.

Table C-1 NetFlow Packet Header

Byte	Parameter	Description
0-1	Version	NetFlow export format version number (5).
2-3	Count	Number of flows exported in this packet (1 to 30).
4-7	Sysuptime	Number of milliseconds since the device was restarted.
8-11	Unix seconds	Number of seconds since 0000 1970 Coordinated Universal Time (UTC).
12-15	Unix nanoseconds	Residual nanoseconds since 0000 1970 UTC.
16-19	Flow number	Sequence counter of total flows seen.
20	Engine type	Not applicable.
21	Engine ID	Not applicable.
22-23	Sampling interval	Not applicable.

Table C-2 describes each traffic flow entry in a NetFlow packet (up to 30 entries per packet).

Table C-2 NetFlow Packet Entry

Byte	Parameter	Description
0-3	Srcaddr	Source IP address.
4-7	Dstaddr	Destination IP address.
8-11	Nexthop	Not applicable.
12-13	Input	SNMP index number of input interface.
14-15	Output	SNMP index number of output interface.
16-19	Packets	Number of packets in the flow.
20-23	Octets	Number of Layer 3 bytes in the flow.
24-27	First	SysUptime at start of flow.
28-31	Last	SysUptime when the last packet in the flow was received.
32-33	Source port	TCP/UDP source port number or equivalent.
34-35	Destination port	TCP/UDP destination port number or equivalent.
36	Pad1	Unused (zero).
37	TCP flags	Cumulative OR of TCP flags.
38	Protocol	IP protocol number (for example, TCP = 6; UDP = 17).
39	ToS	IP type of service.
40-41	Source system	Not applicable.
42-43	Destination system	Not applicable.
44	Source mask	Not applicable.
45	Destination mask	Not applicable.
46-47	Pad2	Unused (zero).

Performance Statistics Export

The following sections describe the performance data that can be extracted in CSV format from each device in the *AllData.csv* file.

- “General Device Information” in the next section
- “Data Section Information” on page 366
- “System Session Statistics” on page 367
- “Reduction Session Statistics” on page 370
- “Application Session Statistics” on page 371
- “WAN Statistics” on page 372
- “Application Flow Acceleration Statistics” on page 372
- “Bandwidth Management Statistics” on page 373
- “WAN Performance Statistics” on page 374
- “Inbound Traffic By Port Statistics” on page 375

General Device Information

Table C-3 describes the general device information.

Table C-3 General Device Information

Parameter	Description
Device IP	IP address of the device.
Software version	Version of SRS software that was running when the statistics were exported.
Serial number	Serial number of the device that exported the statistics.
License speed	Licensed speed of the device.
Monitor applications	Names of the applications being monitored.
Fast Connection applications	Names of the applications using Fast Connection Setup.
Active Flow Pipelining applications	Names of the applications using Active Flow Pipelining.
Prime time enabled	Indicates whether prime time is enabled (Y or N).
Prime time hours	Hours of the day when prime time starts and ends (in 24-hour format).
Prime time days	Days of the week included in prime time.
Operation mode	Indicates whether the device is active (Inline) or in Profile mode.

Data Section Information

Table C-4 describes the data section information that precedes the set of statistic tables for each time range.

Table C-4 Data Section Information

Parameter	Description
<time> data section	Indicates the time range for the statistics tables that follow: <ul style="list-style-type: none"> • This hour • Last hour • Today • Yesterday • This week • Last week
ip=	IP address of the device.
device local time=	Local date and time of the export.

Table C-4 Data Section Information (Continued)

gmt time=	Date and time of the export in Greenwich Mean Time (GMT).
peak interval=5	Peak statistics are calculated over 5 second intervals.

System Session Statistics

Table C-5 describes the system session statistics.

Table C-5 System Session Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Bytes Into AE	Number of bytes that entered the Assembly Engine.
Bytes Out AE	Number of bytes out of the Assembly Engine after assembly.
Packets Into AE	Number of packets into the Assembly Engine.
Packets Out AE	Number of packets out of the Assembly Engine after assembly.
Resvd 1	Reserved
Bytes Out OOB	Number of out-of-band bytes sent to the control channel.
Bytes PT NO AE	Number of bytes that passed through without reduction due to no remote Assembly Engine.
Packets PT NO AE	Number of packets that passed through without reduction due to no remote Assembly Engine.
Bytes PT By Filter	Number of bytes that passed through without reduction due to a manually configured filter (such as an application filter).
Packets PT By Filter	Number of packets that passed through without reduction due to a manually configured filter (such as an application filter).
OfPt Bytes (Overflow Pass-through)	Number of bytes that passed through without reduction due to device buffer overflow.
OfPt Packets (Overflow Pass-through)	Number of packets that passed through without reduction due to device buffer overflow.
Bytes PT NO SR	Number of bytes that passed through without reduction due to a disabled reduction engine on this device.
Packets PT NO SR	Number of packets that passed through without reduction due to a disabled reduction engine on this device.

Table C-5 System Session Statistics (Continued)

Bytes PT NON-IP	Number of non-IP bytes that passed through without reduction (e.g., IPX, etc.).
Packets PT NON-IP	Number of non-IP packets that passed through without reduction (e.g., IPX, etc.).
Bytes PT IP-Other	Number of IP bytes that passed through without reduction because the protocols were not configured for reduction.
Packets PT IP-Other	Number of IP packets that passed through without reduction because the protocols were not configured for reduction.
Bytes PT SR	Number of bytes that passed through without reduction because the source address is the address of another device in the same community.
Packets PT SR	Number of packets that passed through without reduction because the source address is the address of another device in the same community.
Bytes PT SR-Hash	Number of bytes that passed through without reduction because the device is part of a reduction cluster and the data will be processed by another device.
Packets PT SR-Hash	Number of packets that passed through without reduction because the device is part of a reduction cluster and the data will be processed by another device.
Bytes PT IpFrag	Number of bytes that passed through without reduction because the device is not enabled to reduce IP fragments.
Packets PT IpFrag	Number of packets that passed through without reduction because the device is not enabled to reduce IP fragments.
Bytes PT License	Number of bytes that passed through without reduction because the throughput level exceeded the device's license.
Packets PT License	Number of packets that passed through without reduction because the throughput level exceeded the device's license.
Bytes PT Tunneled Only	Number of bytes that passed through without reduction.
Packets PT Tunneled Only	Number of packets that passed through without reduction.
Bytes PT VLAN	Number of bytes of VLAN traffic that passed through without reduction.
Packets PT VLAN	Number of packets of VLAN traffic that passed through without reduction.
Bytes PT L2Mcast	Number of Layer 2 Multicast bytes that passed through the device.
Packets PT L2Mcast	Number of Layer 2 Multicast packets that passed through the device.
TP Bytes In (throughput)	Number of bytes into the Reduction Engine at the peak five-second interval of data input ¹ .

Table C-5 System Session Statistics (Continued)

TP Bytes Out (throughput)	Number of bytes out of the Reduction Engine at the peak five-second interval of data input.
TP Bytes PT (throughput)	Number of bytes that passed through at the peak five-second interval of data input.
TP Packets In (throughput)	Number of packets into the Reduction Engine at the peak five-second interval of data input.
TP Packets Out (throughput)	Number of packets out of the Reduction Engine at the peak five-second interval of data input.
TP Packets PT (throughput)	Number of packets that passed through at the peak five-second interval of data input.
Resvd 2	Reserved
Resvd 3	Reserved
Peak % Rdn	Maximum data reduction rate for any five second interval within the selected time period. Peak percentage reduction is calculated by the following formula: $10^5 \times \left(\frac{\text{Bytes In} - \text{Bytes Out}}{\text{Bytes In}} \right) = \text{Peak \% Reduction}$
Rsv H1 through Rsv H20	Reserved
PkIn1 to PkIn6	Six fields that show the number of packets in each of six packet-size ranges for traffic into the device, as follows: <ul style="list-style-type: none"> • PkIn1 Less than 64 bytes • PkIn2 64 to 127 • PkIn3 128 to 255 • PkIn4 256 to 511 • PkIn5 512 to 1023 • PkIn6 More than 1023 bytes
PkOut1 to PkOut6	Six fields that show the number of packets in each of six packet-size ranges for traffic out of the device, as follows: <ul style="list-style-type: none"> • PkOut1 Less than 64 bytes • PkOut2 64 to 127 • PkOut3 128 to 255 • PkOut4 256 to 511 • PkOut5 512 to 1023 • PkOut6 More than 1023 bytes

T. Data input is the number of IP bytes into the device from the Local port.

Reduction Session Statistics

Table C-6 describes the reduction session statistics.

Table C-6 Reduction Session Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Dst Ip (Destination IP Address)	IP address of the remote device that receives reduced and/or encrypted data from this device.
Packets In	Number of packets into this reduction engine that were intended for the destination IP address.
Packets Out	Number of reduced packets sent to the destination IP address.
Packets Into Ipsec	Number of packets that were identified for encryption and intended for the destination IP address.
Packets Out of Ipsec	Number of encrypted packets sent to the destination IP address.
Packets Dropped by Ipsec	Number of packets intended for the destination IP address that were dropped according to the default IPSec policy.
Ipsec Overhead	Number of bytes added by IPSec processing.

Application Session Statistics

Table C-7 describes the application session statistics.

Table C-7 Application Session Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
App Id	Application from which the data was received (e.g., FTP, HTTP, Lotus Notes).
Dst Ip	IP address of the device that receives reduced data from this device.
Bytes In	Number of bytes into the device that have been identified for reduction, and addressed for the device listed with the destination IP address and application ID.
Bytes Out	Number of bytes out of this device after reduction, and addressed for the device listed with the destination IP address and application ID.
Acc Bytes In	Number of bytes eligible for Active Flow Pipelining.
Est Boost Bytes	Estimated number of bytes accelerated by Active Flow Pipelining.
Active Session time	Number of milliseconds during which data was sent for all Active Flow Pipelining sessions that ended during this time period.
Session Count	Number of all sessions that ended during this time period.
Avg % FC Speedup	Sum of the average percentages of time saved for each session by Fast Connection Setup. To get the average session speedup time shown on the Acceleration report, divide this value by the number of sessions, and then divide by 100.
FP Session Count	Number of Active Flow Pipelining sessions that ended during this time period.
FC Session Count	Number of Fast Connection Setup sessions that ended during this time period.
FC Session Time	Number of milliseconds for all Fast Connection Setup sessions that ended during this time period.
Bytes Out NSM	Number of bytes out of this device after reduction using NSM (Sequence Mirror devices only), and addressed for the device listed with the destination IP address and application ID.

WAN Statistics

Table C-8 describes the WAN statistics.

Table C-8 WAN Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
App Id	Application for which the data was sent or received.
App Type	Type of application (0=Default, 1=HTTP, 3=CIFS, 4=Exchange).
Dst Ip	IP address of the remote device that sent or received data from this device.
Bytes From WAN	Number of bytes received from the WAN for the remote device and application.
Bytes To WAN	Number of bytes sent to the WAN for the remote device and application.

Application Flow Acceleration Statistics

Table C-9 describes the Application Flow Acceleration statistics.

Table C-9 Acceleration Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
App Id	Application for which the traffic was accelerated (HTTP, CIFS, or Exchange).
App Type	Type of application (0=Default, 1=HTTP, 3=CIFS, 4=Exchange).
Tran Id	Transaction ID number (0=All, 1=Bulk read/write).
Dst Ip	IP address of the remote device that received the accelerated traffic.
Time With Accel	Number of seconds required to complete the transaction.
Time Without Accel	Estimated number of seconds required to complete the transaction with no acceleration.

Bandwidth Management Statistics

Table C-10 describes the bandwidth management statistics collected per application class for each reduction tunnel.

Table C-10 Bandwidth Management Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Tunnel	Outbound bandwidth management: The IP address of the destination assembler or the default allocation.
	Inbound bandwidth management: The parameter is Inbound.
Class	Outbound bandwidth management: The bandwidth class ID, which is a collection of applications that a user has mapped to the class.
	Inbound bandwidth management: One of the four pre-defined classes (i.e., Reduced, Intranet, TCP or Default).
Bytes In	Outbound bandwidth management: The total number of application bytes into the device.
	Inbound bandwidth management: The total number of bytes into the Remote interface of the device by class.
Bytes Out	Outbound bandwidth management: The total number of application bytes out of outbound bandwidth management.
	Inbound bandwidth management: the total number of bytes out of inbound bandwidth management.
Bytes Dropped	Outbound bandwidth management: The total number of application bytes dropped by the bandwidth management feature.
	Inbound bandwidth management: The total number of bytes dropped by the bandwidth management feature.
Packets In	Outbound bandwidth management: The total number of application packets into the device.
	Inbound bandwidth management: The total number of packets passed into the device by inbound bandwidth management.

Table C-10 Bandwidth Management Statistics

Packets Out	<p>Outbound bandwidth management: The total number of application packets transmitted by the device. (The total number does not include meta packetization.)</p> <p>Inbound bandwidth management: The total number of packets out of inbound bandwidth management.</p>
Packets Dropped	<p>Outbound bandwidth management: The total number of application packets dropped by the bandwidth management feature.</p> <p>Inbound bandwidth management: The total number of packets dropped by the bandwidth management feature.</p>

WAN Performance Statistics

Table C-11 describes the WAN performance statistics.

Table C-11 WAN Performance Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Tunnel	IP address of a remote device.
Avg Latency	Average round-trip time to the remote device (in milliseconds). For hourly data, the median value is shown for each minute.
Latency Count	Number of minutes for which a latency value was measured.
Latency Above Thresh	Average percentage of minutes that the latency threshold was exceeded. For hourly data, the value is 0 or 1 for each minute (1=above threshold).
Latency Above Thresh Count	Number of minutes for which the median latency exceeded the latency threshold.
Loss Pct	Average percentage of the probes that were lost.
Loss Count	Number of minutes for which a loss value was measured (excludes minutes for which none of the probes were returned).
Event Count	Number of times the loss or latency thresholds were exceeded or returned to normal.
Diversion Count	Number of times traffic was diverted to the alternate path (Multi-Path only).
Return Count	Number of times traffic was diverted back to the preferred path (Multi-Path only).
Last Down	Not used.

Table C-11 WAN Performance Statistics

Unavailable Count	Number of minutes for which none of the probes were returned.
Minute Count	Number of minutes for which performance monitoring was enabled.

Inbound Traffic By Port Statistics

Table C-12 describes the Inbound traffic by port statistics.

Table C-12 Inbound Traffic By Port Data

Parameter	Description
Src Port	Inbound data's source port number.
Bytes In	Number of reduced bytes from the source port for unmonitored applications.
Packets In	Number of reduced packets from the source port for unmonitored applications.
Dst Port	Inbound data's destination port number.
Bytes In	Number of reduced bytes to the destination port for unmonitored applications.
Packets In	Number of reduced packets to the destination port for unmonitored applications.

Top Traffic Export

Table C-13 describes the Traffic statistics retrieved in the *ip-flow.csv* file.

Table C-13 Inbound Traffic By Port Data

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Other Data	Number of bytes and packets sent and received for flows that exceeded the maximum retained by the device (16K for SR-15, 32K for SR-20, 65K for other models).
SrcIp	IP address of the flow source.
DstIp	IP address of the flow destination.
SrcPort	Source port number.
DstPort	Destination port number.
Proto	Traffic flow protocol (TCP, UDP, or protocol number).
Application	Traffic flow application name.
BytesSent	Number of bytes sent by the source.
PktsSent	Number of packets sent by the source.
BytesRcvdt	Number of bytes received by the source.
PktsRcvd	Number of packets received by the source.
TotalSendDelay	Cumulative delay between packets sent (in milliseconds).
TotalRcvDelay	Cumulative delay between packets received (in milliseconds).
Type	Indicates the traffic type: <ul style="list-style-type: none"> • RA. Reduced application • RO. Reduced undefined application • PT. Passed through due to policy setting • U. Unknown passthrough traffic, such as non-TCP/UDP traffic
StartTime	Start date and time of traffic flow.
EndTime	End date and time of traffic flow.

Appendix D Common Application Port Numbers

Table D-1 lists common application port numbers, as listed by the Internet Assigned Numbers Authority (IANA, <http://www.iana.org/assignments/port-numbers>).

NOTE: Port numbers 3577 and 3578 for TCP and UDP are reserved for data transmission.

Table D-1 Common Application Port Numbers

Keyword	Port Number	Protocol	Description
ftp-data	20	TCP/UDP	File Transfer [Default Data]
ftp	21	TCP/UDP	File Transfer [Control]
ssh	22	TCP/UDP	Secure Shell Protocol
telnet	23	TCP/UDP	Telnet
smtp	25	TCP/UDP	Simple Mail Transfer
dns	53	TCP/UDP	Domain Name Server
tftp	69	TCP/UDP	Trivial File Transfer
www-http	80	TCP/UDP	World Wide Web HTTP
kerberos	88	TCP/UDP	Kerberos
pop3	110	TCP/UDP	Post Office Protocol - Version 3
sunrpc	111	TCP/UDP	SUN Remote Procedure Call
nntp	119	TCP/UDP	Network News Transfer Protocol
netbios-ns	137	TCP/UDP	NETBIOS Name Service
netbios-dgm	138	TCP/UDP	NETBIOS Datagram Service
netbios-ssn	139	TCP/UDP	NETBIOS Session Service
imap2	143	TCP/UDP	Interim Mail Access Protocol v2
snmp	161	TCP/UDP	SNMP
snmptrap	162	TCP/UDP	SNMPTRAP
clearcase	371	TCP/UDP	Clearcase
legent-1	373	TCP/UDP	Legent Corporation
legent-2	374	TCP/UDP	Legent Corporation
ldap	389	TCP/UDP	Lightweight Directory Access Protocol

Table D-1 Common Application Port Numbers (Continued)

https	443	TCP/UDP	https MCom
netnews	532	TCP/UDP	readnews
lotusnotes	1352	TCP/UDP	Lotus Notes
ms-sql-s	1433	TCP/UDP	Microsoft-SQL-Server
ms-sql-m	1434	TCP/UDP	Microsoft-SQL-Monitor
watcom-sql	1498	TCP/UDP	Watcom-SQL
orasrv	1525	TCP/UDP	Oracle
ccmail	3264	TCP/UDP	cc:mail/lotus

Glossary

access control list	List of IP addresses from which an administrator can log in to CMS.
assembly	Process by which a device reassembles reduced traffic into its original form.
auto-negotiation	A protocol that enables Ethernet systems at the end of a twisted-pair or optical fiber segment to negotiate configuration parameters such as speed, half or full-duplex mode, and use of flow control.
bandwidth	The amount of data that can be sent through a network connection, measured in bits per second (bps).
bridge	A device that partitions a network into separate segments. The bridge allows a packet to be transmitted from one segment to the other only if it is addressed to a host on the other segment.
endpoint	A device in a community or a virtual device used to apply outbound QoS policies to specific remote subnets.
filter	An operator defined IP address or TCP port number that determines valid addresses or applications for reduction processing. A single filter or a list of filters can be defined for each system.
full-duplex	A mode of operation that enables a pair of systems connected by a link to transmit frames to one another at the same time.
gateway	A device that connects and forwards packets between computers or different networks. See also, <i>router</i> .
half-duplex	A mode of operation that allows only a single station to successfully transmit a frame at a given time.
hardware passthrough	Hardware-driven process by which all traffic is passed through the device at wire-speed. It is invoked automatically upon disruption.
HTTP	HyperText Transfer Protocol. The protocol most often used to transfer information from World Wide Web servers to browsers.
ICMP	Internet Control Message Protocol. An Internet Protocol used to communicate between devices on a network to manage errors and generate control messages.
Interior Gateway Protocol (IGP)	A group of protocols that provide routing information to the routers within an autonomous network.
Internet Protocol (IP)	The protocol that is used to route a data packet from its source to its destination over the Internet.

IP address	A numeric address, such as 10.10.124.22, assigned to every device on the network.
IP subnet mask	A numeric address, such as 255.255.0.0, used to define an IP subnet or to determine membership of an IP address in an IP subnet.
IP subnet	A group of IP addresses defined by the IP address and IP subnet mask pair, such as 10.10.0.0/255.255.0.0.
latency	The time necessary for a packet of data to travel from a source to a destination across a network.
log	A record of CMS activity. Logs are recorded for system information, performance, backup, and recovery.
MIB	Management Information Base. A database containing ongoing configuration information and statistics of a device in a network. MIBs are used with SNMP.
MTU	Maximum Transmission Unit. The largest size packet that can be transmitted by a device on a network.
OSPF	Open Shortest Path First. An interior gateway protocol that routes messages according to the least expensive path.
packet	A unit of data formatted for transmission on a network. Data is broken down into packets for sending over a packet switched network. Each packet has a header containing its source, destination, other control information, and a payload of data to be transmitted.
passthrough mode	A function of devices where all traffic passes through at wire-speed due to device disruption or overflow.
community	Two or more devices that can reduce and assemble data for each other. Initially, all devices belong to the Default community. Each device contacts the registration server to identify the other devices in the same community.
ping	A program 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.
reduction subnets	Subnets that a device can advertise to the other devices in the community. The other devices can then reduce traffic destined for those subnets.
registration server	The device that stores the network information for all devices in each community. Each device periodically contacts the registration server to identify the other devices in the same community.
response time	The time it takes for a host to respond to a user command.

RIP	See <i>Routing Information Protocol</i> .
round-trip time (RTT)	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.
router	Specialized computer that forwards data packets between networks. Routers can exchange information about their network connectivity (or accessibility) with neighboring network routes using standard routing protocols. This information is used by the router to determine an optimal path for a packet being forwarded.
Routing Information Protocol (RIP)	An interior gateway protocol used in IP networks.
Simple Network Management Protocol (SNMP)	The Internet standard protocol for network management software.
Simple Network Time Protocol (SNTP)	A protocol that can synchronize clocks on local computers with radio or atomic clocks on the Internet.
software passthrough	Software-driven process by which a device transparently passes packets through the system in lieu of processing (reducing).
static IP address	A permanent IP address for a client, server, or other network device.
Switch	A networking device that sends packets directly to a port associated with a given network address.
TCP	Transmission Control Protocol. The most common Internet transport layer protocol, defined in RFC 793. TCP is connection-oriented and stream-oriented, and provides for reliable communication over packet-switched networks.
tunneling	Encapsulating one type of packet inside the data field of another packet.
User Datagram Protocol (UDP)	User Datagram Protocol. UDP is connectionless and does not guarantee reliable communication; the application itself must process any errors and check for reliable delivery. Defined in RFC 768.
warm reboot	A reboot of the device without powering off the unit.

Index

Numerics

3DES encryption for IPsec 250

802.1q VLAN support 105

A

AAA settings 116

acceleration

Active Flow Pipelining 201, 203

CIFS 205

Exchange 208

Fast Connection Setup 201, 204

feature/topology setting 215

Forward Error Correction 201

HTTP 210

access levels, CMS 335

access lists 339

access log 351

acknowledging failed tasks 65

Active Flow Pipelining

configuring 201, 203

feature/topology setting 215

report 308

active FTP 134

activity log 351

advertising reduction subnets 94

AES encryption for IPsec 250

aggregate WAN speed

about 160

defining inbound QoS 189

defining outbound QoS 169, 180

analyzing device configurations 45

Application Flow Acceleration

about 193

CIFS and Exchange reports 313

CIFS traffic 205

Exchange traffic 208

feature/topology setting 215

HTTP reports 315

HTTP traffic 210

applications

accelerating

Active Flow Pipelining 203

CIFS traffic 205

Exchange traffic 208

Fast Connection Setup 204

HTTP traffic 210

assigning to traffic classes 136

common port numbers 377

managing 127

monitoring 138

monitoring percent reduction 288

reducing 144

summary statistics

all traffic 291

WAN traffic 282

visibility in tunnels 154

ARP, configuring 93

assemblers

default 149

preferred 151

authentication methods, selecting 117

Authorization Codes

importing 262

matching with devices 264

auto-deployment

about 253

of configurations and software 254

of device licenses 261

status 259

B

backing up

device configurations 52

the database 352

balancing, load

across reduction tunnels 147

across routers 112

- bandwidth management
 - inbound 188
 - outbound 155
- bidirectional view, percentage reduction 288, 290
- boot images
 - downgrading 38
 - rolling back 40
 - upgrading 38
 - uploading to CMS 327
- browser support, CMS 17
- bypass condition, multi-path 239

C

- caching, for HTTP acceleration 196
- cancelling pending tasks 64
- carving out addresses
 - and outbound QoS 166
 - from off-path RIP advertisements 221
- catch-up log, polling 349
- CIFS acceleration
 - about 194
 - configuring 205
 - feature/topology setting 215
 - reports 313
- circuit speeds
 - and router overhead 159
 - configuring 170, 181
- Citrix names, in application definitions 135
- classes, traffic
 - in the QoS Wizard 173
 - inbound QoS 188
 - outbound QoS and Multi-Path 136
- CLI commands, appending 234
- client access, CMS 339
- CMS Web console
 - about 19
 - browser support 17
 - logging in and out 18, 27
 - user accounts 325, 335
 - viewing logged in users 326
- CMS Web server port
 - changing 345
 - default 22, 25
- communities, managing 328

- Configuration window 83
- configuration, initial 27
- configurations
 - about 67
 - analyzing 45
 - backing up 52
 - changing 83
 - CLI commands, appending 234
 - comparing 80
 - deleting 82
 - displaying 81
 - generating
 - creating 78
 - duplicating 77
 - extracting from devices 47, 75
 - management recommendations 73
 - previewing before loading 50, 257
 - restoring 54
 - rolling back 51
 - summaries by device 43
 - verifying running configurations 43
 - version tracking 72
 - viewing history 82
- Configurations page 73
- congestion control 172, 183
- consistency checking 72
- conventions, document 13

D

- data collection and retention 346
- data packets, Forward Error Correction 202
- data reduction statistics 322
- data reduction, monitoring 283
- database
 - backups 352
 - growth estimates 346
- dead-time interval, RADIUS 122
- debug log 351
- dedicated WANs 160
- default gateway, configuring 88
- Default traffic class
 - inbound QoS 188
 - outbound QoS and Multi-Path 136
- deployment groups 255
- deployment records 257

- device groups, managing 331
- device names 89
- device time, viewing reports in 269
- devices
 - about 33
 - accessing the Web console 61
 - adding to scheduled tasks 65
 - cancelling pending tasks 64
 - events
 - list of 359
 - viewing 37
 - exporting information 61
 - failed tasks
 - icon on Devices page 36
 - rescheduling and acknowledging 65
 - groups of 331
 - icons 35
 - monitoring
 - inbound QoS 304
 - outbound QoS 300
 - percentage reduction 283
 - tunnel status 296
 - WAN performance 274
 - polling 346
 - rebooting 42
 - safe mode 60
 - supported by CMS 18
 - verifying running configurations 43
 - viewing 33
- Devices page 33
- DHCP and auto-deployment 253
- diagnostic files
 - CMS 333
 - retrieving from devices 56
- DNS servers
 - and auto-deployment 253
 - configuring for the CMS Traffic report 320
- DNS servers, configuring for the device Traffic report 89
- domain names in the Traffic report 89, 320
- DSCP values, see "ToS/DSCP values"
- duplicating configurations 77
- dynamic routes
 - OSPF and RIP 110
 - router polling 99

E

- encryption, see "IPSec"
- endpoints
 - IPSec 247
 - Multi-Path 240
 - NSM 141
 - outbound QoS 180
 - packet flow acceleration 198
 - reduction 139
 - WAN performance monitoring 231
- events
 - enabling CMS Syslog 342
 - list of device 359
 - viewing device 37
- Exchange acceleration
 - about 194
 - configuring 208
 - feature/topology setting 215
 - reports 313
- Executive report 321
- exporting
 - Active Flow Pipelining data 310
 - community and device information 61
 - Fast Connection Setup data 313
 - percentage reduction, device 287, 290
 - QoS data 301, 302, 305
 - schedule logs 66
- external routing for packet interception 221
- extracting configurations 47, 75

F

- failed tasks
 - icon on Devices page 36
 - rescheduling and acknowledging 65
- failure log, polling 349
- Fast Connection Setup
 - configuring 201, 204
 - report 311
- features, CMS 15
- features/topology, configuring 215
- file locations, CMS and JRE 25
- files
 - diagnostic, CMS 333
 - retrieving from devices 56

- filters, reduction
 - application 144
 - source/destination 216
- flow statistics, retrieving 56
- Forward Error Correction, configuring 201
- front panel access, device 126
- FTP application type 134
- FTP server
 - installing 23
 - on CMS server 56, 341

G

- gateways, configuring
 - default 88
 - in Multi-Path configurations 102
- global configurations 67
- groups
 - CMS device 331
 - CMS user 337
- guaranteed bandwidths
 - configuring 174, 180
 - overriding 178

H

- hardware requirements 21
- high-availability support 104
- HMAC/SHA-1 authentication for IPsec 250
- HTTP acceleration
 - about 196
 - configuring 210
 - feature/topology setting 215
 - reports 315
- HTTPS 17
- Hub and Spoke topology 212

I

- IANA port map 320
- icons on Devices page 35
- idle user timeout 123
- inbound QoS
 - configuring 188
 - monitoring 304
- installing CMS 21

- interface
 - link failure propagation 105
 - settings, configuring 104
- Intranet traffic class 188
- IP address
 - configuring 88
 - secondary address for Multi-Path 102
- IP compression tunnel mode, configuring 153
- IPSec
 - configuration procedure 247
 - defining templates 249
 - icon on Devices page 35

J

- Java Runtime Environment (JRE)
 - location of files 25
 - version 22

K

- key lifetimes, IPSec 250
- keys, RADIUS 122

L

- latency threshold
 - Multi-Path 241
 - WAN performance monitoring 233
- Layer 2 multicast traffic 294
- license keys
 - CMS
 - about 357
 - entering 343
 - device
 - deployment procedure 261
 - generating and applying 264
 - importing Authorization Codes 262
 - viewing status 266
- License Server, accessing 266
- lifetimes, IPSec key 250
- link failure propagation 105
- load balancing
 - across reduction tunnels 147
 - across routers 112
- local domain name 89
- local routes, adding static 97
- local users, device 123

- locations of CMS and JRE files 25
- logging in and out 18, 27
- login retries, SSH 119
- logs
 - polling catch-up and failure 349
 - system 351
- logs, retrieving from devices 56

M

- MAC addresses 93
- maximum bandwidths
 - inbound 190
 - outbound
 - configuring 174, 180
 - overriding 178
- max-mem topology setting 213
- MD5
 - for IPSec 250
 - for OSPF 111
- Mesh topology 212
- meta packets
 - IP compression 153
 - ToS/DSCP values in 184
- minimum WAN speed 172, 183
- monitoring
 - applications 138
 - inbound QoS 304
 - outbound QoS 300
 - percentage reduction 283
 - tunnel status 296
 - WAN performance 274
 - configuring 231
- Monitoring pages 269
- multi-flow emulation 154
- Multi-Path, configuring
 - about 235
 - addresses 101
 - defining endpoints 240
 - defining templates 238
 - icon on Devices page 35
 - router configuration 243
- My Peribit page 270

N

- NetFlow records, generating 115
- network
 - interfaces, configuring 104
 - settings, configuring 88
- Not Applicable icon 285
- NSM
 - configuring applications 144
 - defining endpoints 141
 - icon on Devices page 35
- NTP
 - configuring 107
 - icon on Devices page 35

O

- off-path deployments, configuring 220
- operator access, device 125
- OSPF 110
- outbound QoS
 - about 155
 - and packet flow acceleration 198
 - configuration procedure 166
 - congestion control 172, 183
 - dedicated and oversubscribed WANs 160
 - defining endpoints 169, 170, 180
 - defining settings by endpoint 176, 247
 - defining templates 179
 - defining traffic classes 136, 173
 - excluding LAN/WAN addresses 96
 - monitoring 300
 - outbound WAN speed
 - about 160
 - defining 169, 180
 - running the Setup Wizard 168
 - starting and stopping 187
 - ToS/DSCP values 184
 - virtual endpoints 184
- overflow, traffic volume 294
- oversubscribed WANs 160
- overviews
 - CMS 17
 - configurations 67

P

- Packet Flow Acceleration
 - Active Flow Pipelining 201, 203
 - Fast Connection Setup 201, 204
 - Forward Error Correction 201
 - icon on Devices page 35
- packet interception
 - configuring 220
 - icons on Devices page 36
- packet size distribution statistics 295
- pages
 - Auto-Deployment 254
 - Configurations 73
 - Devices 33
 - License Management 261
 - Monitoring 269
 - Schedules 62
- partial configurations 67
- passthrough statistics 293
- passwords
 - CMS users 325
 - device 123
 - OSPF 111
 - registration server
 - applying to devices 58
 - updating in CMS 330
 - RIP 112, 222
- peak data reduction 322
- pending tasks
 - adding devices to 65
 - cancelling 64
- performance monitoring, WAN
 - configuring 231
 - viewing reports 275
- policy-based routing for packet interception 221
- polling
 - catch-up and failure logs 349
 - intervals 346
- port numbers
 - application 377
 - in application definitions 134
 - RADIUS server 122
- port, CMS Web server
 - changing 345
 - default 22, 25

- preferred path 239
- pre-fetch for HTTP acceleration 196
- pre-installation tasks 22
- prime time, defining 218
- privilege levels
 - for CMS accounts 335
 - for device accounts 123
- protocols, in application definitions 134

Q

- QoS
 - inbound 188
 - outbound, see "outbound QoS"
- quick setup 27

R

- RADIUS servers and server groups 121
- read-only access 335
- rebooting devices 42
- recommended tasks 30
- recovery packets, Forward Error Correction 202
- recurring tasks, adding devices to 65
- Reduced traffic class 188
- reduction percentage, monitoring 283
- reduction subnets
 - configuring 94
 - filtering source/destination 216
- reduction tunnels 139
- registration server
 - designating 114
 - password
 - applying to devices 58
 - updating in CMS 330
- remote circuit speeds
 - and router overhead 159
 - configuring 170, 181
- remote routes 146
- reports
 - about 269
 - Active Flow Pipelining 308
 - Application Summary
 - all traffic 291
 - WAN traffic 282
 - CIFS and Exchange acceleration 313
 - data reduction 283

- Executive 321
- Fast Connection Setup 311
- HTTP acceleration 315
- inbound QoS 304
- My Peribit 270
- outbound QoS 300
- packet size distribution 295
- Passthrough Data 293
- throughput, WAN traffic 280
- traffic 317
- tunnel status 296
- viewing in device time 269
- WAN 274
- rescheduling failed tasks 65
- restoring
 - backup databases 352
 - configurations 54
- retransmissions, RADIUS 122
- retries, SSH login 119
- retrieving device files and statistics 56
- RIP
 - for dynamic routes 110
 - for packet interception 221
- rolling back
 - boot images 40
 - configurations 51
- root user account 18
- route injection 221
- router balancing, route-based 112
- router configuration
 - for Multi-Path 243
 - for packet interception 224
- routes
 - adding static 97
 - remote 146
 - router polling 99

S

- safe mode, devices 60
- schedule log 66
- scheduled tasks
 - failed tasks 65
 - viewing details 62
- scheduler, stopping and restarting 344
- Schedules page 62

- secondary IP address for Multi-Path 102
- secret key, RADIUS 122
- security
 - CMS
 - access lists 339
 - user accounts 335
 - device
 - defining local users 123
 - front panel access 126
 - operator access 125
- security features 116
 - defining RADIUS servers and server groups 121
 - selecting authentication methods 117
- server time, viewing reports in 269
- servers
 - DHCP 253
 - DNS 89, 253, 320
 - NetFlow 115
 - RADIUS 121
- Setup Wizard, outbound QoS 168
- SMB signing, disabling 206
- SNMP 108
- software requirements 21
- source/destination filters 216
- Spoke topology 212
- static routes, adding 97
- statistics
 - acceleration 308
 - application
 - all traffic 291
 - WAN traffic 282
- executive summary 321
- inbound QoS 304
- outbound QoS 300
- packet size distribution 295
- passthrough traffic 293
- reduction percentage 283
- retrieving from devices 56
- throughput, WAN traffic 280
- top traffic 317
- understanding retrieved data 363
- WAN 274

- status
 - of applied licenses 266
 - of auto-deployment 259
- subnet mask, configuring 88
- subnets
 - advertising for reduction 94
 - defining whether encryption is required 251
 - excluding from outbound QoS 96
 - excluding from reduction 216
 - filtering the Traffic report 319
 - unadvertised subnets and outbound QoS 166
- subnets, excluding from default assemblers 150
- summary of device configurations 43
- support
 - browser 17
 - generating CMS diagnostic files 333
 - retrieving diagnostic files from devices 56
 - technical 13
- Syslog
 - enabling CMS 342
 - enabling on devices 109
 - list of events 359
 - retrieving from devices 56
- system log 351

T

- tasks, managing scheduled 62
- TCP traffic class 188
- technical support 13
- templates
 - IPSec 249
 - Multi-Path 238
 - outbound QoS 179
- thresholds, loss and latency
 - Multi-Path 241
 - WAN performance monitoring 233
- throughput statistics, WAN traffic 280
- time settings
 - NTP server 107
 - time zone and daylight savings 92, 103
- timeout
 - idle user 123
 - RADIUS server 122
- topology settings 212

- ToS/DSCP values
 - defining by QoS traffic class 184
 - in application definitions 135
 - in Multi-Path configurations 237
- traffic classes
 - in the QoS Wizard 173
 - inbound QoS 188
 - outbound QoS and Multi-Path 136
- traffic statistics 317
- tunnel mode 153
- tunnels
 - monitoring 296
 - reduction 139
- types of applications 133

U

- UDP
 - and meta packets 153
 - in application definitions 134
- unadvertised subnets and outbound QoS 166
- uninstalling CMS 27
- upgrading from a previous release 23
- URLs, in application definitions 135
- user accounts
 - CMS
 - about 18
 - changing passwords 325
 - defining 335
 - device 123
 - user groups 337
- users, logged in 326

V

- validating remote routes 147
- verifying running configurations 43
- versions, configuration 72
- virtual endpoints, outbound QoS 184
- VLAN 802.1q support 105

W

- WAN circuit speeds
 - about 159
 - congestion control 172, 183

- WAN performance monitoring
 - configuring devices 231
 - configuring thresholds for reports 348
 - reports 274
- WAN reduction subnet
 - for off-path devices 96
 - for VLAN environments 106
- WCCP for packet interception 221
- Web console
 - CMS
 - about 19
 - logging in and out 18, 27
 - device, accessing 61
- Web server port, CMS
 - changing 345
 - default 22, 25
- Weighted Fair Queuing 175, 187
- Weighted Strict Priority 175, 187
- Wizard, outbound QoS 168

