



**WXOS Software for  
Application Acceleration Platforms**

# **WX/WXC Operator's Guide**

*Release 5.3*

**Juniper Networks, Inc.**  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089  
USA  
408-745-2000  
**[www.juniper.net](http://www.juniper.net)**

Part Number: 530-016333-01

This product includes a modified copy of the traceroute software developed by the University of California and its contributors. © 1990, 1993 The Regents of the University of California. A copy of the University of California copyright notice, license terms and disclaimer is available in the *WX/WXC Operator's Guide* on page 437.

This product includes a modified version of OpenSSL. © 2001-2006 Juniper Networks, Inc. All Rights Reserved. © 1998-2000 The OpenSSL Project. © 1995-1998 Eric Young. A copy of the Eric Young copyright notice, license terms and disclaimer is available in the *WX/WXC Operator's Guide* on page 438.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org>). A copy of the OpenSSL Project copyright notice, license terms and disclaimer is available in the *WX/WXC Operator's Guide* on page 438.

This product contains a modified version of the IPsec software developed by the KAME Project. A copy of the KAME copyright notice, license terms and disclaimer is available in the *WX/WXC Operator's Guide* on page 446.

This installation includes a modified version of ospfd. © 2001-2006 Juniper Networks, Inc. All Rights Reserved. ospfd © 1998 John T. Moy. ospfd is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

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.: Active Flow Pipelining, AFP, Application Flow Acceleration, AppFlow, Central Management System, CMS, ERX, E-series, ESP, Fast Connection Setup, Instant Virtual Extranet, Internet Processor, J2300, J4300, J6300, J-Protect, J-series, J-Web, JUNOS, JUNOScope, JUNOScript, JUNOSe, M5, M7i, M10, M10i, M20, M40, M40e, M160, M320, M-series, MMD, Molecular Sequence Reduction, MSR, NetScreen-5GT, NetScreen-5XP, NetScreen-5XT, NetScreen-25, NetScreen-50, NetScreen-204, NetScreen-208, NetScreen-500, NetScreen-5200, NetScreen-5400, NetScreen-IDP 10, NetScreen-IDP 100, NetScreen-IDP 500, NetScreen-Remote Security Client, NetScreen-Remote VPN Client, NetScreen-SA 1000 Series, NetScreen-SA 3000 Series, NetScreen-SA 5000 Series, NetScreen-SA Central Manager, NetScreen Secure Access, NetScreen-SM 3000, NetScreen-Security Manager, Network Sequence Caching, NSC, NMC-RX, Packet Flow Acceleration, PFA, Policy-Based Multipath, PBM, SDX, Stateful Signature, T320, T640, T-series, TX Matrix, and WX, WXC, and WXOS. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. All specifications are subject to change without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Copyright © 2006, Juniper Networks, Inc. All rights reserved. Printed in USA.

Copyright © 2001-2005, Peribit Networks, Inc. All rights reserved. Printed in USA.

*WX/WXC Operator's Guide*, Release 5.3

Revision History  
May 2006 —Rev 1

The information in this document is current as of the date listed in the revision history.

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.

## Year 2000 Notice

Juniper Networks hardware and software products are Year 2000 compliant. The JUNOS software has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

## Software License

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

For complete product documentation, please see the Juniper Networks Web site at [www.juniper.net/techpubs](http://www.juniper.net/techpubs).

## End User License Agreement

**READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE.** BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

**1. The Parties.** The parties to this Agreement are Juniper Networks, Inc. and its subsidiaries (collectively "Juniper"), and the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

**2. The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, and updates and releases of such software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller.

**3. License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use the Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller, unless the applicable Juniper documentation expressly permits installation on non-Juniper equipment.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

**4. Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use the Software on non-Juniper equipment where the Juniper documentation does not expressly permit installation on non-Juniper equipment; (j) use the Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; or (k) use the Software in any manner other than as expressly provided herein.

**5. Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

**6. Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

**7. Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

**8. Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

**9. Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

**10. Taxes.** All license fees for the Software are exclusive of taxes, withholdings, duties, or levies (collectively "Taxes"). Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software.

**11. Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

**12. Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

**13. Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

**14. Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.

**15. Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and

contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

# Table of Contents

	<b>About This Guide</b>	<b>15</b>
	Audience .....	15
	Package Contents .....	15
	Operator's Guide Contents .....	16
	Document Conventions .....	17
	Commonly Used Terms .....	17
	Typographical Conventions .....	18
	Technical Support .....	18
	Obtaining Additional Product Information .....	18
<b>Chapter 1</b>	<b>Introduction</b>	<b>19</b>
	About the WX and WXC Devices .....	19
	Features and Benefits .....	20
	What's New in Version 5.3 .....	21
	Sample Topologies .....	22
	Typical Inline Deployment .....	22
	Off-Path Deployment .....	22
	Point-to-Multipoint Topology .....	23
	Virtual Private Network (VPN) Topology .....	24
	Basic Concepts .....	24
	Communities and Registration Servers .....	25
	Reduction Tunnels .....	26
	Local Routes and Reduction Subnets .....	26
	Remote Routes .....	26
	Community Topologies .....	27
	High Availability Support .....	27
	Profile Mode .....	28
	Central Management System (CMS) .....	28
	Where to Go Next .....	28
<b>Chapter 2</b>	<b>Installation</b>	<b>29</b>
	Before You Begin .....	29
	Battery Warning .....	30
	Manual and Automatic Installations .....	30
	Inline and Off-path Installations .....	30
	Interface Speeds and Modes .....	31
	Installing the SR-100 .....	31
	Hardware Installation .....	32
	Copper-wire Interfaces .....	32
	Fiber-optic Interfaces .....	34
	Disconnecting Power from the SR-100 .....	34
	Configuring Network Settings .....	35
	Installing the WX 100 .....	37

Hardware Installation .....	38
Copper-wire Interfaces .....	38
Fiber-optic Interfaces .....	39
Disconnecting Power from the WX 100.....	40
Configuring Network Settings .....	41
Connecting Client Devices to the Server .....	43
Disconnecting Client Devices from the Server .....	44
Running Quick Setup through the Web Console.....	45
Post-Installation Tasks.....	53
Where to Go Next .....	53

**Chapter 3    Configuring Basic Setup Policies** **55**

Using the Web Console .....	55
Logging In.....	55
Understanding the WXOS Web Console Interface .....	56
Using Special Characters.....	57
Configuring Basic Setup Policies.....	57
Configuring Device Address and Contact Information .....	57
Configuring the Interface Settings.....	59
Configuring 802.1Q VLAN Support .....	61
Configuring Time Settings .....	62
Obtaining a Permanent License.....	62
Enabling SNMP.....	64
Enabling Syslog Reporting.....	65
Configuring Local Routes .....	66
Adding Static Routes .....	68
Enabling RIP and OSPF Support.....	69
Enabling Route Polling .....	70
Importing a Routing Table.....	71
Enabling Route-Based Router Balancing.....	73
Configuring Registration Servers and Communities.....	75
Defining Registration Servers and Passwords.....	75
Defining Communities .....	77
Configuring AAA .....	79
Selecting Authentication Methods.....	80
Enabling Authorization Checking.....	81
Defining RADIUS Servers and Server Groups.....	82
Defining Local Users.....	84
Securing Operator Access .....	86
Securing Front Panel Access .....	87
Changing the Packet Capture Password.....	88
Managing Applications .....	88
About Application Definitions .....	89
Viewing the Application Overview .....	91
Configuring Application Definitions .....	92
Testing New Application Definitions .....	95
Assigning Applications to Traffic Classes .....	95
Monitoring Applications.....	97

<b>Chapter 4</b>	<b>Configuring Advanced Setup Policies</b>	<b>99</b>
	Setting Community Feature and Topology Parameters .....	99
	Filtering Data Reduction by Source and Destination .....	101
	Configuring the ARP Table .....	104
	Defining the Prime Time .....	105
	Configuring Packet Interception .....	106
	Configuring Packet Interception for Off-Path Devices .....	106
	RIP Router/Switch Configuration Commands .....	109
	Single Layer 3 Switch .....	109
	Dual Off-Path Devices on Two Layer 3 Switches .....	111
	WCCP Router Configuration Commands .....	112
	External Policy-Based Router Commands .....	113
	Alternatives to Packet Interception .....	113
	Layer 2 Switch Sandwich .....	113
	Layer 3 Switch Sandwich .....	114
	Configuring Policy-Based Multi-Path .....	115
	Procedure for Configuring Multi-Path .....	116
	Enabling Multi-Path and Defining Marking Methods .....	117
	Defining Multi-Path Templates .....	119
	Defining Multi-Path Endpoints .....	121
	Configuring Routers to Support Multi-Path .....	123
	Configuring WAN Performance Monitoring .....	124
<b>Chapter 5</b>	<b>Configuring Reduction Policies</b>	<b>129</b>
	Configuring Basic Reduction Policies .....	129
	Configuring Endpoints for Reduction Tunnels .....	129
	Advertising Reduction Subnets .....	131
	Configuring Network Sequence Caching .....	134
	Reducing Applications .....	135
	Configuring Advanced Reduction Policies .....	137
	Viewing and Fetching Remote Routes .....	138
	Configuring Tunnel Load Balancing Policies .....	139
	Defining Default Assemblers .....	141
	Defining Preferred Assemblers .....	144
	Configuring Tunnel Mode Settings .....	145
	Configuring Pre-Synchronization for Network Sequence Caching .....	146
	Configuring Tunnel Switching .....	148
	Tunnel Switching Between Communities .....	148
	Procedure for Configuring Tunnel Switching Between Communities .....	149
	Tunnel Switching Between Hub and Spoke Devices .....	150
<b>Chapter 6</b>	<b>Applying Quality of Service (QoS) Policies</b>	<b>153</b>
	Using Outbound QoS to Enhance Performance .....	153
	Understanding Outbound Bandwidth Management .....	154
	Traffic Classes and Bandwidths .....	155
	QoS Templates and Endpoints .....	155
	WAN Circuit Speeds and Router Overhead .....	156
	Dedicated, Oversubscribed, and Variable Rate WANs .....	157
	Direct Setup Versus Wizard Configuration Results .....	159
	Class Priorities and Excess Bandwidth Allocation .....	161
	ToS/DSCP Values .....	162
	Unadvertised Subnets .....	162

Configuring Outbound QoS Policies .....	162
Procedure for Configuring Outbound QoS Policies.....	163
Using the Outbound QoS Setup Wizard .....	164
Defining Outbound QoS Settings by Endpoint .....	171
Defining Traffic Classes .....	173
Defining Outbound QoS Templates .....	174
Defining Outbound QoS Endpoints.....	176
Changing Outbound ToS/DSCP Values.....	181
Starting and Stopping Outbound QoS .....	183
Processing Queues Based on Incoming ToS/DSCP Values .....	184
Configuring Inbound QoS Policies.....	184
Summary of Key Terms .....	187
<b>Chapter 7   Accelerating WAN Traffic</b>	<b>189</b>
Packet Flow Acceleration .....	189
Overview of Packet Flow Acceleration.....	189
Active Flow Pipelining.....	190
Forward Error Correction .....	192
Fast Connection Setup .....	192
Requirements for Using Packet Flow Acceleration.....	193
Enabling Packet Flow Acceleration by Endpoint .....	194
Enabling Active Flow Pipelining by Application .....	197
Enabling Fast Connection Setup by Application .....	198
Application Flow Acceleration.....	200
Overview of Application Flow Acceleration .....	200
Microsoft CIFS and Microsoft Exchange Acceleration .....	201
HTTP Acceleration .....	202
Enabling Microsoft CIFS Acceleration .....	204
Enabling Microsoft Exchange Acceleration .....	207
Enabling HTTP Acceleration .....	209
<b>Chapter 8   Configuring IP Security (IPSec)</b>	<b>213</b>
Overview of IPSec.....	213
Default IPSec Policy.....	213
IPSec Implementation Details.....	214
Procedure for Configuring IPSec Policies.....	215
Using the IPSec Setup Wizard .....	215
Defining IPSec Settings by Endpoint .....	219
Defining IPSec Templates .....	222
Defining the Default IPSec Policy .....	224
<b>Chapter 9   Monitoring and Reporting</b>	<b>227</b>
Viewing and Printing Reports.....	227
WAN Statistics .....	228
WAN Throughput Statistics .....	228
WAN Application Summary .....	230
WAN Performance Statistics.....	231
Reduction Statistics.....	235
Device Throughput Statistics.....	235
Data Reduction Statistics .....	237
Application Summary Statistics .....	240
Passthrough Statistics .....	242
Packet Size Distribution Statistics .....	243



Outbound Bandwidth Statistics .....	244
Inbound Bandwidth Statistics .....	246
Acceleration Statistics .....	248
Active Flow Pipelining Statistics .....	248
Fast Connection Setup Statistics .....	250
Forward Error Correction Statistics .....	252
CIFS and Exchange Acceleration Statistics .....	253
HTTP Acceleration Statistics .....	254
Traffic Statistics .....	256
Endpoints Summary .....	258
Executive Summary .....	260
<b>Chapter 10    Maintaining WX Devices</b>	<b>263</b>
Maintaining Configurations and Software .....	263
Saving the Device Configuration .....	263
Displaying the Running Configuration .....	265
Loading a Device Configuration File .....	266
Loading a Boot Image .....	267
Clearing Application Monitoring Statistics .....	268
Setting the Device to the Factory Default Configuration .....	268
Rebooting the Device .....	270
Using Maintenance Tools .....	271
Pinging a Network Device .....	271
Running a Traceroute to a Network Device .....	272
Running a Packet Capture .....	273
Generating NetFlow Records .....	274
Entering CLI Commands from the Web Console .....	275
Viewing and Saving System Logs .....	277
Viewing and Saving the Access Control Log .....	278
Exporting Performance Data .....	279
Creating a Diagnostic File .....	280
Viewing the WX 100 Server/Client Summary .....	281
<b>Chapter 11    Using the Command Line Interface (CLI)</b>	<b>283</b>
Accessing the CLI .....	283
Using a Secure Shell Program from a Remote Workstation .....	284
Using a Terminal Connected to the Serial Port .....	284
Logging In Using the CLI .....	284
CLI Basics .....	285
Command Modes .....	286
CLI Command Summary .....	287
System-Level Commands .....	290
commit .....	290
configure .....	290
copy .....	290
embed .....	291
import-route-table .....	291
list .....	292
load-config .....	292
packet-capture .....	293
ping .....	294
reboot .....	295
remove .....	296

rollback .....	296
save-config .....	297
set .....	298
shutdown .....	298
source .....	299
support .....	299
upgrade .....	299
traceroute .....	300
Configuration Commands .....	300
configure aaa .....	300
configure acceleration .....	302
configure application .....	309
configure arp .....	312
configure backup .....	312
configure clock .....	315
configure console .....	315
configure dns .....	316
configure filter .....	316
configure interface .....	318
configure ip .....	320
configure ipsec .....	320
configure license .....	324
configure mon-apps .....	325
configure multi-node .....	326
configure multi-path .....	326
configure ospf .....	330
configure packet-interception .....	331
configure prime-time .....	333
configure profile-mode .....	334
configure qos inbound .....	335
configure qos outbound .....	337
configure radius .....	342
configure reduction .....	343
configure reduction-subnet .....	350
configure reg-server .....	352
configure remote-routes .....	355
configure rip .....	356
configure route .....	356
configure route-poll .....	359
configure security .....	360
configure snmp .....	361
configure sntp .....	362
configure stack-group .....	363
configure syslog .....	364
configure top-talker .....	365
configure wan-performance-monitor .....	366
Show Commands .....	368
show aaa .....	368
show acceleration .....	368
show access-log .....	369
show all .....	369
show application .....	370
show arp .....	370
show backup-sr .....	370

show clock.....	370
show connection .....	370
show console.....	371
show contact .....	371
show dns .....	371
show filter .....	371
show flow-details .....	372
show import-route-table .....	373
show interface.....	373
show ip.....	374
show ipsec.....	374
show license.....	374
show location .....	375
show log.....	375
show mon-apps.....	375
show multi-node.....	375
show multi-node-status.....	375
show multi-path.....	375
show ospf.....	376
show packet-capture.....	376
show packet-interception.....	376
show prime-time .....	377
show profile-mode.....	377
show qos excl-filter.....	377
show qos inbound .....	377
show qos outbound .....	377
show radius .....	378
show reduction.....	378
show reduction-subnet .....	379
show reg-detail .....	379
show reg-server.....	380
show reg-summary.....	380
show remote-routes.....	380
show rip .....	381
show route .....	381
show route-poll.....	381
show security .....	382
show snmp.....	382
show snmp.....	382
show stack-group.....	382
show syslog .....	383
show system.....	383
show system-name.....	383
show top-talker.....	384
show uptime.....	384
show version .....	384
show wan-performance-mon.....	384

<b>Appendix A</b>	<b>WX Device Specifications</b>	<b>385</b>
	WX 15 Specifications .....	385
	WX 20 Specifications .....	387
	WX 50 and WX 60 Specifications.....	388
	WX 100 Specifications .....	390
	WXC 250 Specifications .....	391
	WXC 500 Specifications .....	393
	DB9 Console Port Pin-Outs .....	395
<b>Appendix B</b>	<b>SNMP Traps and Syslog Messages</b>	<b>397</b>
	SNMP Traps .....	397
	Syslog Messages.....	399
<b>Appendix C</b>	<b>Understanding Exported Data Results</b>	<b>405</b>
	NetFlow Version 5 Export .....	405
	Performance Statistics Export .....	406
	General Device Information.....	407
	Data Section Information.....	407
	System Session Statistics .....	408
	Reduction Session Statistics.....	410
	Application Session Statistics .....	410
	WAN Statistics .....	411
	Application Flow Acceleration Statistics.....	411
	Bandwidth Management Statistics.....	412
	WAN Performance Statistics.....	413
	Inbound Traffic By Port Statistics.....	413
	Top Traffic Export .....	414
<b>Appendix D</b>	<b>Common Application Port Numbers</b>	<b>415</b>
<b>Appendix E</b>	<b>Safety and EMC Certifications</b>	<b>417</b>
	Product Reclamation and Recycling Program .....	418
<b>Appendix F</b>	<b>Safety Recommendations and Warnings</b>	<b>419</b>
	Power Cable Warning (Japanese) .....	419
	VCCI Compliance.....	419
	Lightning Activity Warning .....	420
	Jewelry Removal Warning .....	420
	Installation Warning .....	420
	IT Power Statement .....	420
	SELV Circuit Warning .....	420
	Circuit Breaker (15A) Warning.....	420
	Grounded Equipment Warning.....	421
	Class 1 Laser Product Warning .....	421
	Laser Beam Warning .....	421
	Battery Warning .....	421
	Rack Mounting of Systems .....	422
	Anti-static Precautions .....	422

<b>Glossary</b>	<b>423</b>
<b>Index</b> .....	<b>427</b>
<b>Copyrights</b>	<b>437</b>
Traceroute Copyright License.....	437
OpenSSL Copyright License .....	438
GNU GENERAL PUBLIC LICENSE.....	440
KAME Copyright License.....	446



# About This Guide

Welcome to the operator's guide for the Juniper® WX and WXC Application Acceleration Platforms. With their patented Molecular Sequence Reduction™ (MSR) technology and Network Sequence Caching™ (NSC), the WX devices provide instant WAN capacity to your existing network.



---

**NOTE:** This manual describes WXOS 5.3, which is supported only by the WX 100 in a stack configuration (a WX 100 server with one or more client devices).

---

The following sections describe the audience, organization, and typographical conventions used in this manual.

## Audience

---

This manual is intended for administrators responsible for configuring and managing WX and WXC devices. It is assumed that readers of this manual are familiar with their network architecture and devices, and can perform basic network configuration procedures.

## Package Contents

---

WX 100 devices are shipped with the following:

- 1 WX 100
- 1 Female/female DB-9 crossover cable
- 2 Rack-mount flanges for rack mount installation
- 6 Screws for the rack-mount flanges
- 4 Rubber feet for desktop placement
- 2 Power Cords
- 1 Quick Start Card
- 1 Documentation/Utilities CD
- 1 Release notes document

## Operator's Guide Contents

---

- Chapter 1, “Introduction” on page 19  

This chapter introduces the WX and WXC devices, describes the new features, and provides sample topologies for deployment.
- Chapter 2, “Installation” on page 29  

This chapter describes how to install and initially configure WX and WXC devices.
- Chapter 3, “Configuring Basic Setup Policies” on page 55  

This chapter describes how to configure basic policies through the Web console, such as IP parameters, security settings, and discovery of local routes.
- Chapter 4, “Configuring Advanced Setup Policies” on page 99  

This chapter describes how to configure advanced policies, such as topology parameters, packet interception, and Policy-Based Multi-Path™ (PBM).
- Chapter 5, “Configuring Reduction Policies” on page 129  

This chapter describes how to configure policy settings for data reduction, and the communication links with other devices in the community.
- Chapter 6, “Applying Quality of Service (QoS) Policies” on page 153  

This chapter describes how to configure outbound and inbound Quality of Service (QoS) policy settings, including traffic classes, WAN circuit speeds, and guaranteed bandwidths.
- Chapter 7, “Accelerating WAN Traffic” on page 189  

This chapter describes how to configure Packet Flow Acceleration™ (PFA™) for TCP applications, and Application Flow Acceleration™ (AppFlow™) for CIFS, Exchange, and HTTP traffic.
- Chapter 8, “Configuring IP Security (IPSec)” on page 213  

This chapter describes how to configure IPSec to encrypt the traffic between two WX and WXC devices.
- Chapter 9, “Monitoring and Reporting” on page 227  

This chapter describes the detailed graphs and reports that you use to monitor network performance.
- Chapter 10, “Maintaining WX Devices” on page 263  

This chapter describes how to maintain and manage the WX and WXC devices, and covers topics such as saving configuration files and displaying system log files.



- Chapter 11, “Using the Command Line Interface (CLI)” on page 283  
This chapter describes how to set up and configure the WX and WXC device using the Command Line Interface (CLI).
- Appendix A, “WX Device Specifications” on page 385  
This appendix lists the specifications for each type of WX and WXC device.
- Appendix B, “SNMP Traps and Syslog Messages” on page 397  
This appendix describes SNMP Trap and Syslog messages generated by the WX and WXC devices.
- Appendix C, “Understanding Exported Data Results” on page 405  
This appendix describes the details of exported data results. After exporting the reduction statistics to a comma-separated values file, use this appendix to interpret the data.
- Appendix D, “Common Application Port Numbers” on page 415  
This appendix provides a listing of common application port numbers that you can use when defining new applications.
- Appendix E, “Safety and EMC Certifications” on page 417  
This appendix lists the safety and EMC certifications for each type of WX and WXC device.
- Appendix F, “Safety Recommendations and Warnings” on page 419  
This appendix lists hardware safety recommendations and warnings.
- “Glossary” on page 423  
The glossary provides definitions of terms used throughout this manual.

## Document Conventions

---

This section describes conventions used throughout this manual.

### Commonly Used Terms

WX and WXC devices can be configured through a Graphical User Interface (GUI) Web console or Command Line Interface (CLI). When referring to these configuration methods, the following terminology is used:

- Web console — Web-based console.
- CLI — Command Line Interface.

## Typographical Conventions

The following table lists the typographical conventions used in this manual.

Convention	Meaning	Example
courier font	Text that you enter from your keyboard.	Enter the following command: <code>a:\setup</code>
Angle brackets	Encloses variables that you must substitute another value for.	set ip < IP address >
Square brackets	Encloses optional parameters.	show log [ < n > ]
Curved brackets	Encloses related parameters.	set mode { on   off }
italics	Names of manuals, directories, files, or Uniform Resource Locators (URLs).	The address of Juniper's web site is <i><a href="http://www.juniper.net">http://www.juniper.net</a></i> .

## Technical Support

For technical support, use the following methods:

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

## Obtaining Additional Product Information

In addition to this operator's guide, a printed Quick Start card and a copy of the Release Notes are enclosed with each device. Refer to the Quick Start card for product installation instructions, and the Release Notes for the latest product information.

For additional product information, please visit our web site at <http://www.juniper.net>.

## Chapter 1

# Introduction

The following sections provide an overview of the WX and WXC devices, including a description of the new features in this release:

- About the WX and WXC Devices on page 19
- Features and Benefits on page 20
- What's New in Version 5.3 on page 21
- Sample Topologies on page 22
- Basic Concepts on page 24
- Central Management System (CMS) on page 28

### About the WX and WXC Devices

---

The WX and WXC devices are LAN-based network devices that enhance the throughput of WAN circuits by addressing the three constraints on WAN performance—bandwidth, latency, and application contention. Installed on each side of a WAN circuit, the WX and WXC devices use the following technologies to compress, accelerate, and manage WAN traffic:

- **Molecular Sequence Reduction (MSR).** Based on algorithms used to find repeating patterns in DNA molecules, MSR locates repeated data patterns at the byte level, in real time, across all IP application sessions. Repeated patterns are sent as symbols, which the receiving device assembles (restores) from a shared dictionary. The reduction in traffic effectively increases the WAN bandwidth, reduces network congestion, and improves overall data flow.
- **Network Sequence Caching (NSC).** An enhanced disk-based version of MSR available between WXC devices. NSC 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). NSC is most effective where large files are often sent over the WAN, such as for database backups.

- **Quality of Service (QoS).** Application contention for available WAN bandwidth can be tightly controlled by assigning applications to traffic classes, and setting guaranteed and maximum bandwidths for each class. Class priorities can be set to ensure that time-sensitive applications, like VoIP, receive a sufficient amount of bandwidth. WX and WXC devices can also honor and set the ToS/DSCP values used by QoS devices in your network.
- **Packet Flow Acceleration (PFA).** While MSR effectively increases available bandwidth, PFA provides several methods to improve TCP application performance in networks where the use of available bandwidth is constrained by network latency.
- **Application Flow Acceleration (AppFlow).** Provides application-level acceleration for Microsoft CIFS, Microsoft Exchange, and HTTP traffic.
- **Policy-Based Multi-Path (PBM).** Directs traffic to one of two paths based on the performance needs of an application and the performance of the path. When loss and/or latency exceed the specified thresholds, traffic can be directed to the alternate path.
- **Encryption.** IPSec encryption can be enabled on specific paths to protect traffic in environments that are not secure (such as the Internet and satellite links).

The various WX and WXC devices support Ethernet speeds up to 1 Gbps, and can process IP WAN traffic up to 45 Mbps (T3 speeds). Higher WAN speeds (up to OC-3/STM-1) can be supported by connecting client devices to the WX 100 (stack-group configuration).

You can monitor and manage WX and WXC devices through a secure Web console, a command line interface (CLI), or the Central Management System (CMS). You can also monitor device performance through an SNMP-based management system. For the specifications of each type of device, refer to “WX Device Specifications” on page 385.

## Features and Benefits

---

WX and WXC devices enable networks to achieve maximum capacity over wide-area network (WAN) links. The primary features and benefits include:

- **Substantial throughput gain** — Greatly improves WAN capacity, accelerates TCP applications in high-latency environments, and reduces the load on other network devices.
- **Scalable** — All remote WX and WXC devices can be managed and monitored at a central point using the Central Management System (CMS).
- **Immediate impact** — Gains are realized immediately when WX or WXC devices are installed in the network. No time-consuming build-out.
- **Transparent** — Operates transparently to existing network equipment, topologies, and WAN interfaces (such as Frame Relay, ATM). No network or application modifications are required.

- **Application independent** — Works on any application over IP (such as e-mail, database, Web, ERP, and so on). Uses open standard protocols.
- **QoS Interoperable** — Honors, retains, and sets QoS priority levels within your network. Can maintain application visibility for data flows, enabling WAN probes and WFQ to work effectively.
- **Intelligent Bandwidth Management** — Can allocate operator-defined bandwidth ranges by traffic classes for greater control of newly created bandwidth.
- **Failsafe non-stop operation** — Switch-to-wire on any hardware or software disruption, including power loss. A single device can be installed as a backup for multiple primary devices.
- **Easily managed** — Administrative access through an intuitive Web user interface (SSL), a command line interface (CLI) using SSH. Users can be authenticated and authorized locally or through a RADIUS server.
- **VPN and firewall friendly** — WX and WXC devices installed on the LAN side of encryption devices work seamlessly with VPNs and firewalls.
- **Secure** — Provides confidentiality and message integrity for WAN traffic.

### What's New in Version 5.3

---

WXOS 5.3 enhances the performance of the WX 100 in stack mode (a WX 100 server with one or more clients). Note the following:

- WXOS 5.3 is supported only on WX 100 servers in stack mode— it is not supported on any WX device in standalone mode. If you install WXOS 5.3, and then want to change the WX 100 server to standalone mode (no clients), you must roll back to the previous release.
- WXOS 5.3 supports only two topology ranges: Hub range 5 and Mesh range 0, both with all features enabled. Use Mesh range 0 in a point to point topology where two WX 100 servers communicate with each other.
- Application Flow Pipelining (AFP), outbound QoS, and congestion control are enabled by default.
- The WX 100 now supports up to 320 virtual endpoints.



**NOTE:** The WX 50 is no longer supported as a WX 100 client. Client devices must have a Gigabit Ethernet interface (WX 55, WX 60, WX 80, or WXC 500), and all clients on the same WX 100 server must be the same device type.

---

## Sample Topologies


The following topics provide sample deployment topologies for WX and WXC devices:

- Typical Inline Deployment on page 22
- Off-Path Deployment on page 22
- Point-to-Multipoint Topology on page 23
- Virtual Private Network (VPN) Topology on page 24

### Typical Inline Deployment

WX and WXC devices must be installed on both sides of the WAN. They are typically deployed in the data path between the LAN and the edge routers (Figure 1).

**Figure 1: Typical Inline Deployment**




When two or more devices are installed in the same community, a reduction tunnel is formed between them.

### Off-Path Deployment

WX and WXC devices are usually deployed in the physical data path between a LAN switch and a WAN edge router, with no changes to layer 3 routing. When interrupting the data path is not practical, such as in collapsed backbone environments where the switch and the router are the same physical device, you can deploy the device “off path” (Figure 2).

In an off-path deployment, the device’s 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).


**Figure 2: Off-Path Deployment**



## Point-to-Multipoint Topology

WX and WXC devices support multi-point configurations of both “hub and spoke” and “mesh” configurations between multiple enterprise sites (Figure 3).

**Figure 3: Deploying WX and WXC Devices in a Point-to-Multipoint Configuration**



In this example, a hub (located at headquarters) is accessed by workgroups in remote sites. Data reduction tunnels, which are automatically established and managed by the WX and WXC devices at the various corporate sites, continuously process and reduce the data traveling through these tunnels thereby reducing traffic on the WAN circuits and creating more bandwidth.

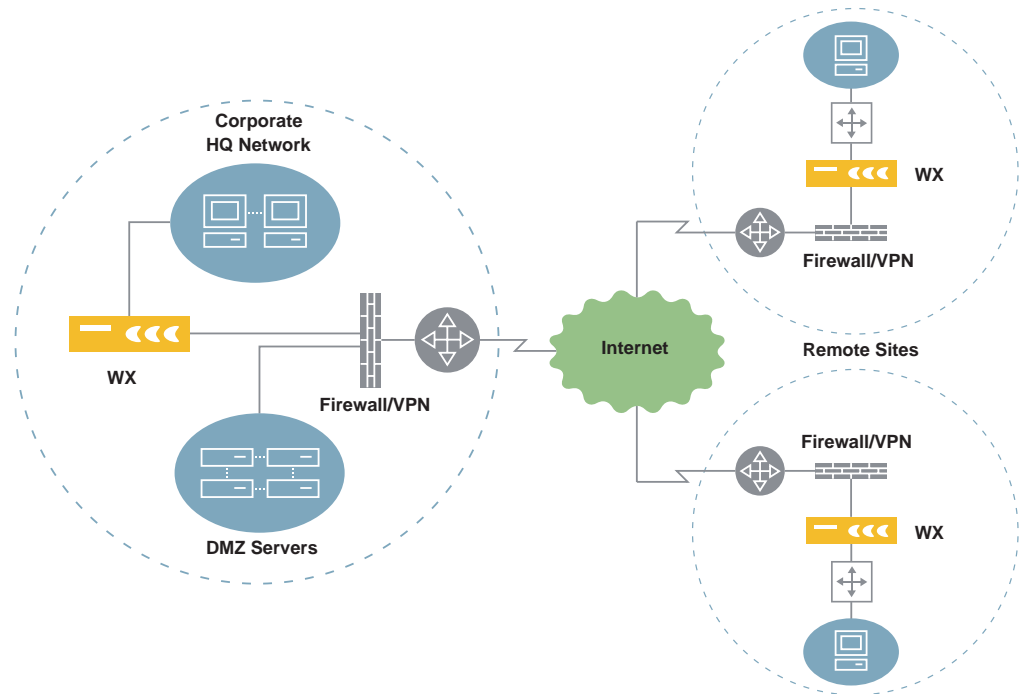
Note that it is not necessary to deploy a WX or WXC device for every remote site that links into the corporate headquarters network. In such instances, data from the hub is passed through without reduction.

In addition, Figure 3 shows four remote sites with dedicated connections to the Corporate HQ network. Since the WX and WXC devices are protocol and interface neutral, any of the four links could be any type of public or private packet-based service interface, such as Frame Relay or ATM.

## Virtual Private Network (VPN) Topology

WX and WXC devices operate transparently relative to existing network equipment, including firewalls and Virtual Private Network (VPN) devices (Figure 4).

**Figure 4: Deploying WX and WXC devices in a VPN Configuration**



By reducing data before it enters the VPN tunnel, the WX and WXC devices reduce the workload for the VPN devices. The same bandwidth multiplication effect is achieved for VPN encapsulated traffic as for unencapsulated traffic.

## Basic Concepts

The following topics provide an overview of key terms and concepts:

- Communities and Registration Servers on page 25
- Reduction Tunnels on page 26
- Local Routes and Reduction Subnets on page 26
- Remote Routes on page 26
- Community Topologies on page 27
- High Availability Support on page 27




## Communities and Registration Servers

At least two WX or WXC devices are required to perform data reduction (one on each side of the WAN). Two or more devices that can reduce and assemble data for each other are said to be in the same community. You can selectively enable or disable data reduction between any two devices in the same community.

When you install a WX or WXC device, you must specify the IP address of a registration server. The registration server is a WX or WXC device that stores the network information for all the other WX and WXC devices that report to it. Each device periodically contacts the registration server to identify the other devices in the same community. Initially, all devices reporting to the same registration server are in the Default community.

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 send reduced traffic between communities, you can create a hierarchical structure where selected devices reside in multiple communities (Figure 5).

**Figure 5: Example of Hierarchical Communities**



In most cases, one registration server can manage all devices and communities in the network. A secondary registration server can be specified to act as a backup if the primary server is unavailable.

## Reduction Tunnels

When you install a new WX or WXC device and specify a registration server, the device attempts to form a reduction tunnel with each registered device, or “endpoint,” in the same community. The existing devices also attempt to form tunnels with the new device, so that each device can have two types of tunnels—OUT tunnels that convey reduced data to remote devices, and IN tunnels that convey reduced data to be assembled.

At any time, you can disable data reduction from all other devices and/or reduce data only for specific devices in the community.

## Local Routes and Reduction Subnets

Local routes are the routes defined in the device's routing table. When you first install a WX or WXC 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
- Add dynamic routes using one of the following methods:
  - Enable the Open Shortest Path First (OSPF) and/or the Routing Information Protocol (RIPv1 or RIPv2)
  - Periodically poll the routing table of a Cisco router
  - Import a file of routes from an FTP server

Reduction subnets are the LAN subnets for which the local device can assemble the data reduced by other devices. Static routes and routes discovered dynamically on the Local interface are added to the list of reduction subnets, which can then be advertised to the other devices in the community. By default, only the subnets you select are advertised.

In some cases, such as in VLAN environments, some routes on the Local interface may be discovered only on the Remote interface. To advertise these subnets, you must enable the WAN reduction subnet option through the CLI so that routes discovered on the Remote interface are included on the list of reduction subnets.



**NOTE:** For off-path devices, where only the Local interface is connected to the network, all routes are listed as reduction subnets because the device cannot distinguish between local and remote routes. In this case, you must be careful to advertise only the routes on the LAN side of the device.

## Remote Routes

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


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

## Community Topologies

For each device in a community, you can select a community topology. The community topology setting ensures that each device's resources are allocated efficiently. There are two basic topologies:

- **Mesh.** Multiple devices are interconnected and each one can reduce and assemble data for all the others.
- **Hub and Spoke.** A central device (Hub) can reduce and assemble data for all other devices in the community (Figure 6). By default, the spoke devices reduce data only for the hub. A community can have multiple hubs. Each device attempts to form a reduction tunnel with a hub before creating tunnels to other devices.

**Figure 6: Deploying Devices in a Hub and Spoke Topology**



For Hub and Mesh devices, you can specify the maximum number of devices so that sufficient resources are allocated for the potential number of reduction tunnels.

## High Availability Support

For critical WAN links, you can install backup devices that take over when a primary device is unavailable. Each backup can support one or more primary devices.

In addition, the WX and WXC devices transparently operate in high-availability (HA) environments. The Local and Remote interfaces can be configured so that when a failure occurs on one interface, the other interface is disabled. 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.

For WX 15, WX 60, and WX 100 devices installed in HA environments, you can disable the hardware passthrough so that a power failure on either device will block all traffic, thus allowing the failure to be detected and traffic routed to the other device.

## **Profile Mode**

Profile Mode lets you see how a WX or WXC device performs in your network without affecting network traffic. In Profile Mode, the device passively calculates potential data reduction statistics for all traffic and for individual applications.

In addition, you can view the performance for specific remote subnets by defining “virtual” devices and associating one or more subnets with each virtual device. On the reduction reports, you can then select a virtual device from the Destination menu to view the performance for the associated remote subnets (refer to “Monitoring and Reporting” on page 227).

To use Profile Mode, the Local interface must be connected to a mirrored port on the LAN switch, and the Remote interface must be disconnected. For more information about setting up and using a device in Profile Mode, refer to “Profile Mode” on page 417.

## **Central Management System (CMS)**

---

The Central Management System (CMS) is a Web-based tool that lets you centrally manage the configuration and software upgrades for geographically dispersed WX and WXC devices. From the secure CMS Web console you can view the performance of all devices, and apply configuration changes and software upgrades to selected devices. You can also schedule such tasks as upgrades and reboots to occur during off-peak hours.

## **Where to Go Next**

---

Refer to “Installation” on page 29 for complete installation instructions, or “Configuring Basic Setup Policies” on page 55 for information on setting up WX and WXC devices through the Web console.

## Chapter 2

# Installation

This chapter describes how to install WX and WXC devices and perform the initial configuration. It covers the following topics:

- Before You Begin on page 29
- Manual and Automatic Installations on page 30
- Inline and Off-path Installations on page 30
- Running Quick Setup through the Web Console on page 45
- Post-Installation Tasks on page 53

### Before You Begin

---

Before you begin, complete the following pre-installation tasks.

- Ensure that sufficient power is available. Supply circuits should be protected by a maximum 20A circuit breaker.
- Ensure there is ample space and lighting. You need enough space to connect one or two CAT-5 UTP Ethernet data cables and two power cords to the WX 100, and the proper lighting to see the LEDs on the Ethernet data ports.
- Provide a minimum of six inches clearance in the front and back of the chassis. Do not install one device directly behind another where warm or hot air may be recirculated. There are no ventilation requirements above or below the device.
- Do not stack paper materials or heavy equipment on top of a device.
- For rack-mount installations, reserve space for a 2U form factor device.
- Identify a 10/100/1000 LAN port where you can connect the WX 100. This port is typically on an aggregation switch or other LAN device connected directly to the WAN router. The WX 100 is also available with two 1000 Base-SX fiber-optic Ethernet interfaces.
- Log in to the router that will be on the WAN side of the WX 100 and note the interface speed and duplex mode.

- Verify that all firewalls between WX and WXC devices allow traffic on TCP/UDP ports 3577 and 3578, and for the IPComp protocol (protocol number 108). WXOS 5.1 uses IPComp as the default tunnel mode. For WXOS 5.0 and earlier, the default tunnel mode is UDP.
- Reserve an IP address and identify the default gateway for the WX 100. The default gateway is the next hop on the WAN side of the device.

## Battery Warning



**WARNING:** WX and WXC devices have no user serviceable parts. Opening the device voids the warranty. As a safety caution, note that opening the chassis exposes a lithium battery. If you attempt to remove or replace the lithium cell, do not use a conductive instrument, as a short-circuit may cause the cell to explode. A replacement cell must be of the same type (CR2032). Dispose of a spent cell promptly—do not recharge, disassemble, or incinerate. Keep cells away from children.

## Manual and Automatic Installations

A manual installation consists of the following steps for each type of device:

1. Install the hardware and apply power
2. Configure network settings (such as IP address)
3. Run Quick Setup to define required configuration settings
4. Perform post-installation tasks for optional configuration settings

Step 2 through 4 can be performed automatically if you have the Central Management System (CMS) 5.0 (or later) and a DHCP server. Entire configurations, including network settings, can be predefined in CMS, and then downloaded automatically when power is first applied to the device. For more information, refer to the CMS administrator's guide.



**NOTE:** Automatic installation cannot be used for multi-node configurations (refer to “Multi-Node Configurations” on page 433).

## Inline and Off-path Installations

WX and WXC devices are usually installed in the data path (inline) between a LAN switch (or other aggregation device) and the WAN edge router. If interrupting the data path is not practical, such as in collapsed backbone environments, you can deploy the device “off path.” Installing a device off path is similar to an inline installation, except for the following:

- Do not disconnect any cables—simply connect the Local interface of the device to the switch or the router. Connecting directly to the router is recommended. The Local interface should be set to full-duplex (half-duplex may cause excessive collisions).
- Do not connect the Remote interface to the router. The Remote interface is not used, so you can apply power to the device without first verifying connectivity between the LAN and the router.
- After you run Quick Setup, use RIP, WCCP, or policy-based routing to route traffic to the off-path device, as described in “Configuring Packet Interception” on page 106.

The following sections describe how to install the WX 100 in the data path.

- Interface Speeds and Modes on page 31
- Installing the SR-100 on page 31
- Installing the WX 100 on page 37

## Interface Speeds and Modes

---

Interface speed and duplex settings should be the same across all devices: the switch, the WX 100 Local and Remote interfaces, and the router. This ensures connectivity through the device in the event of a power loss or a condition that causes a hardware bypass. Note that this is not an issue for fiber-optic version of the WX 100 because fiber does not support hardware bypass.

## Installing the SR-100

---

The SR-100 is an earlier version of the WX 100. The SR-100 has an LCD on the front panel, and the power and data ports are on the back panel (the WX 100 has no LCD, and the data ports are on the front panel). Like the WX 100, the SR-100 can be used as a standalone device or as a server to distribute the processing load to up to six client devices. Client devices are connected directly to the SR-100, and can be WX 55s, WX 60s, WX 80s, or WXC 500s (all clients must be the same device type).

The following sections describe the installation process for the SR-100.

- Hardware Installation on page 32
- Disconnecting Power from the SR-100 on page 34
- Configuring Network Settings on page 35
- Connecting Client Devices to the Server on page 43
- Disconnecting Client Devices from the Server on page 44

## Hardware Installation

To install the SR-100 in your network:

1. Set up the chassis.
  - To install the SR-100 in a 19-inch device rack, install the supplied brackets (front panel forward) to the sides of the device with the screws provided in the kit. Next, install the chassis in your network device rack. Leave adequate space to install additional client devices.
  - To install the SR-100 on a desktop, place the chassis upside down on a smooth, flat surface, and install the supplied rubber feet on the bottom of the chassis. Place the chassis on a desktop or on top of another device so that all four rubber feet are securely mounted to the flat surface.

The subsequent steps depend on whether the SR-100 has standard copper-wire or fiber-optic interfaces.

### Copper-wire Interfaces


2. Connect the network cables and verify connectivity.

The standard SR-100 has two 10/100/1000 auto-sensing, Ethernet interfaces. These REMOTE and LOCAL interfaces are on the back panel (Figure 7).



**NOTE:** Do not connect power to the device until Step 4.

**Figure 7: SR-100 Ethernet Ports on Back Panel**



To connect the network cables to the SR-100:

- a. Locate the cable that connects the switch (or other aggregating device) to the router.
- b. Disconnect this cable from the router port and connect it to the SR-100's LOCAL port.



- c. Using a straight-through cable (not provided), connect one end to the SR-100's REMOTE port and the other end to the router port.
3. Use one of the following methods to verify connectivity across the SR-100 when the power is off. This step ensures that the correct cables are used and that traffic will bypass (pass through) the SR-100 in the event of a power loss.
  - Ping a host on the remote side of the SR-100 from a host on the local side of the SR-100.
  - Observe the link status LEDs (if available) on the interfaces of the adjacent network devices (switch and router).
4. After you verify network connectivity across the SR-100, connect the supplied power cords to the dual power supplies on the back of the chassis, and then connect the power cords to the local power source.

The SR-100 maximum power usage is 250 Watts or 850 BTU/hour.



**WARNING:** The appliance is designed to work with IT power systems.

**Waarschuwing** Het apparaat is ontworpen om te functioneren met IT energiesystemen.

**Varoituis** Koje on suunniteltu toimimaan IT-sähkövoimajärjestelmien yhteydessä.

**Attention** Ce dispositif a été conçu pour fonctionner avec des systèmes d'alimentation IT.

**Warnung** Das Gerät ist für die Verwendung mit IT-Stromsystemen ausgelegt.

**Avvertenza** Il dispositivo è stato progettato per l'uso con sistemi di alimentazione IT.

**Advarsel** Utstyret er utfomet til bruk med IT-strømsystemer.

**Aviso** O dispositivo foi criado para operar com sistemas de corrente IT.

**¡Atención!** El equipo está diseñado para trabajar con sistemas de alimentación tipo IT.

**Varning!** Enheten är konstruerad för användning tillsammans med elkraftssystem av IT-typ.

---

Now that the SR-100 is installed and powered on, configure the network settings, as described in “Configuring Network Settings” on this page.

## Fiber-optic Interfaces

The fiber-optic SR-100 is installed in the same way as the copper-wire version, except that you can apply the power before you connect the cables. Fiber-optic technology does not support a hard-wire passthrough connectivity in the event of a power loss.



**NOTE:** The fiber-optic SR-100 should be installed in a high-availability environment. Data transmission stops during a reboot or a power failure.

## Disconnecting Power from the SR-100



**WARNING:** The appliance has more than one power supply connection. All connections must be removed completely to remove power from the unit completely.

**Waarschuwing** Deze eenheid heeft meer dan één stroomtoevoerverbinding; alle verbindingen moeten volledig worden verwijderd om de stroom van deze eenheid volledig te verwijderen.

**Varoitus** Tässä laitteessa on useampia virtalähdekytkentöjä. Kaikki kytkennät on irrotettava kokonaan, jotta virta poistettaisiin täysin laitteesta.

**Attention** Cette unité est équipée de plusieurs raccordements d'alimentation. Pour supprimer tout courant électrique de l'unité, tous les cordons d'alimentation doivent être débranchés.

**Warnung** Diese Einheit verfügt über mehr als einen Stromanschluß; um Strom gänzlich von der Einheit fernzuhalten, müssen alle Stromzufuhren abgetrennt sein.

**Avvertenza** Questa unità ha più di una connessione per alimentatore elettrico; tutte le connessioni devono essere completamente rimosse per togliere l'elettricità dall'unità.

**Advarsel** Denne enheten har mer enn én strømtilkobling. Alle tilkoblinger må kobles helt fra for å eliminere strøm fra enheten.

**Aviso** Este dispositivo possui mais do que uma conexão de fonte de alimentação de energia; para poder remover a fonte de alimentação de energia, deverão ser desconectadas todas as conexões existentes.

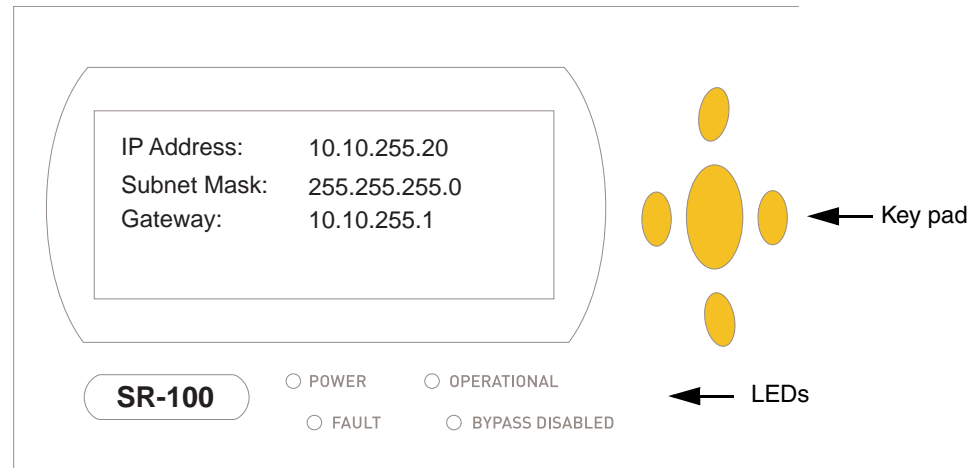
**¡Atención!** Esta unidad tiene más de una conexión de suministros de alimentación; para eliminar la alimentación por completo, deben desconectarse completamente todas las conexiones.

**Varning!** Denna enhet har mer än en strömförsörjningsanslutning; alla anslutningar måste vara helt avlägsnade innan strömtillförseln till enheten är fullständigt bruten.

## Configuring Network Settings

The configuration can be downloaded automatically from CMS when you first apply power to the device (refer to “Manual and Automatic Installations” on page 30). To configure the network settings manually, use the front-panel keypad and LCD as described below. Figure 8 shows the front panel keypad and LCD of the SR-100.

**Figure 8: SR-100 Front Panel Keypad and LCD**



1. Press the Enter button (center button) to display the menu in the LCD.
2. Press Enter to select Setup.
3. Use the front-panel keypad to assign an IP address, subnet mask, and default gateway:
  - Use the up or down arrow to select “IP Address” and press Enter.
  - Use the up and down arrow buttons to display a number (between 0-9).
  - Use the left and right arrow buttons to move to the previous or next character.
  - When the IP address is correct, press Enter.
  - Repeat this procedure for the subnet mask and default gateway. If you enter an invalid IP address, default gateway, or subnet mask, the values are not accepted when you press Enter, and the cursor returns to the IP menu. No error message is displayed.
  - After you enter the gateway address, use the left arrow to select “Save and Reboot” and press Enter.

The default gateway is typically the next hop on the Remote side of the WX 100. You may want to change the default gateway if you designate the device as a Default Assembler. After installing the SR-100, refer to “Defining Default Assemblers” on page 141 for more information.

4. After the device reboots, specify the speed and mode of each interface. By default, the Local and Remote interfaces are set to auto-negotiate. However, to avoid problems when the switch or router speed and duplex mode are set manually, it is **strongly recommended** that you manually configure the Local and Remote interface settings. Also, some routers may fail to auto-negotiate with fiber interfaces, so set the router and SR-100 interfaces to 1000 Mbps.

To manually configure the interfaces from the front panel:

- a. Press Enter to display the menu in the LCD.
- b. Use the down arrow to select the “Interfaces” option, and press Enter.
- c. Use the up arrow to select “Remote” and press Enter.
- d. Use the up arrow to select “Manual”.
- e. Use the right arrow to move to the speed setting.
- f. Use the up arrow to select one of the following settings to match the speed and mode of the router, and press Enter:

**Copper-wire options:** 10 half, 10 full, 100 half, 100 full, 1000 full

**Fiber-optic options:** 1000 full


- g. Use the up arrow to select “Local”, press Enter, and repeat Steps **d** to **f**. Select a setting to match the speed and mode of the switch.
- h. Select **Commit and Save** and press Enter.



**NOTE:** After installation, you can change the interface settings from the Web console or CLI.

5. Check the LEDs next to the Ethernet ports (Figure 9). Note that the fiber-optic version has only the top two LEDs (fiber-optic interfaces always run at 1 Gbps).

**Figure 9: Checking the Link LEDs for the SR-100**



- The LINK LEDs indicate the port is connected properly.
  - The 1000 LEDs indicate the port is running at 1000 Mbps.
  - The 100/10 LEDs indicate the port is running at 100 Mbps (not shown for fiber-optic interfaces). If the 1000M and 100/10 LEDs are off, the port is running at 10 Mbps.
6. If you install the copper-wire SR-100 in a high-availability environment, you can press the **Bypass Disable** button to disable the hardware passthrough feature, which will block all traffic through the device during a power failure. This allows power failures to be detected and the traffic routed to an alternate device. When passthrough is disabled, the BYPASS DISABLED LED on the front panel is illuminated.
7. Check the LEDs on the front panel:

Front Panel LED	Description
POWER	Indicates that power is on.
FAULT	Indicates a system failure (hardware passthrough).
OPERATIONAL	Indicates normal operation. During a reboot, a system failure, or a power failure, the light is turned off, and traffic is passed through without any processing (hardware passthrough).
BYPASS DISABLE	Indicates that a power failure will block all traffic through the device (hardware passthrough disabled). To enable or disable hardware passthrough, press the <b>Bypass Disable</b> button on the back panel.

The installation is complete. You can now run Quick Setup, as described in “Running Quick Setup through the Web Console” on page 45.

## Installing the WX 100

The WX 100 and SR-100 are very similar, except that the latest model of the WX 100 has the network connectors on the front panel. The WX 100 can be used as a standalone device or as a server to distribute the processing load to up to six client devices. The client devices are connected directly to the WX 100, and can be WX 55s, WX 60s, WX 80s, or WXC 500s (all clients must be the same device type).

The following sections describe the installation process for the WX 100.

- Hardware Installation on page 38
- Disconnecting Power from the WX 100 on page 40
- Configuring Network Settings on page 41
- Connecting Client Devices to the Server on page 43
- Disconnecting Client Devices from the Server on page 44

## Hardware Installation

To install the WX 100 in your network:

1. Set up the chassis.
  - To install the WX 100 in a 19-inch device rack, install the supplied brackets (front panel forward) to the sides of the device with the screws provided in the kit. Next, install the chassis in your network device rack. Leave adequate space to install additional client devices.
  - To install the WX 100 on a desktop, place the chassis upside down on a smooth, flat surface, and install the supplied rubber feet on the bottom of the chassis. Place the chassis on a desktop or on top of another device so that all four rubber feet are securely mounted to the flat surface.

The subsequent steps depend on whether the WX 100 has standard copper-wire or fiber-optic interfaces.

### Copper-wire Interfaces


2. Connect the network cables and verify connectivity.

The standard WX 100 has two 10/100/1000 auto-sensing, Ethernet interfaces (Figure 10).



**NOTE:** Do not connect power to the device until Step 4.

**Figure 10: WX 100 Ethernet Ports on Front Panel**



To connect the network cables to the WX 100:

- a. Locate the cable that connects the switch (or other aggregating device) to the router.
- b. Disconnect this cable from the router port and connect it to the WX 100's LOCAL port.
- c. Using a straight-through cable (not provided), connect one end to the WX 100's REMOTE port and the other end to the router port.

3. Use one of the following methods to verify connectivity across the WX 100 when the power is off. This step ensures that the correct cables are used and that traffic will bypass (pass through) the WX 100 in the event of a power loss.
  - Ping a host on the remote side of the WX 100 from a host on the local side of the WX 100.
  - Observe the link status LEDs (if available) on the interfaces of the adjacent network devices (switch and router).
4. After you verify network connectivity across the WX 100, connect the supplied power cords to the dual power supplies on the back of the chassis, and then connect the power cords to the local power source.

The WX 100 maximum power usage is 300 Watts or 1025 BTU/hour.



**WARNING:** The appliance is designed to work with IT power systems.

**Waarschuwing** Het apparaat is ontworpen om te functioneren met IT energiesystemen.

**Varoitus** Koje on suunniteltu toimimaan IT-sähkövoimajärjestelmien yhteydessä.

**Attention** Ce dispositif a été conçu pour fonctionner avec des systèmes d'alimentation IT.

**Warnung** Das Gerät ist für die Verwendung mit IT-Stromsystemen ausgelegt.

**Avvertenza** Il dispositivo è stato progettato per l'uso con sistemi di alimentazione IT.

**Advarsel** Utstyret er utfomet til bruk med IT-strømsystemer.

**Aviso** O dispositivo foi criado para operar com sistemas de corrente IT.

**¡Atención!** El equipo está diseñado para trabajar con sistemas de alimentación tipo IT.

**Varning!** Enheten är konstruerad för användning tillsammans med elkraftssystem av IT-typ.

---

Now that the WX 100 is installed and powered on, configure the network settings, as described in “Configuring Network Settings” on this page.

### Fiber-optic Interfaces

The fiber-optic WX 100 is installed in the same way as the copper-wire version, except that you can apply the power before you connect the cables. Fiber-optic technology does not support a hard-wire passthrough connectivity in the event of a power loss.



**NOTE:** The fiber-optic WX 100 should be installed in a high-availability environment. Data transmission stops during a reboot or a power failure.

---

## Disconnecting Power from the WX 100



**WARNING:** The appliance has more than one power supply connection. All connections must be removed completely to remove power from the unit completely.

**Waarschuwing** Deze eenheid heeft meer dan één stroomtoevoerverbinding; alle verbindingen moeten volledig worden verwijderd om de stroom van deze eenheid volledig te verwijderen.

**Varoitus** Tässä laitteessa on useampia virtälähdekytkentöjä. Kaikki kytkennät on irrotettava kokonaan, jotta virta poistettaisiin täysin laitteesta.

**Attention** Cette unité est équipée de plusieurs raccordements d'alimentation. Pour supprimer tout courant électrique de l'unité, tous les cordons d'alimentation doivent être débranchés.

**Warnung** Diese Einheit verfügt über mehr als einen Stromanschluß; um Strom gänzlich von der Einheit fernzuhalten, müssen alle Stromzufuhren abgetrennt sein.

**Avvertenza** Questa unità ha più di una connessione per alimentatore elettrico; tutte le connessioni devono essere completamente rimosse per togliere l'elettricità dall'unità.

**Advarsel** Denne enheten har mer enn én strømtilkobling. Alle tilkoblinger må kobles helt fra for å eliminere strøm fra enheten.

**Aviso** Este dispositivo possui mais do que uma conexão de fonte de alimentação de energia; para poder remover a fonte de alimentação de energia, deverão ser desconectadas todas as conexões existentes.

**¡Atención!** Esta unidad tiene más de una conexión de suministros de alimentación; para eliminar la alimentación por completo, deben desconectarse completamente todas las conexiones.

**Varning!** Denna enhet har mer än en strömförsörjningsanslutning; alla anslutningar måste vara helt avlägsnade innan strömtillförseln till enheten är fullständigt bruten.



## Configuring Network Settings

The configuration can be downloaded automatically from CMS when you first apply power to the device (refer to “Manual and Automatic Installations” on page 30). To manually configure the network settings for the WX 100, you must connect an ANSI compatible terminal to the serial console port, and use a terminal emulation program (such as HyperTerminal) to enter the CLI commands described here.



**NOTE:** The serial console port is of type RS-232 (AT-compatible) with a male, DB-9 connector. You should use a female/female DB-9 crossover cable (such as a null-modem cable) when connecting directly to a PC serial port. The pin-outs for the console port are shown in “DB9 Console Port Pin-Outs” on page 395.

1. Connect an ANSI compatible terminal to the serial port on the front of the WX 100 (Figure 10 on page 38).
2. Verify the serial port settings are as follows:
  - Baud rate: 9600 bps
  - Data bits: 8
  - Parity: none
  - Stop bits: 1
  - Flow control: none
3. Start the terminal emulation program (such as HyperTerminal), and choose to connect via the serial port. The device will attempt to auto-deploy by downloading a configuration from CMS. To configure the device manually, press Enter.
4. At the User name and Password prompts, type **admin** for the user name and **peribit** for the password.



**NOTE:** This is a factory-configured password for the device. You will be asked to change the default password during the Quick Setup.

5. Press Enter and enter the following network information at the prompts:
  - a. Type an IP address for the device, and then press Enter.
  - b. Type the subnet mask for the network, and then press Enter.
  - c. Type the default gateway address for the device, and then press Enter.

The default gateway is typically the next hop on the Remote side of the WX 100. You may want to change the default gateway if you designate the device as a Default Assembler. After installing the WX 100, refer to “Defining Default Assemblers” on page 141 for more information.

Press Enter to confirm the network settings.

6. By default, the Local and Remote interfaces are set to auto-negotiate the speed and duplex mode. However, to avoid problems when the switch or router speed and duplex mode are set manually, it is **strongly recommended** that you manually configure the Local and Remote interface settings.


To manually configure the interface settings:

- a. At the prompt to configure the interface settings, type “y” and press Enter.
- b. Enter a number (0 to 5) for the speed and mode of the Local interface (fiber-optic interfaces have only the “auto” and “1000-full” options).
  - 0 - 10-full
  - 1 - 10-half
  - 2 - 100-full
  - 3 - 100-half
  - 4 - 1000-full
  - 5 - auto

Press Enter to confirm the setting, and then repeat for the Remote interface.

7. You can continue the Quick Setup or just press Enter at each prompt, and later run Quick Setup from the Web console. Note that the last prompt is to save the configuration as *startup.cfg*, which is used when you reboot the device.
8. Check the LEDs below the Ethernet ports (Figure 11). Note that the fiber-optic version has one LED for each interface to indicate 1 Gbps.

**Figure 11: Checking the Link LEDs for the WX 100**



- The LINK LEDs indicate the port is connected properly.
- The 100 and 1000 LEDs indicate the interface speed in Mbps.
- If the 100 and 1000 LEDs are off, the port is running at 10 Mbps.
- If all four 100 and 1000 LEDs are on, the device is in passthrough mode.

9. If you install the copper-wire WX 100 in a high-availability environment, you can disable hardware passthrough (refer to the “embed” CLI command on page 291), which will block all traffic through the device during a power failure. This allows power failures to be detected and the traffic routed to an alternate device.

The installation is complete. You can now run Quick Setup, as described in “Running Quick Setup through the Web Console” on page 45.



**NOTE:** After you run Quick Setup, you can connect client devices to the WX 100 (refer to “Connecting Client Devices to the Server” on page 43). However, to run the WX 100 in standalone mode (no clients), you must roll back to the previous release. WXOS 5.3 is supported only on WX 100 servers with one or more clients (stack mode).

## Connecting Client Devices to the Server

After you install the WX 100 or SR-100 and run Quick Setup, you can add more processing capacity by connecting up to six WX 55s, WX 60s, WX80s, or WXC 500s as client devices (all clients must be the same device type).

The following table shows the number of WX 60 or WXC 500 clients that are recommended for the licensed speed of the WX 100/SR-100.

WX 100 Speed	WX 60	WXC 500
Up to 45 Mbps	2	2
Up to 60 Mbps	3	3
Up to 80 Mbps	4	4
Up to 100 Mbps	5	5
Up to 125 Mbps	6	6
Up to 155 Mbps	6	6

To connect one or more client devices to the server:

1. If you are setting up the server and client devices for the first time, disconnect the WX 100 or SR-100 server from the network.
2. If necessary, upgrade the client devices to WXOS 5.0 or later, specify the IP information, and run Quick Setup. Note that if you use a console connected to the device, you can enter any IP address and gateway (the server will change them to internal addresses).
3. Log in to each client device and enter the following commands:
 

```
config stack-group set client-mode on
commit
save-config
reboot
```

 Type “y” to confirm the save and the reboot.
4. Mount the client devices near the server.

5. On each client device, connect a straight-through cable from the LOCAL port on the client to one of the ports numbered 1 to 6 on the WX 100 (see Figure 10 on page 38) or SR-100 (see Figure 7 on page 32). The port number becomes the client ID, and is shown on the client's front panel.
6. Plug in the supplied power cord to the back of each client, and then connect the power cord to the local power source.

The server assigns an IP address to each client. The client addresses are internal, so the clients cannot be accessed by other devices. If the WXOS version on the client and server are not the same, the client downloads the WXOS image from the server.

7. If the client devices are WXC 500s, configure the WX 100 or SR-100 as follows:
  - a. Use the following CLI commands to enable support for WXC devices:
 

```
config stack-group set sequence-mirror-server on
commit
save-config
```
  - b. Enable Network Sequence Caching (NSC) for the remote WXC devices (refer to “Configuring Network Sequence Caching” on page 134).
8. If you disconnected the server from the network in Step 1, reconnect the server now. If you added a client to a server that is already running with one or more clients, restart reduction and assembly on the server to ensure that tunnels are distributed across all clients. Click REDUCTION, clear the two check boxes at the top of the Endpoints page, click Submit, and then reselect the two check boxes, and click Submit again.

The client configuration is complete. A client device can be accessed only through the command console (no Web or SSH interface). On the WX 100 or SR-100 server, the number of client devices is shown in the banner of the Web console (unless hidden by a license expiration warning). The reduction Endpoints page on the WX 100 indicates which tunnels are handled by each client (refer to “Configuring Endpoints for Reduction Tunnels” on page 129).

### ***Disconnecting Client Devices from the Server***

To return a client device to stand-alone operation, you must reload the factory default settings:

1. Disconnect the cable from the client device to the WX 100 or SR-100 server.
2. Reload the factory default settings from the front panel of the client device. Alternatively, connect a terminal to the console port, log in, and enter the following command:

```
load-config factory-default
```

When the factory defaults are reloaded, unplug the power cable from the back of the client, plug the cable back in, and then specify the IP address, subnet mask, and default gateway for the device.

3. If you disconnect all clients from the server, and you want to run the server in standalone mode (no clients), enter the following CLI command to allow tunnels to be hosted on the server:

```
config stack-group set host-session server-only
```

If you disconnected all WXC 500 clients, enter the following command to disable the “sequence-mirror-server” mode:

```
config stack-group set sequence-mirror-server off
```

## Running Quick Setup through the Web Console

---

After starting the WX 100 and configuring network settings, the next step is to run the Quick Setup program. The first time you log in to the Web console, the Quick Setup program starts automatically and guides you through the required configuration steps. All settings made during Quick Setup can be changed later.

You can log in to the WXOS Web console from any workstation in your network. Data is securely transmitted through HTTPS. The WXOS Web console has the following requirements:

- Microsoft Internet Explorer browser version 6.0 or later
- Monitor display settings of 1024 x 768 or higher
- A Java Virtual Machine (JVM) version 5.0.0.3802 or later. To check your JVM version, open a command prompt and type "jview". If reports in the Monitor page are blank, or the graphs are not displayed correctly, go to <http://www.java.com> to install the latest Java Runtime Environment (JRE), which contains the JVM.

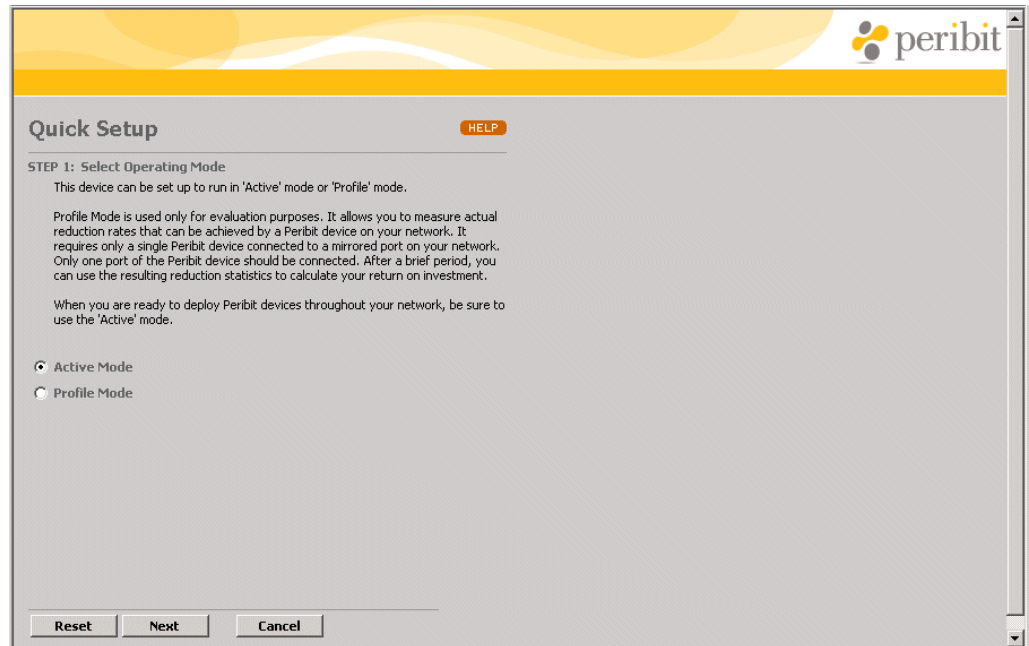
To run Quick Setup from the Web console:

1. Verify that the browser accepts cookies (required to log in), and that the server is always checked for the latest configuration information:
  - a. Select Tools > Internet Options.
  - b. Click Settings under Temporary Internet Files, select Every visit to the page, and click OK.
  - c. Click the Privacy tab and verify that the setting is Medium High or lower.
  - d. Click the Security tab, click Default Level, and verify that the setting is Medium or lower.
2. Enter the following URL in the browser:
 

```
https:// < IP address of the device >
```
3. Depending on your browser settings, the Security Alert dialog box may appear, click Yes to proceed.

- In the Login page, type **admin** for the user name and **peribit** for the password, and click Login. You will be asked to change the default password at the end of Quick Setup.

**Figure 12: Select Active or Profile Mode**



- Select the operating mode of the device:

Active Mode	Active operation where the device can reduce data, accelerate TCP applications, and manage WAN bandwidth.
Profile Mode	<p>Passive operation where the device can calculate potential reduction and acceleration results for all traffic, individual applications, and specific remote subnets. You can view the statistics on the standard reports. The actual traffic is not affected.</p> <p><b>Note:</b> To use Profile Mode, the device's Local interface must be connected to a mirrored port on the switch, and the Remote interface must be disconnected. Do not enable Profile Mode if the device is installed in the data path.</p> <p>For more information about setting up and using a device in Profile Mode, refer to "Profile Mode" on page 417.</p>



**CAUTION:** Enabling Profile Mode disables the Remote interface. If the device is installed in the data path, all transmission through the device will stop.

If you select Profile Mode, only the pages shown in Steps 7, 9, and 11 are displayed.

- Click Next to open the Registration Server Setup page (Figure 13).

**Figure 13: Registration Server Setup**

The screenshot shows the 'Quick Setup' window for Peribit. The title bar includes the Peribit logo. The main heading is 'Quick Setup' with a 'HELP' button. Below this, it says 'STEP 2: Set up a registration server'. A paragraph explains that Peribit devices exchange information through a designated registration server and must provide a password. It also notes that if the device is the registration server, a password must be set, and if it's not, the server's address and password must be entered. There are two radio button options: 'Designate this device as the registration server' (selected) and 'Direct this device to an existing registration server'. The first option has fields for 'Set password' and 'Verify password'. The second option has fields for 'Registration server' (labeled 'IP address') and 'Enter password'. There is also a checkbox for 'Delay reduction tunnel formation' with a descriptive text below it. At the bottom are 'Back', 'Next', and 'Cancel' buttons.

In Active Mode, one device must be designated as a registration server. Each device periodically contacts the registration server to find the other devices in the same community. Two devices can reduce and assemble data for each other only if they belong to the same community. Initially, all devices reporting to the same registration server are in the Default community. For more information, refer to “Configuring Registration Servers and Communities” on page 75.

To specify the registration server, do one of the following:

- a. Select Designate this device as the registration server, and enter a registration server password in both fields. The password is used to authenticate each device, and should be different from the administrator password.
- b. Select Direct this device to an existing registration server, and enter the IP address and password of the current (or future) registration server.




**NOTE:** If the registration server has not been installed, you can enter its IP address and password in advance.

Optionally, select Delay reduction tunnel formation to prevent this device from reducing or assembling data until you enable reduction tunnel formation. For more information, refer to “Configuring Endpoints for Reduction Tunnels” on page 129.

7. Click Next to open the Time Setup page (Figure 14).

### Figure 14: Set the Time



## Quick Setup

HELP

---

STEP 3 : Set the time

The time must be set on the Peribit device in order to display the correct time in reports. If you have access to a time (NTP) server, the Peribit device can be synchronized with it. Otherwise, you can enter the local time manually.

☒ Use NTP Server

Primary:  IP address

Secondary:  IP address (optional)

☐ Enter Local Time:

Time:  HH:MM ☐ AM ☒ PM

Date:  MM/DD/YYYY

Time Zone:

Daylight Saving: ☐ Automatically adjust time for daylight saving

Back

Next

Cancel

Do the following:

- a. Enter the IP address of your NTP server in the Primary field (a secondary NTP server is optional) or select Enter Local Time and enter the current date and time.
  - b. Select the local time zone, and select the Daylight Saving checkbox (if applicable).
8. Click Next to open the Features/Topology page (Figure 15).



**Figure 15: Specify the Feature Set and Community Topology**

**Quick Setup** HELP

STEP 4: Specify Features/Topology

A: Indicate the features that will be enabled on this device.

☒ All features  
☐ All features except Application Flow Acceleration  
☐ All features except Application Flow Acceleration and Active Flow Pipelining (AFP)

B: Indicate the topology in which this device will be deployed.

☒ Hub      Number of associated spokes      5 : Up to 300  
☐ Spoke  
☐ Mesh      Number of meshed devices      0 : Up to 60

C: Enter your adjacent router's WAN circuit speed      Kbps

Back Next Cancel

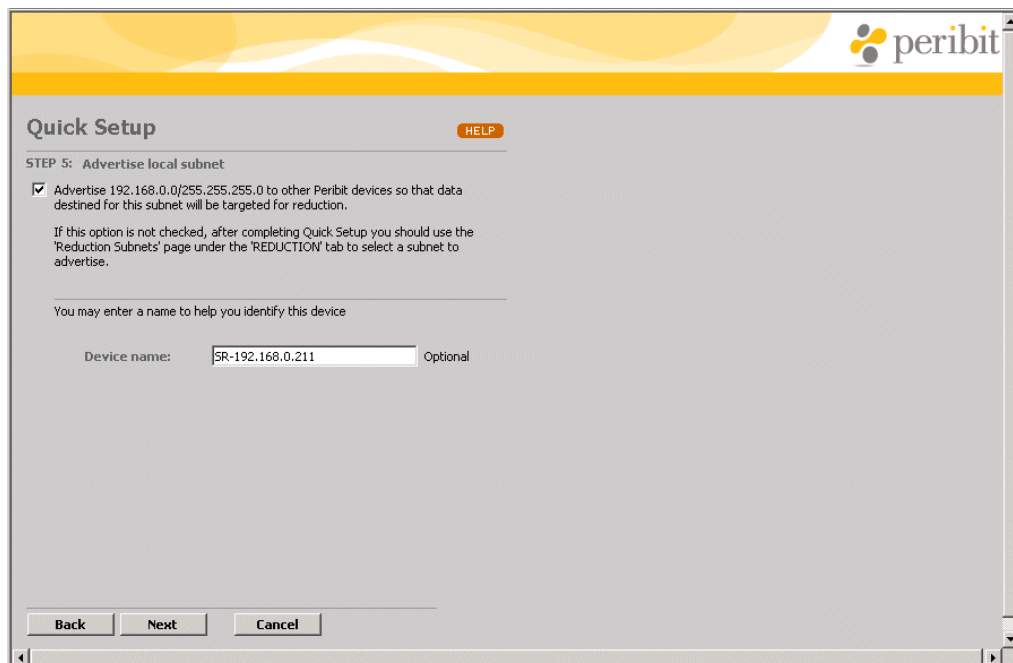
WXOS 5.3 supports only two topology ranges: Hub range 5 and Mesh range 0, both with all features enabled.

In Active mode, specify the following:

- a. **Features.** Select “All features” (the default).
- b. **Topology.** Select Hub with range 5 (the default) or Mesh with range 0. Use Mesh range 0 in a point to point topology where two WX 100 servers communicate with each other.
- c. **WAN circuit speed.** Enter the sum of the speeds for all the outbound circuits on the adjacent router that carry traffic from the WX 100 to the WAN. If the speed is variable, enter the maximum circuit speed (do not exceed the Ethernet speed of the device).

The WAN circuit speed is advertised to the remote WX devices in the community. The advertised speed becomes the maximum circuit speed for this endpoint in the outbound QoS configuration of the remote WXs. Outbound QoS, congestion control, and Application Flow Pipelining (AFP) are enabled automatically on the WX 100.

9. Click Next to open the Advertise Local Subnet page (Figure 16).

**Figure 16: Advertise Local Subnet and Enter a Device Name**

**Quick Setup** HELP

STEP 5: Advertise local subnet

☒ Advertise 192.168.0.0/255.255.255.0 to other Peribit devices so that data destined for this subnet will be targeted for reduction.

If this option is not checked, after completing Quick Setup you should use the 'Reduction Subnets' page under the 'REDUCTION' tab to select a subnet to advertise.

You may enter a name to help you identify this device

Device name:  Optional

Back Next Cancel

By default, the local subnet where the device is installed is advertised to the other devices in the community. As a result, data destined for this subnet is reduced by other devices in the path. Clear the check box to delay advertising the local subnet. To advertise the local subnet after you complete the Quick Setup, refer to “Advertising Reduction Subnets” on page 131.

Optionally, enter a unique name for the device (up to 30 characters) in the Device name field.



**NOTE:** If you plan to use CMS to manage your devices, entering a unique name for each device is strongly recommended.

10. Click Next to open the License Key page (Figure 17).

**Figure 17: Entering a License Key**

**Quick Setup** HELP

**STEP 6: License key**

The maximum throughput of this device is determined by the license key.  
A license key can be obtained by calling Peribit Networks at 866-737-4248 (866-PERIBIT).  
Please be prepared to provide the product serial number shown below.

Serial number:

License key:

If you have an Internet connection, you can take advantage of Peribit Networks' Online License Service by clicking this button. Online Service...

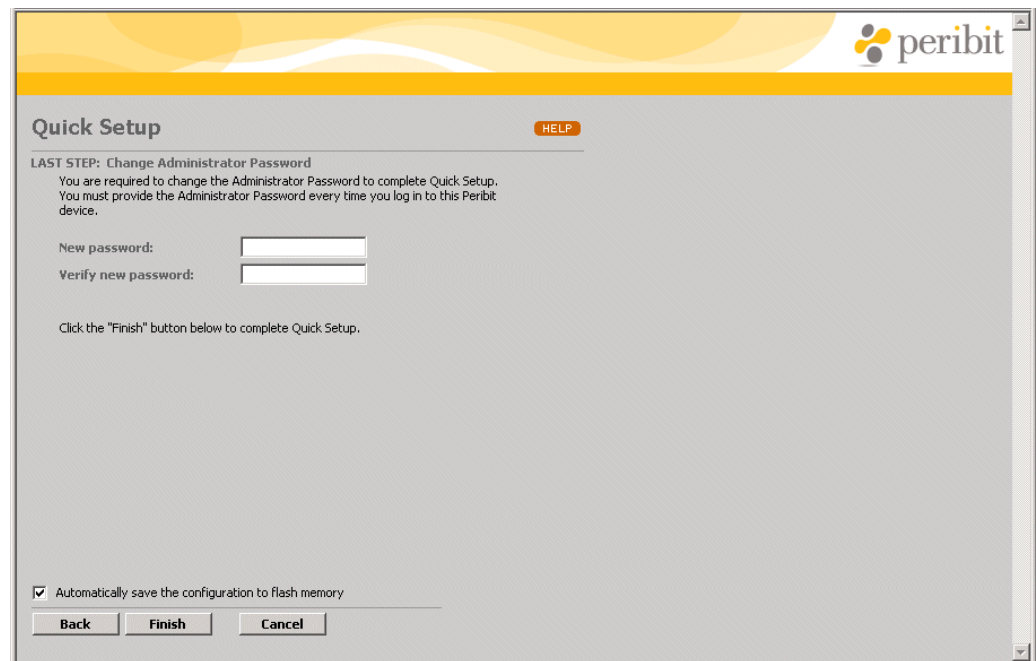
If you do not enter a license key, this device will operate normally for 30 days. It will then stop assembling/reducing data if a valid license key has not been entered.

Back Next Cancel

In Active Mode, each device requires a license key, which is based on its serial number. By default, each device has a 30-day evaluation license. When the evaluation license expires, data will pass through the device without reduction.

Serial number	If the serial number is not shown, get the "S/N" number from the back of the device.
License key	Enter your permanent license key or leave this field blank to use the 30-day evaluation license.

- Click Next to open the Change Administrator Password page (Figure 18).

**Figure 18: Change Administrator Password**

The screenshot shows a web browser window with the Peribit logo in the top right corner. The main heading is "Quick Setup" with a "HELP" button to its right. Below this, the text reads: "LAST STEP: Change Administrator Password. You are required to change the Administrator Password to complete Quick Setup. You must provide the Administrator Password every time you log in to this Peribit device." There are two input fields: "New password:" and "Verify new password:". Below these fields, it says: "Click the 'Finish' button below to complete Quick Setup." At the bottom, there is a checkbox labeled "Automatically save the configuration to flash memory" which is checked. Below the checkbox are three buttons: "Back", "Finish", and "Cancel".

Enter a new Administrator password in the New password and Verify new password fields, and then click Finish. You will be prompted to log in using the new password.



**NOTE:** If you deselect the “automatic save” option, the configuration settings will be lost the next time you restart the device, and Quick Setup will have to be run again. You can save the configuration later by clicking SAVE in the menu frame of the Web console.

Initial configuration is complete. Refer to the next section for a list of key configuration tasks.

## Post-Installation Tasks

---

After you run Quick Setup, you can continue configuring the device through the Web console or through the command line interface (CLI).

- To use the Web console, refer to “Configuring Basic Setup Policies” on page 55.
- To use CLI, refer to “Using the Command Line Interface (CLI)” on page 283.

Be sure to review the following key configuration tasks. The references are to instructions for using the Web console.

- Configure the local routes for the device, as described in “Configuring Local Routes” on page 66. To use RIP and/or OSPF to discover routes, you need the following information for your network:
  - OSPF Area ID, and the password or the MD5 authentication key and key ID
  - RIP password (if any)
- Select the local subnets that you want to advertise to other devices for data reduction, as described in “Advertising Reduction Subnets” on page 131.
- Review the available security features, such as limiting operator access to specific IP addresses or subnets, as described in “Configuring AAA” on page 79.
- Review the application definitions provided and add any new ones needed for your network, as described in “Managing Applications” on page 88.
- Configure inbound and outbound bandwidth management, as described in “Applying Quality of Service (QoS) Policies” on page 153.
- Enable traffic acceleration, as described in “Accelerating WAN Traffic” on page 189.

## Where to Go Next

---

After installing an WX or WXC device and running Quick Setup, proceed to one of the following chapters depending on your preference for configuring the device:

- Configuring Basic Setup Policies on page 55.
- Using the Command Line Interface (CLI) on page 283.



## Chapter 3

# Configuring Basic Setup Policies

The following topics describe the basic setup procedures:

- Using the Web Console on page 55
- Configuring Basic Setup Policies on page 57
- Configuring AAA on page 79
- Managing Applications on page 88



---

**NOTE:** You can also set up WX and WXC devices through the Command Line Interface (CLI). Refer to “Using the Command Line Interface (CLI)” on page 283 for more information.

---

## Using the Web Console

---

The WXOS Web console is a portal for accessing and configuring WX and WXC devices. Using the Web console, you can log in to a device from anywhere in your network and securely access configuration and management information, as well as reduction, acceleration, and QoS statistics.

The WXOS Web console supports the Microsoft Internet Explorer browser, version 6.0 and later (browser privacy settings must be configured to accept cookies). The WXOS Web console is designed to be viewed at 1024 x 768 pixels. To ensure secure transmission of configuration and management data, the WXOS Web console uses the Secure Sockets Layer protocol (SSL/HTTPS).

## Logging In

To log in to a device through the WXOS Web console:

1. Using a supported Web browser, enter the IP address of a WX or WXC device as follows:  
  
`https://<IP address of a device>`
2. Depending on your browser settings, a Security Alert dialog box may appear, click Yes to proceed.
3. In the Enter Network Password dialog box, type your user name and password.



**NOTE:** When a new device is accessed for the first time, use admin and peribit for the user name and password, and run Quick Setup (refer to “Running Quick Setup through the Web Console” on page 45).

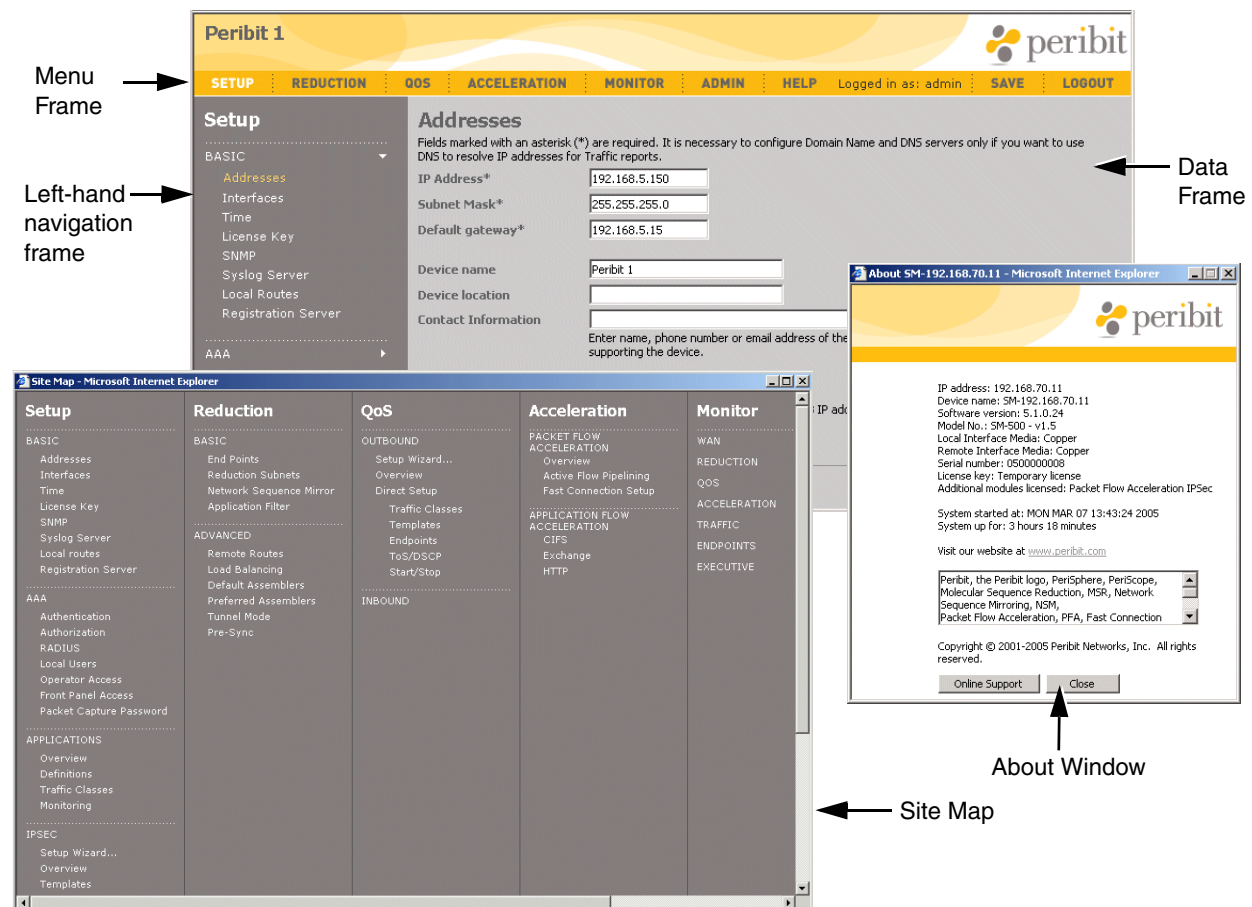
- Continue to the next section for a description of the WXOS Web console interface.

To log out of the WXOS 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).

## Understanding the WXOS Web Console Interface

The WXOS Web console contains a menu frame of seven administrative functions, a left-hand navigation frame of various sub-menu items, and a data frame for configuring and viewing policies and performance data.

**Figure 19: WXOS Web Console Interface**



Click HELP > About to view hardware and software information for the device, such as the IP address, the software and hardware versions, and the license key assigned to the device. Click HELP > Site Map to view a list of the options available under each menu frame selection.



## Using Special Characters

In general, use only letters, numbers, and blanks when assigning names to devices and other objects. If necessary, you can also use the following special characters:

# \$ & \_ - + . ( ) ' ,



**NOTE:** You can also use colons (:), but not in device names.

## Configuring Basic Setup Policies

The following topics describe the basic configuration procedures:

- “Configuring Device Address and Contact Information” on page 57
- “Configuring the Interface Settings” on page 59
- “Configuring 802.1Q VLAN Support” on page 61
- “Configuring Time Settings” on page 62
- “Obtaining a Permanent License” on page 62
- “Enabling SNMP” on page 64
- “Enabling Syslog Reporting” on page 65
- “Configuring Local Routes” on page 66
- “Configuring Registration Servers and Communities” on page 75

## Configuring Device Address and Contact Information

The device’s IP address, subnet mask, and default gateway are specified during the installation process. The Addresses page of the Web console lets you change these settings, as well as add device and administrator contact information, and specify DNS servers used to resolve IP addresses on the Traffic report.

To change the network address and contact information:

1. Click SETUP in the menu frame.

**Figure 20: Configuring Network Address and Contact Information**

## 2. Specify the following information:

IP address	Enter the IP address of the device. <b>NOTE:</b> If you change the IP address or subnet mask, you must reboot the device. If this device is also a registration server, you must first transfer the registration server to another device before changing the IP address (refer to “Configuring Registration Servers and Communities” on page 75).
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 WX or WXC device).
Device name	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.  A device name change is propagated to the other WX and WXC devices in the community the next time the device checks in with the registration server.

## 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

Enter the local DNS domain name of the WX or WXC device (up to 256 characters). The domain name must include at least one period, but not as the first or last character.  
  
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.  
  
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.
- 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 (refer to “Traffic Statistics” on page 256).
4.

Click Submit to activate the changes, or click Reset to discard them.
5.

To retain your changes when the device is restarted, click SAVE in the menu frame.

Configuring the Interface Settings

Each WX and WXC device has two Network Interface Controllers (NICs) for its Local and Remote interfaces. By default, these interfaces are set to auto-negotiate the link speed and mode (half- or full-duplex).



**NOTE:** The WX 100 has two 10/100/1000 NICs. The fiber WX 100 supports only 1 Gigabit speeds at full-duplex.


The Web console lets you do the following:

- View the status, MAC address, and negotiated speed and mode of each interface.
- Run a test to detect a mode mismatch on the Local or Remote interface, and manually configure the speed and mode when necessary.

By default, a passive test runs periodically and displays a message above the menu frame if a mismatch is detected (refer to “configure interface” on page 318). The passive test can detect a mismatch only when data is sent and received at the same time.

In addition, you can enable high-availability support so that when a failure is detected on one interface, the other interface is 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 21).

Figure 21: 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 copper-wire WX 100, you can also disable hardware passthrough so that the router detects the loss of traffic if the WX 100 fails (refer to the “embed” CLI command on page 291).

To view and/or configure the interface settings:

1. In the Setup page, click Interfaces in the left-hand navigation frame.

**Figure 22: Configuring Interface Speed and Duplex Mode Settings**

The Status fields show the current speed and mode parameters for each interface. By default, the Local and Remote interfaces are set to auto-negotiate. In addition, each interface's Media Access Control (MAC) address is listed.

2. To change the speed and mode for the Local or Remote interfaces, select Manual, and then choose 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 activate the changes, or click Reset to discard them.

5. To test for mismatched duplex settings between the WX or WXC device and another device, click Test Settings, select the Local or Remote interface, enter the IP address of any device on the selected interface segment, and click Submit. The test results are displayed in a popup window.
6. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring 802.1Q VLAN Support

WX and WXC devices support reduction of VLAN traffic that conforms to the IEEE 802.1Q specification. However, VLAN traffic that uses ISL encapsulation is passed through without reduction.

To enable reduction of 802.1Q VLAN traffic:

7. Click SETUP in the menu frame, and then click Interfaces in the left-hand navigation frame.
8. 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 WX or WXC device is installed.
  - VLAN ID. Enter a VLAN ID (1 through 4095) for the port where the Local interface of the WX or WXC 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.

Click Submit to activate the changes, or click Reset to discard them.

When a WX or WXC 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 page and cannot be advertised for reduction.

To include WAN-side routes on the list of reduction subnets, you must enter the following CLI commands:

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

After you configure the local routes (refer to “Configuring Local Routes” on page 66), verify that the appropriate subnets are discovered and advertised (refer to “Advertising Reduction Subnets” on page 131). Since both LAN and WAN-side subnets will be shown on the Reduction Subnets page, be sure to advertise only the true LAN-side subnets.

9. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring Time Settings

WX and WXC devices support the Network Time Protocol (NTP). If your network uses NTP, you can specify a primary and secondary NTP server to maintain the current time. You can also set the time manually. Entries in the system log files include the current time to assist with device administration.

To configure the time settings:

1. In the Setup page, click Time in the left-hand navigation frame.

**Figure 23: Configuring the Time Settings for a Device**

2. Do one of the following:
  - If you have an NTP server in your network, select Use NTP Server and enter the IP address of the NTP server in the Primary field. A secondary NTP server is optional.
  - If you do not have an NTP server, select Enter Local Time and enter the current time and date.
3. Select the time zone of the device, and then select Automatically adjust time for daylight savings if applicable.
4. Click Submit to activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Obtaining a Permanent License

Each non-backup device requires a permanent license key for operation. The license key determines the licensed modules and throughput for the device, and properly registers the product. Initially, each device has a temporary 30-day license with access to all features. When the temporary license expires, all traffic will pass through without reduction.

For backup devices, temporary licenses are sufficient because only the active device time is counted against the 30-day limit (WXOS 5.1 or later required).

To obtain a permanent license key, you need:

- Device serial number displayed in the License Key page (also displayed in the About box and on the back of the device)
- One or more Right To Use (RTU) keys that were emailed to you in a PDF file
- User ID and password to access the License Key server at:

[http://www.juniper.net/generate\\_license](http://www.juniper.net/generate_license)

If you do not have Internet access, please call Technical Support at +1-888-314-JTAC (U.S, Canada, and Mexico) or +1-408-745-9500.

The speed RTU key specifies the licensed speed and level of support for the device. A separate RTU is needed for each optional feature (such as IPSec encryption). If you do not enter an RTU key, the device is licensed for the base speed with no customer support. If you lose the license key, you can use the License Key server to retrieve your current license key.

To install a permanent license key:

1. In the Setup page, click License Key in the left-hand navigation frame.

**Figure 24: Replacing the Temporary License with a Permanent License**

The screenshot shows the 'Peribit 1' web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left-hand navigation frame is expanded to show the 'Setup' section, with 'License Key' selected. The main content area is titled 'License Key' and contains the following information:

- The maximum throughput through the reduction engine is determined by the license key.
- Current license key: Temporary license
- Maximum throughput licensed: Unlimited
- Additional modules licensed: Packet Flow Acceleration, IPSec
- Expiration: 28 days, 19 hours, 8 minutes

A note states: 'A license key can be obtained by calling Peribit Networks at 866-737-4248 (866-PERIBIT). Please be prepared to provide the product serial number shown below.'

Below this note, there are two input fields: 'Serial number:' with the value '0100000006' and 'License key:' which is empty. To the right of the 'License key:' field is a button labeled 'Online Service...'. At the bottom of the form are 'Submit' and 'Reset' buttons.

The License Key page displays the status of the current license, including the licensed modules and the maximum throughput for the device.

2. If you have obtained a registered license key, enter it the License key field. If you do not have a registered license key, you can obtain one as follows:
  - a. Click Online Service.

- b. Enter your contact information and the device serial number, and click Submit.
  - c. Enter an RTU key for the desired device speed and level of support, and click Yes. If you omit the RTU key and click No, the device is licensed for the base speed with no support.
  - d. Copy the displayed license key into the License key field in the Web console.
3. Click Submit to activate the changes, or click Reset to discard them.

## Enabling SNMP

WX and WXC devices provide the following SNMP support:

- 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 performance statistics from a Network Management System (NMS). In addition, the SNMP traps can be sent to the NMS and other network devices. For a description of the SNMP traps, refer to “SNMP Traps and Syslog Messages” on page 397.

To enable SNMP:

1. In the Setup page, click SNMP in the left-hand navigation frame.

**Figure 25: Enabling SNMP**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, OOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left-hand navigation pane is expanded to show the 'Setup' section, with 'SNMP' selected. The main content area is titled 'SNMP' and contains the following configuration options:

- SNMP Enabled:** A checkbox labeled 'Yes' is checked.
- Read Community String:** A text input field containing '\*\*\*\*\*'.
- Write Community String:** A text input field containing '\*\*\*\*\*'.
- Trap Enabled:** A checkbox labeled 'Yes' is unchecked.
- Trap Community String:** A text input field containing '\*\*\*\*\*'.
- Trap Destinations:** A text area for entering IP addresses, one per line. A tooltip indicates 'Enter IP addresses, one per line.'.
- Authentication Failure Trap Enabled:** A checkbox labeled 'Yes' is unchecked.

At the bottom of the configuration area are two buttons: 'Submit' and 'Reset'.



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 on the device. The default community strings 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 activate the changes, or click Reset to discard them.
6. To retain your changes when the device is restarted, click SAVE in the menu frame.

### Enabling Syslog Reporting

WX and WXC devices 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 interface status, security alerts, and environmental conditions. For a description of Syslog messages, refer to “SNMP Traps and Syslog Messages” on page 397.

To enable Syslog reporting:

1. In the Setup page, click Syslog Server in the left-hand navigation frame.

**Figure 26: Enabling Syslog Reporting**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left-hand navigation menu is expanded, showing 'Setup' as the active section. Under 'Setup', 'Syslog Server' is highlighted. The main content area is titled 'Syslog Server' and contains the following configuration options:

- Syslog enabled:** A checkbox labeled 'Yes' is checked.
- Syslog servers:** A text input field is empty. To its right, a note says: 'Enter IP addresses, one per line with a maximum of 5 servers.'
- Syslog message severity:** Three checkboxes are present: 'Critical' (checked), 'Error' (checked), and 'Informational' (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'.

2. Select the Syslog enabled 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 message, such as License expired.

- **Informational:** Informational messages, such as reload requests and low-process stack messages.
4. Click Submit to activate the changes, or click Reset to discard them.
  5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring Local Routes

Local routes are the routes defined in the device's routing table. When you first install a WX or WXC 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
- Add dynamic routes using one of the following methods:
  - Enable the Open Shortest Path First (OSPF) and/or the Routing Information Protocol (RIPv1, RIPv2)
  - Periodically poll the routing table of a Cisco router (not supported on off-path devices that use RIP for packet interception)
  - Import a file of routes from an FTP server

A total of 8192 IP routes (static and dynamic) are supported (the WX 15 is limited to 1000). Also, each device can balance the load across up to four routers that have equal cost paths to the same destination.

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 131). By default, only the subnets you select are advertised.



**NOTE:** In some environments, such as VLAN, some routes on the Local interface may be discovered only on the Remote interface. To advertise these subnets, you must enable the WAN reduction subnet option through the CLI to display all routes on the list of reduction subnets (refer to "configure reduction-subnet" on page 350).

---

To configure local routes:

1. In the Setup page, click Local Routes in the left-hand navigation frame.

**Figure 27: Configuring Local Network Routes**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The left-hand navigation frame is expanded to show the Setup section, with Local Routes highlighted. The main content area is titled 'Local Routes' and shows a table of 3 local routes defined on the device. The table has columns for IP Address, Subnet Mask, Gateway, and Route Type. Below the table is a checkbox labeled 'Static routes take precedence over dynamic routes'. At the bottom are buttons for Submit, Static..., Dynamic..., Import..., and Router Balancing....

IP Address	Subnet Mask	Gateway	Route Type
0.0.0.0	0.0.0.0	192.168.5.15	Static
127.0.0.1	0.0.0.0	127.0.0.1	Dynamic
192.168.5.0	255.255.255.0	192.168.5.150	Dynamic

☐ Static routes take precedence over dynamic routes

Buttons: Submit, Static..., Dynamic..., Import..., Router Balancing...

2. If you want static routes to take precedence over dynamically discovered routes, select the check box at the bottom of the page.



**NOTE:** The default route (0.0.0.0/0.0.0.0) and manually entered routes are static; all other routes are dynamic. The WX device is shown as the “gateway” for its local subnet, which is also labeled as dynamic.

3. To remove a manually-defined static route, click Delete next to the route. The route is also removed from the list of reduction subnets.
4. Refer to the following sections to add routes or enable router balancing:
  - “Adding Static Routes” in the next section
  - “Enabling RIP and OSPF Support” on page 69
  - “Enabling Route Polling” on page 70
  - “Importing a Routing Table” on page 71
  - “Enabling Route-Based Router Balancing” on page 73
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Adding Static Routes

To manually add static network routes:

1. On the Local Routes page, click Static.

**Figure 28: Adding a New Local Static Route**

2. Enter the IP address, subnet mask, and the IP address of the gateway to this subnet.
3. Click Submit to activate the new route. The route is added to the list of local routes and reduction subnets. Note that ICMP redirect routes take precedence over static routes.

New static routes are advertised automatically to other WX and WXC devices, except when the WAN reduction subnets option is enabled (refer to “Advertising Reduction Subnets” on page 131).

## Enabling RIP and OSPF Support

If your network uses OSPF or RIP, you can enable these protocols on the WX or WXC device so that local routes are discovered dynamically.



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

To enable RIP and/or OSPF routing:

1. On the Local Routes page, click Dynamic.

**Figure 29: Enabling Dynamic Routing**

2. Click Use OSPF/RIP.
3. To enable support for OSPF:
  - a. Click OSPF to configure the dynamic route settings.
  - b. On the Local routes > Dynamic > OSPF page, enter the Area ID for OSPF.
  - c. 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).
  - d. Click Submit on the Local routes > Dynamic > OSPF page.
  - e. On the Local routes > Dynamic page, select Start next to OSPF.
4. To enable support for RIP:
  - a. Click RIP to configure the dynamic route settings.
  - b. On the Local routes > Dynamic > RIP page, select the version of RIP used in your network (1 or 2).

- c. If your network uses RIP authentication, select Password and enter the password (up to 15 characters).
  - d. Click Submit on the Local routes > Dynamic > RIP page.
  - e. On the Local routes > Dynamic page, select Start next to RIP.
5. Click Submit to start the enabled protocols.

All discovered routes are added to the list of local routes. Routes discovered on the Local interface are added to the list of reduction subnets, but they are not advertised automatically to the other WX and WXC devices. If the WAN reduction subnets option is enabled, all routes discovered on the Remote interface are also added to the list of reduction subnets (refer to “Advertising Reduction Subnets” on page 131).

### Enabling Route Polling

Routes can be discovered by periodically polling a Cisco router on the same subnet. The router must be configured to allow Remote Shell (*rsh*) access by the WX device. The *rsh* protocol allows a user or device to execute commands on a remote system without having to log in. The BGP routes are included only if you enable the BGP option using the CLI (refer to “configure route-poll” on page 359).



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

### Configuring a Cisco Router for Route Polling

The following sample Cisco router commands enable Remote Shell access for the WX device at IP address 172.16.5.3. The local and remote user names are “juniper” and “wxdevice,” respectively. On the WX device, the names must be reversed (specify “juniper” as the remote name, and “wxdevice” as the local name).

```
config terminal
ip rcmd rsh-enable
ip rcmd remote-host juniper 172.16.5.63 wxdevice enable
no ip rcmd domain-lookup
end
```

### Configuring Route Polling

To periodically obtain the routing table from a Cisco router:

1. On the Local routes page, click Dynamic.
2. Click Router to configure the dynamic route import settings.

**Figure 30: Dynamically Obtaining a Routing Table from a Cisco Router**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left sidebar shows the 'Setup' menu with options like BASIC, AAA, and APPLICATIONS. The main content area is titled 'Local routes > Dynamic > Router'. It contains several configuration fields: 'Poll router' with 'IP address' and 'Port' (514), 'Secondary router' with 'IP address' and 'Port' (514), 'Local user name', 'Remote user name', and 'Polling interval' (5 minutes). There are also 'Submit', 'Reset', and 'Cancel' buttons at the bottom.

### 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). <b>NOTE:</b> The IP address must be on the same subnet as the WX 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

### 4. Click Submit to save the settings and return to the Local routes > Dynamic page.

### 5. Select Obtain routing table from router, and click Submit.

All discovered routes are added to the list of local routes. Routes discovered on the Local interface are added to the list of reduction subnets, but they are not advertised automatically to the other WX devices. If the WAN reduction subnets option is enabled, routes discovered on the Remote interface are also added to the list of reduction subnets (refer to “Advertising Reduction Subnets” on page 131).

## Importing a Routing Table

If you export a routing table from a Cisco router to a file, and then save the file to an FTP server, you can import the routes file to the WX device. The imported routes are the routes listed when you enter a “show ip route” command on the Cisco router. To import the routes file from a TFTP server, use the CLI command (refer to “import-route-table” on page 291).

The router must be in the same subnet as the WX device, and it is preferable to use the router that is connected to the Remote port. The following types of imported routes are recognized:

**B** - BGP routes, **C** - Connected routes, **D** - EIGRP routes, **E** - EGP derived, **I** - IGRP derived, **i** - IS-IS derived, **O** - OSPF derived, **R** - RIP derived, **S** - Static routes



**NOTE:** On an off-path device that uses RIP for packet interception, be careful not to import RIP routes that were advertised by the off-path device (traffic to those destinations will be dropped).

To import routes from an FTP server:

1. On the Local Routes page, click Import.

**Figure 31: Importing a Routing Table**

2. In the Import from FTP Server section, enter the IP address of the FTP server, the directory path and file name of the file, the user name and password for the FTP server, and the Cisco router's IP address.



**NOTE:** You should not import a routing table if dynamic routing is enabled (RIP, OSPF, or route polling).

3. Select Delete last imported file if you do not want to reload the file from Flash memory the next time the device is restarted.
4. Click Submit to import the file and store a copy of it in flash memory. You return to the Local Routes page.




All discovered routes are added to the list of local routes. Routes discovered on the Local interface are added to the list of reduction subnets, but they are not advertised automatically to the other WX devices. If the WAN reduction subnets option is enabled, routes discovered on the Remote interface are also added to the list of reduction subnets (refer to “Advertising Reduction Subnets” on page 131).

### Enabling Route-Based Router Balancing

To balance the reduced traffic load across multiple routers, you can configure the WX device to distribute traffic across equal-cost paths (route-based balancing) and/or configure the local router to distribute traffic based on ToS values set by the WX device (ToS marking for router-based balancing). To configure ToS marking for router-based balancing, refer to “configure route” on page 356.

Using route-based balancing, the device can distribute reduced traffic across up to four different gateways. In Figure 32, WX 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 32: Configuring Router Balancing Policies**

If two or more gateways (up to four) have equal cost paths to the same IP address, the routes are grouped together in the Local Routes page (Figure 33).

**Figure 33: Common Routes with Equal Cost Paths**

IP Address	Subnet Mask	Gateway	Route Type
0.0.0.0	0.0.0.0	192.168.53.130	Static
127.0.0.1	0.0.0.0	127.0.0.1	Dynamic
173.16.4.0	255.255.255.0	192.168.0.1	Dynamic
192.168.53.128	255.255.255.192	192.168.53.180	Dynamic

4 local routes are defined on this Peribit device

Static routes take precedence over dynamic routes

Submit Static... Dynamic... Import... Router Balancing...

OSPF: Stopped  
RIP: Stopped  
Router Polling: None

Equal cost paths to the same destination

To enable router balancing:

5. On the Local Routes page, click Router Balancing.

**Figure 34: Configuring Router Balancing**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

- Addresses
- Interfaces
- Time
- License Key
- SNMP
- Syslog Server
- Local Routes
- Registration Server

AAA

APPLICATIONS

- IPSEC

ADVANCED

**Local routes > Router balancing**

The rule selected below determines how traffic is directed when more than one gateway exists for a given subnet.

☒ Off All traffic is directed to one of the available routers.

☐ 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, i.e. round robin.

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

Submit Reset Cancel

6. 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.

7. Click Submit to activate the changes, or click Reset to discard them.

## Configuring Registration Servers and Communities

At least one WX 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 WX 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 (also called “endpoints”).

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 send reduced traffic between communities, you can create a hierarchical structure where selected devices reside in multiple communities (refer to “Configuring Tunnel Switching” on page 148).

Initially, all WX devices are in the Default community. The registration server can reside in any community, and in most cases only one registration server is required. Also, you can enable or disable data reduction between any two devices in the same community, as described in “Configuring Endpoints for Reduction Tunnels” on page 129.

The following sections describe how to configure registration servers and communities:

- “Defining Registration Servers and Passwords” in the next section
- “Defining Communities” on page 77

### Defining Registration Servers and Passwords

When you log in to a registration server, you can change the password of the registration server, assign devices to communities, or designate a different device as the registration server. You can also specify a secondary registration server to act as a backup if the primary server is unavailable. On all other WX and WXC devices, you can change only the primary registration server used by the device.

To configure registration server settings:

1. Log in to the device that acts as the registration server. For any other device, you can specify only the IP address and password of the registration server used by that device.
2. Click Setup in the menu frame, and then click Registration Server in the left-hand navigation frame.

**Figure 35: Configuring Registration Server Settings**

3. To change the password of the registration server, select **Change registration server password**, and then enter the old and new passwords in the appropriate fields.



**NOTE:** Changing the password disrupts communication with all WX devices that use the registration server. To restore communication with the registration server, you must update the registration server password on each WX device. If you have the Central Management System (CMS), you can schedule an update for all devices.

4. To designate a secondary registration server that acts as a backup should the primary fail, select **Change SECONDARY registration server**, select **Use IP address**, and then enter the IP address of the device. To remove a secondary registration server, select **No secondary registration server**.

Ideally, the primary and secondary registration servers should be located on a link with relatively high bandwidth and low congestion to facilitate communication between the two servers and the other WX devices.

5. To designate a different device as the registration server, select **Transfer registration server designation to another device**, and then enter the IP address of the device. All devices that used the old registration server are updated with the new address (same password).

If the primary registration server fails, you can promote the secondary server to be the primary registration server, as follows:

- a. Log in to the secondary registration server.
- b. Click **Registration Server** and enter the IP address of the secondary server in the **Registration server** field.

6. To retain your changes when the device is restarted, click **SAVE** in the menu frame. If you transfer the registration server to another device, the change is saved automatically.

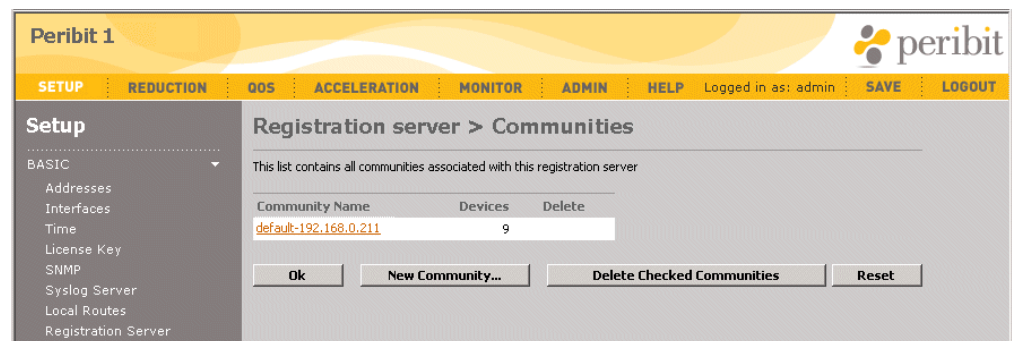
## Defining Communities

Data reduction occurs only between WX and WXC devices in the same community, so you can optimize performance by limiting the number of devices in each community. To send reduced traffic between communities, you can create hierarchical communities (refer to “Configuring Tunnel Switching” on page 148).

To configure the communities on a registration server:

1. Log in to the device that acts as the registration server.
2. Click **Setup** in the menu frame, click **Registration Server** in the left-hand navigation frame, and then **Communities**.

**Figure 36: Viewing the Communities for a Registration Server**



On the Registration server > Communities page you can:

- Add a community. Click **New Community**, enter a community name (up to 31 characters), and click **Submit**.
  - Delete a community. Click the check box next to the appropriate names and click **Delete Checked Communities**. The devices in a deleted community are moved to the Default community if they do not belong to any other user-defined communities.
3. To define the devices in a community, click the community name.

Figure 37: Viewing a List of Devices in a Community



The Registration server > Community Devices page lists all of the devices for the community selected at the top of the page. Note the following:

- The **Reduction** and **Assembly** columns indicate whether the reduction and assembly engines are activated. To activate or deactivate the reduction and assembly engines, refer to “Configuring Endpoints for Reduction Tunnels” on page 129.
- The **Duties** column can contain the following icons:

Duty icon	Description
	<b>Hub</b> — The device is designated as a Hub. Each device attempts to form a reduction tunnel with a hub before creating tunnels to other WX devices (refer to “Setting Community Feature and Topology Parameters” on page 99).
	<b>Spoke</b> — The device is designated as a spoke in a Hub and Spoke topology. By default, a spoke reduces and assembles data only for the hub device(s).
	<b>Mesh</b> — The device is designated as part of a mesh topology.
	<b>Registration Server</b> — Indicates that this device is the Registration Server for the community.
	<b>Secondary Registration Server</b> — Indicates that this device is the Secondary Registration Server for the community.
	<b>Backup and Backup (Active)</b> — The device is designated as backup for one or more primary devices. The icon flashes when the backup device is active. To configure a device as a backup, refer to “configure backup” on page 312.

- The **Last registration** column displays the date and time the device last contacted the registration server for configuration and policy information.

4. From the Community Devices page, you can:

- Add devices to a community. Click Add New Devices, select a community, enter the device IP addresses (one per line), and click Submit.

- Copy devices to another community (a device can belong to multiple communities):
    - a. Select the “copy from” community at the top of the page, and click the check box next to the appropriate devices,
    - b. Select the “copy to” community at the bottom of the page, and click Add Checked Devices.
  - Delete devices from a community. Select the community at the top of the page, click the check box next to the appropriate devices, and click Remove Checked Devices. Deleted devices are moved to the Default community if they do not belong to any other user-defined communities.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring AAA

---

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 WX devices 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.

A WX or WXC 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 topics describe how to define the authentication and authorization settings, RADIUS servers, and other security features:

- “Selecting Authentication Methods” in the next section



- “Enabling Authorization Checking” on page 81
- “Defining RADIUS Servers and Server Groups” on page 82
- “Defining Local Users” on page 84
- “Securing Operator Access” on page 86
- “Securing Front Panel Access” on page 87
- “Changing the Packet Capture Password” on page 88

## Selecting Authentication Methods

For each user interface—the Web, the SSH (CLI), and the console— you can specify the order in which the local database and RADIUS server groups are accessed to authenticate each user. 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 and Server Groups” on page 82. To define user accounts locally, refer to “Defining Local Users” on page 84.

To select the authentication methods for each user interface:

1. In the Setup page, click AAA in the left-hand navigation frame, and click Authentication.

**Figure 38: Selecting Authentication Methods**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

- Addresses
- Interfaces
- Time
- License Key
- SNMP
- Syslog Server
- Local Routes
- Registration Server

AAA

- Authentication**
- Authorization
- RADIUS
- Local Users
- Operator Access
- Front Panel Access
- Packet Capture Password

APPLICATIONS

- IPSEC

ADVANCED

**Authentication**

**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  failed attempts ☐ Never

**Web**

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

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 followed by the 'None' method.

Submit Reset



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> <p><b>RADIUS: <i>group_name</i>.</b> 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.</p> <p><b>Local.</b> Attempts to authenticate users locally.</p> <p><b>None.</b> Login not required. Can be used alone or after the last RADIUS group. Cannot be used directly after Local.</p> <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 <b>Never</b>.</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 activate the changes.

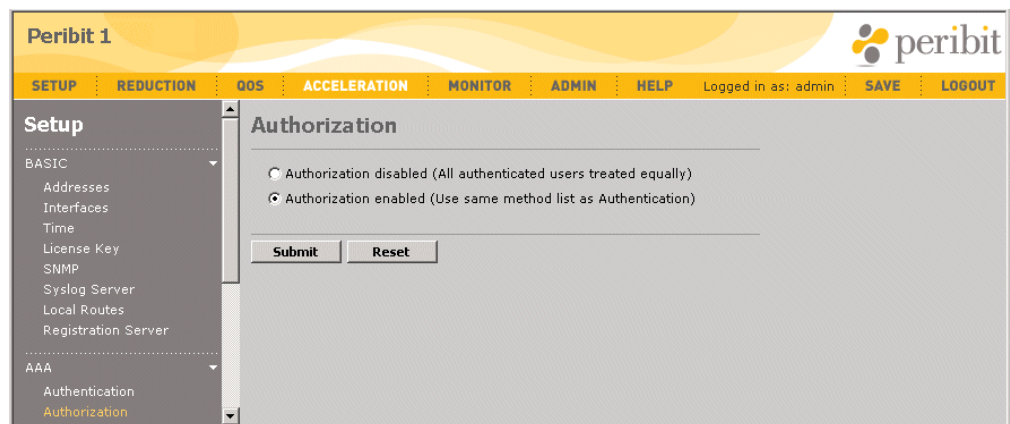
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

### 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 Setup page, click AAA in the left-hand navigation frame, and then click Authorization.

**Figure 39: Enabling Authorization Checking**

2. Select one of the following, and click Submit.
  - **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. To retain your changes when the device is restarted, click SAVE in the menu frame.

### **Defining RADIUS Servers and Server Groups**

To use RADIUS servers to authenticate users, you must define one or more RADIUS servers and assign 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 80.

To define RADIUS servers and server groups:

1. In the Setup page, click AAA in the left-hand navigation frame, and then click RADIUS.

**Figure 40: Defining RADIUS Servers and Server Groups**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

- Addresses
- Interfaces
- Time
- License Key
- SNMP
- Syslog Server
- Local Routes
- Registration Server

AAA

- Authentication
- Authorization
- RADIUS**
- Local Users
- Operator Access
- Front Panel Access
- Packet Capture Password

**RADIUS**

RADIUS Client

Source IP Address: 192.168.0.211

RADIUS Servers	IP Address	Auth. Port	Time-out (sec)	Re-transmit	Dead Time (min)	Delete
Main	10.20.30.100	1812	3	3	0	<input type="checkbox"/>

New Server...

RADIUS Server Groups

Central ☐

New Group...

Submit Reset

From the RADIUS page, you can:

- Add new servers and assign them to groups, as described in Step 2 through Step 4.
- Change a server or server group. Click the server or group name, make any needed changes, and click Submit.
- Change the IP address in the **Source IP Address** field (defaults to the device's IP address). Replies from the RADIUS server are sent to the source address.
- 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 and specify the following information:

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 WX 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 and specify the following information:

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.

4. Click Submit to activate the changes.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

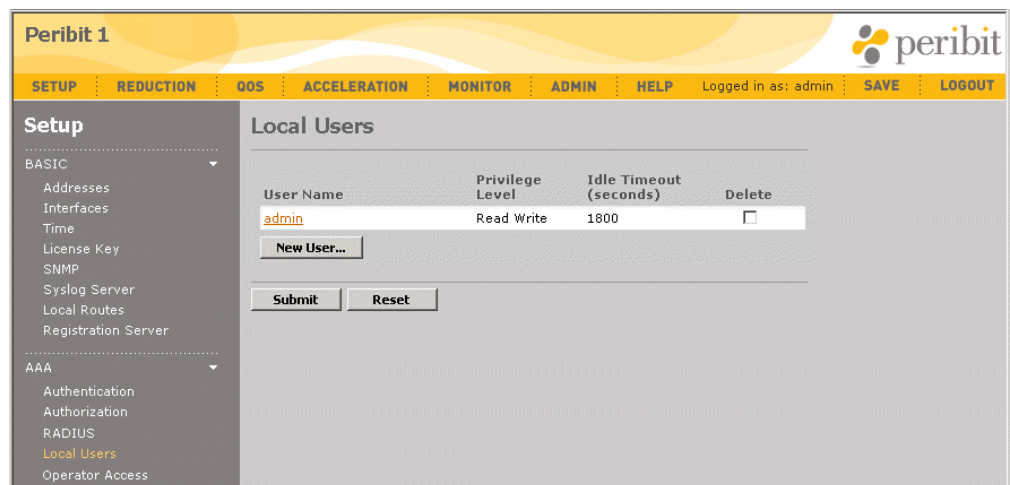
## Defining Local Users

You can define up to 25 users that can be authenticated locally. Each user can have full (admin) or read-only access privileges. The predefined **admin** account has a default password of **peribit**. To ensure secure access to the device, you should change the password periodically. To specify how users are authenticated (locally and/or through RADIUS), refer to “Selecting Authentication Methods” on page 80.

To define local user accounts:

1. In the Setup page, click AAA in the left-hand navigation frame, and then click Local Users.

**Figure 41: Defining Local Users**



From the Local Users page, you can:

- Add a new user account, as described in Step 2 through Step 4.
- Change a user account. Click the user name, make any needed changes, and click Submit.
- Delete user accounts. Select the check box next to the accounts you want to delete, and click Submit.

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

User Name	Enter the account name (up to 32 characters).
Privilege Level	Select administrator or read-only privileges.
Idle Timeout	Enter the number of minutes before an idle user is logged out (the default is 30) or select <b>Never</b> .
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 81).

---

3. Click Submit to activate the changes.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Securing Operator Access

You can create an Include or Exclude list to allow or deny access to the device from specific IP addresses or subnets. For example, if you enter one address in the Include list, 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. By default, the Include and Exclude lists are empty.

To restrict operator access:

1. In the Setup page, click AAA in the left-hand navigation frame, and then click Operator Access.

**Figure 42: Controlling Operator Access**

2. To allow access 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 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 activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## 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.

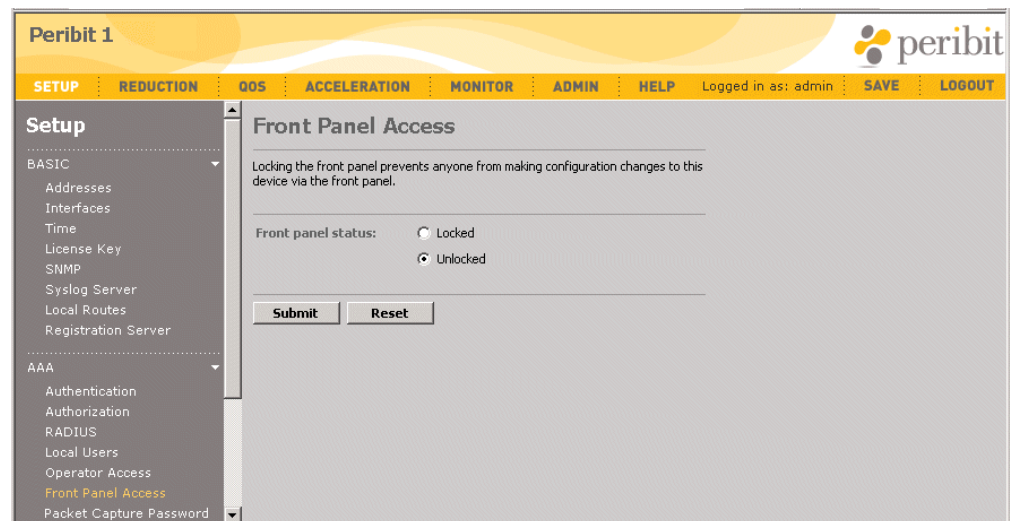


**NOTE:** The latest model of the WX 100 does not have a front-panel keypad. Also, locking the front panel on the earlier WX 100s does not lock the front panels of the client devices.

To lock the front panel keypad:

1. In the Setup page, click AAA in the left-hand navigation frame, and then click Front Panel Access.

**Figure 43: Controlling Front Panel Access**



2. To lock the front-panel keypad, select Locked.
3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Changing the Packet Capture Password

The default packet capture password for the Web console and CLI is **peribit**. To ensure secure access to the packet capture feature, you should change the password periodically.

To change the packet capture password:

1. In the Setup page, click AAA in the left-hand navigation frame, and then click Packet Capture Password.

**Figure 44: Changing the Packet Capture Password**

The screenshot shows the Peribit 1 web console interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left-hand navigation menu is expanded to the 'Setup' section, and the 'AAA' sub-menu is selected, showing options like Authentication, Authorization, RADIUS, Local Users, Operator Access, Front Panel Access, and Packet Capture Password. The main content area is titled 'Packet Capture Password' and contains a text box explaining that the password is required to start a packet capture or copy it to a local disk. Below this, there are three input fields: 'Old password:', 'New password:', and 'Verify new password:'. A 'Submit' button is located at the bottom of the form.

2. In the Packet Capture 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 changes.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Managing Applications

The following topics describe how to manage applications.

- “About Application Definitions” in the next section
- “Viewing the Application Overview” on page 91
- “Configuring Application Definitions” on page 92
- “Assigning Applications to Traffic Classes” on page 95
- “Monitoring Applications” on page 97



## About Application Definitions

Application definitions allow WX devices to identify the traffic of up to 256 applications (the WX 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”. For each additional application you define, you can:

- Enable or disable data reduction, as described in “Reducing Applications” on page 135.
- Accelerate the application’s traffic (if data reduction is enabled), as described in “Accelerating WAN Traffic” on page 189.
- Assign the application to a traffic class, as described in “Assigning Applications to Traffic Classes” on page 95. Traffic classes are used by outbound QoS to allocate WAN bandwidth, and by Multi-Path to direct traffic to a specific path between WX devices.
- Enable or disable data application monitoring for reports, as described in “Monitoring Applications” on page 97.
- View reduction and acceleration statistics for monitored applications, as described in “Monitoring and Reporting” on page 227.

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.

The following table lists the default application definitions. Each definition has rules to match any traffic that has the specified port number(s) as the source or destination.

**Table 1: Application Definitions**

Application	Order	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.

Application	Order	Port Numbers
Filenet	40	32768-32774
FTP	1	20-21 <b>Note:</b> 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
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

For each defined application, the Application Overview page shows the application's traffic class, and whether reduction, acceleration, and monitoring are enabled for the application.

To view the application overview:

1. In the Setup page, click APPLICATIONS in the left-hand navigation frame, and then click Overview.

**Figure 45: Application Overview Page**

Application Name	Traffic Class	Reduction	NSM	Active Flow Pipelining	Fast Connection Setup	CIFS Acceleration	Exchange Acceleration	Monitoring
AOL	Default	✓	✓	○	○	○	○	○
CIFS	Default	✓	✓	○	○	○	○	✓
Clearcase	Default	✓	✓	○	○	○	○	○
CVS	Default	✓	✓	○	○	○	○	○
DNS	Default	✓	✓	○	○	○	○	✓
Exchange	Default	✓	✓	○	○	○	○	○
Filenet	Default	✓	✓	○	○	○	○	○
from_port_7777	Default	✓	✓	○	○	○	○	✓
from_port_8888	Default	✓	✓	○	○	○	○	✓
from_port_9999	Default	✓	✓	○	○	○	○	✓
FTP	Default	✓	✓	○	○	○	○	✓

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 95.
Reduction	Indicates whether the application's traffic is reduced (refer to “Reducing Applications” on page 135.).
NSC	Indicates whether Network Sequence Caching is used for data reduction (refer to “Reducing Applications” on page 135). NSC requires a hard disk, and applies only to WXC 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 198).
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 197).
CIFS Acceleration	Indicates whether CIFS traffic for the application is accelerated (refer to “Enabling Microsoft CIFS Acceleration” on page 204).

Exchange Acceleration Indicates whether Exchange traffic for the application is accelerated (refer to “Enabling Microsoft Exchange Acceleration” on page 207).

Monitoring Indicates whether you can view statistics for the application (refer to “Monitoring Applications” on page 97).

## Configuring Application Definitions



**NOTE:** To add an application definition by selecting an undefined application from the Traffic report, refer to “Traffic Statistics” on page 256.

To add or change application definitions:

1. In the Setup page, click APPLICATIONS in the left-hand navigation frame, and then click Definitions.

**Figure 46: Application Definitions Page**

Order	Application name	Type	Definition	Delete
36	<a href="#">AOL</a>	Default	• Src Port: 5190-5193 • Dst Port: 5190-5193	<input type="checkbox"/>
6	<a href="#">CIFS</a>	CIFS	• Src Port: 445,139 • Dst Port: 445,139	<input type="checkbox"/>
23	<a href="#">Clearcase</a>	Default	• Src Port: 371 • Dst Port: 371	<input type="checkbox"/>
33	<a href="#">CVS</a>	Default	• Src Port: 2401 • Dst Port: 2401	<input type="checkbox"/>
15	<a href="#">DNS</a>	Default	• Src Port: 53 • Dst Port: 53	<input type="checkbox"/>
20	<a href="#">Exchange</a>	Exchange	• Src Port: 135 • Dst Port: 135	<input type="checkbox"/>
40	<a href="#">Filenet</a>	Default	• Src Port: 32768-32774 • Dst Port: 32768-32774	<input type="checkbox"/>
43	<a href="#">from_port_7777</a>	Default	• Src Port: 7777	<input type="checkbox"/>
45	<a href="#">from_port_8888</a>	Default	• Src Port: 8888	<input type="checkbox"/>
47	<a href="#">from_port_9999</a>	Default	• Src Port: 9999	<input type="checkbox"/>
1	<a href="#">FTP</a>	FTP	• Src Port: 20-21 • Dst Port: 20-21	<input type="checkbox"/>
29	<a href="#">Groupwise</a>	Default	• Src Port: 1677 • Dst Port: 1677	<input type="checkbox"/>

From the Application Definitions page, you can:

- Add a new application definition, as described in Step 2 through Step 6.
- 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 order. 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 47: Defining New Applications

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC  
Addresses  
Interfaces  
Time  
License Key  
SNMP  
Syslog Server  
Local Routes  
Registration Server

AAA

APPLICATIONS  
Overview  
Definitions  
Traffic Classes  
Monitoring

IPSEC

ADVANCED

**Application Definitions > New**

Application Name:

Application Type:

Application traffic will be identified using the following rules

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="Any"/>	<input type="text" value="Advanced"/> <input type="button" value="CLEAR"/>

Enter IP Address or subnet.  
Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0

Use commas to enter multiple ports.  
Use hyphen (-) to specify a range.  
Example: 25, 27, 125-135

To match any value, leave the field blank. Do not use asterisk (\*).

### 3. Specify the following information:

Application name Enter a name for the application (up to 63 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 204). 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 207). 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 209). 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 have a total of 512 rules for all applications.

Source Address Enter a source IP address or subnet. The general format is:

address/subnetmask

A blank or an asterisk (\*) with no subnet mask indicates any source IP address.

Source Port	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 ports, refer to Appendix , “Common Application Port Numbers”.
Destination Address	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 source and destination addresses will match only the traffic between those addresses.
Destination Port	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.
Protocol	<p>Select an application protocol or select <b>Any</b> 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><b>NOTE:</b> Any protocol defined by number is added to the <b>Any</b> list of defaults that applies to each rule that does not specify a protocol. To use application pattern matching (described below), select <b>TCP</b>.</p>

- 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> <p><b>IP Precedence.</b> Select an IP precedence value (0 through 7).</p> <p><b>DSCP.</b> Enter a DSCP value (0 through 63).</p> <p>For more information about ToS and DSCP, refer to “Changing Outbound ToS/DSCP Values” on page 181.</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> <p>&lt; host &gt; / &lt; uri &gt;</p> <p>Where:</p> <p><b>&lt;host&gt;</b> is up to eight strings separated by periods. You can use an asterisk (*) by itself to indicate any string. For example:</p> <p>www.juniper.*.net/</p> <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><b>&lt;uri&gt;</b> is up to eight strings separated by slashes. You can use an asterisk (*) by itself to indicate any string. For example:</p> <p>www.juniper.*.net/*index.htm</p> <p>Note that an asterisk is treated as a single character (not a wildcard) when it is part of a string, such as “www.juniper*.net”.</p>

Click Continue to return to the Application Definition page.

- Click Submit to activate the definition, or click Cancel to discard it. To erase an entire rule, including the advanced settings, click CLEAR.

6. To retain your changes when the device is restarted, click SAVE in the menu frame.

### **Testing New Application Definitions**

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

If you do not see any traffic for the application (refer to “Monitoring and Reporting” on page 227), 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).

If the new application is encrypted or already compressed, you should disable data reduction, as described in “Reducing Applications” on page 135. If you are accelerating traffic, verify that the new application is enabled or disabled (as appropriate) for each acceleration method, as described in “Accelerating WAN Traffic” on page 189.

### **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 WX 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 WX devices.

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

- “Understanding Outbound Bandwidth Management” on page 154
- “Configuring Policy-Based Multi-Path” on page 115

To define traffic classes and assign applications to each class:

1. In the Setup page, click APPLICATIONS in the left-hand navigation frame, and then click Traffic Classes.

**Figure 48: Assigning Applications to Traffic Classes**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

- BASIC
- AAA
- APPLICATIONS
  - Overview
  - Definitions
  - Traffic Classes**
  - Monitoring
- IPSEC
- ADVANCED

**Traffic Classes** Edit Classes... HELP

For each application, select the Traffic Class that best reflects the relative importance of the application.

Application Name	Traffic Class
AOL	Default
CIFS	Default
Clearcase	Default
CVS	Default
DNS	Default
Exchange	Default
Filenet	Default
from_port_7777	Default
from_port_8888	Default
from_port_9999	Default
FTP	Default

- To change the applications assigned to each traffic class, select the appropriate traffic class for each application, and click Submit.
- To add or change the current traffic classes, click Edit Classes.

From the Traffic Classes > Edit Classes page, you can:

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



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

- To retain your changes when the device is restarted, click SAVE in the menu frame.



## Monitoring Applications

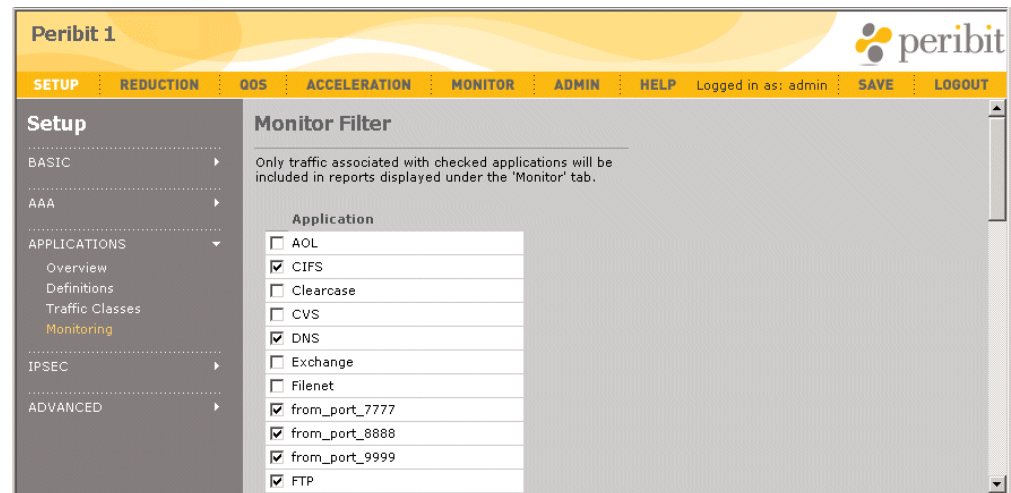
Monitoring an application lets you view reduction and acceleration statistics for the application. You can monitor up to 40 applications. All unmonitored applications are placed in the “Others” category on reports. For more information about monitoring statistics, refer to “Monitoring and Reporting” on page 227.

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 92.

To select applications to be monitored:

1. In the Setup page, click APPLICATIONS in the left-hand navigation frame, and then click Monitoring.

**Figure 49: 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 unmonitored applications are placed in the “Others” category on reports.



**NOTE:** If you disable monitoring for an application, its historical monitoring statistics are permanently moved to the “Others” application category on the reduction and acceleration reports.

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.



## Chapter 4

# Configuring Advanced Setup Policies

The following topics describe the advanced setup procedures:

- Setting Community Feature and Topology Parameters on page 99.
- Filtering Data Reduction by Source and Destination on page 101
- Configuring the ARP Table on page 104
- Defining the Prime Time on page 105
- Configuring Packet Interception on page 106
- Configuring Policy-Based Multi-Path on page 115
- Configuring WAN Performance Monitoring on page 124

### Setting Community Feature and Topology Parameters

---

The Features/Topology page specifies the device's topology setting, and whether Active Flow Pipelining (AFP) and Application Flow Acceleration are used to accelerate traffic. The topology setting and acceleration features affect the maximum number of tunnels each device can support.

WXOS 5.3 is only for the WX 100 in stack mode (one or more clients), so it supports only two topology settings: Hub range 5 and Mesh range 0, both with all features enabled. Use Mesh range 0 in a point to point topology where two WX 100 servers communicate with each other.

Table 2 shows the supported ranges of devices for the WX 100 and its client devices (the WX 60 and WXC 500),



**NOTE:** Ranges 0 through 4 can be selected, but they are not supported by WXOS 5.3. Use only range Mesh range 0 or Hub range 5.

---

**Table 2: Device Ranges by Model and Topology**

Device	Mesh Ranges – All Features	Hub Ranges – All Features
WXC 500	0 = Up to 10	5 = Up to 50
WX 60	0 = Up to 40	5 = Up to 110
WX 100 (with WX 60 or WXC 500 clients)	When a WX 100 has one or more client devices (stack configuration), the client model type is detected, and the topology ranges displayed in the Web console are for the maximum number of devices when all six clients are connected to the server. For example, if a WX 100 hub using all features has one WXC 500 client, the highest topology range is 300 devices, which is the number of devices supported by six WXC 500 clients (6 x 50).	

There are some configuration combinations that can result in a decreased number of tunnels. Use the following guidelines to determine if a tunnel penalty exists:

- A WX 100 stack running WXOS 5.3 configured for Hub range 5 can speak to any other device set to "spoke" with no tunnel penalty
- A WX 100 stack running WXOS 5.3 configured for Hub range 5 can speak to any other WX 100 stack running WXOS 5.3 configured for Hub range 5 with no tunnel penalty.
- A WX 100 stack running WXOS 5.3 configured for Mesh range 0 can speak to any other WX 100 stack running WXOS 5.3 configured for Mesh range 0 with no tunnel penalty.
- Any other combination results in a 2 for one tunnel penalty.

### Examples

- A WXC stack with WXC 500 clients in Hub range 5 running WXOS 5.3 can speak to any combination of remote WXC 250s or WXC 500s configured as spokes, up to a maximum of 300 devices (a "no penalty" case)
- A WXC stack with WXC 500 clients in Hub range 5 running WXOS 5.3 can speak to 250 remote WXC 250s or WXC 500s configured as spokes, and an additional 25 WX-50s or WX-20s configured as Mesh, for a total of 275 tunnels (a 2 for 1 tunnel penalty is paid for the WCs set to a non-spoke topology).
- A WXC stack running WXOS 5.3 with 6 WXC 500 clients configured for Hub range 5 with multi-tunnel enabled can speak to 5 other identically configured WXC stacks (for a total tunnel consumption of 30 tunnels) and any combination of 270 remote WXC 250s or WXC 500s configured as spokes, for a total of 300 tunnels (a "no penalty" case)
- A WXC stack running WXOS 5.3 with 6 WXC 500 clients configured for Mesh range 0 with multi-tunnel enabled can speak to 3 other identically configured WXC stacks (for a total tunnel consumption of 18 tunnels) and 21 other WXC remote devices configured for Mesh range 0 (no penalty for the stack to stack tunnels, but a 2 for 1 penalty for a stack in Mesh range 0 speaking to remote non-stack devices).

To review or change the topology settings:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, and then click Features/Topology.

**Figure 50: Reviewing and Changing the Feature/Topology Settings**

2. For a WX 100 running WXOS 5.3, specify the following:
  - a. **Features.** Select “All features” (the default).
  - b. **Topology.** Select Hub with range 5 (the default) or Mesh with range 0. Use Mesh range 0 in a point to point topology where two WX 100 servers communicate with each other. Note that if you change this setting, all reduction tunnels will be reset when you click Submit.
3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Filtering Data Reduction by Source and Destination

You can enable or disable data reduction between specific sources and destinations by creating a list of source and destination addresses or subnet pairs that are either included or excluded from data reduction. A 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 135. The source/destination filter is applied before the application filter, and is more efficient.

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 by all WX and WXC devices, disable the advertisement of the subnet (refer to “Advertising Reduction Subnets” on page 131).

Note the following:

- If you disable data reduction between a source and destination, traffic between those points cannot be accelerated. Also, for an oversubscribed WAN 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 Setup page, click ADVANCED in the left-hand navigation frame, and then click Source/Destination Filter.

**Figure 51: 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 other WX and WXC 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, for an oversubscribed WAN, is managed by the outbound QoS policies defined for the Default traffic class under the “Other traffic” endpoint).

3. Specify the following information:

Source	<p>Enter a source IP address or subnet. The general format is: address/subnetmask</p> <p>The default subnet mask is “255.255.255.255”. An asterisk (*) with no subnet mask indicates any source IP address.</p>
Destination	<p>Enter a destination IP address or subnet (same format as the source address). An asterisk (*) indicates any destination IP address.</p>
Bidirectional	<p>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.</p> <p>In Profile Mode, you should exclude all traffic sent to the subnet where the WX device is installed. For more information about Profile Mode, refer to “Profile Mode” on page 417.</p>

4. Click Submit to activate the changes. To restore the original parameters, click Reset.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring the ARP Table

The Address Resolution Protocol (ARP) is used to:

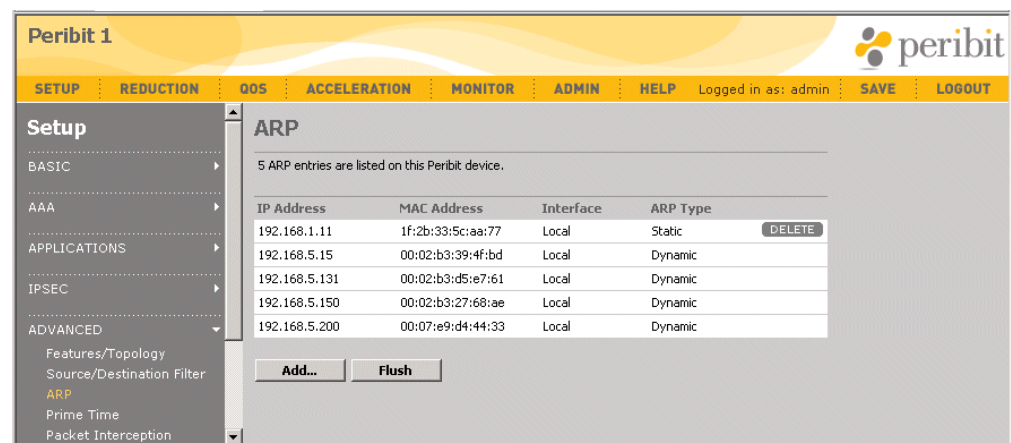
- Determine whether the gateway for a route is on the Local or Remote interface
- Discover the hardware (MAC) addresses of devices that are directly addressable on the Local and Remote interfaces

For devices that do not respond to ARP requests, you can add static ARP entries that map their IP addresses to their MAC addresses. You can also clear the dynamic ARP entries if you suspect some entries are out of date.

To configure the ARP table:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, and then click ARP.

**Figure 52: Viewing the ARP Table**



2. To delete all dynamic ARP entries, click Flush. This forces new ARP requests to be issued as needed.
3. To delete a static ARP entry, click DELETE next to the entry.
4. 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. The format of the MAC address is:  
xx:xx:xx:xx:xx:xx.

Click Submit to activate the new entries, or click Cancel to discard them.

5. To retain your changes when the device is restarted, click SAVE in the menu frame.



## 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. 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 Setup page, click ADVANCED in the left-hand navigation frame, and then click Prime Time.

**Figure 53: Defining the Prime Time**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

- BASIC
- AAA
- APPLICATIONS
- IPSEC
- ADVANCED
  - Features/Topology
  - Source/Destination Filter
  - ARP
  - Prime Time**
  - Packet Interception
  - WAN Performance Monitor
  - Multi-Path

**Prime Time**

This page allows you to modify the definition of prime time periods. This definition can be used to filter statistical reports based on traffic that occurs during prime time periods only. In addition, bandwidth management policies can be optimized for prime time vs. non-prime time periods.

☐ Use Prime Time

From To

Hours: 12 AM 12 AM

Days: ☒ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat


Submit Reset

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 activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring Packet Interception

WX and WXC 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 54). 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).

**Figure 54: Off-Path Deployment**



The following topics describe how to configure packet interception on a WX device and on the local switch or router. 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 109
- “WCCP Router Configuration Commands” on page 112
- “External Policy-Based Router Commands” on page 113
- “Alternatives to Packet Interception” on page 113

### Configuring Packet Interception for Off-Path Devices

In an off-path deployment, the traffic to be reduced must be routed to the WX device using packet interception. Both the router and the WX 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 WX devices in the community. Note the following:
  - If a remote WX 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 WX 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 109.


- **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 112.
- **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 113.

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, inbound QoS is not supported, and outbound QoS 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 WX 100 with up to six client devices can be installed off path.

To configure packet interception for an off-path device:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, and then click Packet Interception.

**Figure 55: 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 <b>Password</b> and enter the RIP password. This is the same password used to discover dynamic routes (refer to “Enabling RIP and OSPF Support” on page 69).
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 “configure packet-interception” on page 331.

- 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: <a href="http://www.iana.org/assignments/protocol-numbers">http://www.iana.org/assignments/protocol-numbers</a>
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. <b>In high-availability environments, where two WX devices use the same router, they must use different IDs for the same protocol.</b>  Heartbeat packets are sent to the router every 10 seconds for each service. If the WX 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 113.

3. Click Submit to activate 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 131). 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.

5. To retain your changes when the device is restarted, click SAVE in the menu frame.

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 WX 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 56). Installing the WX device on a dedicated VLAN is recommended to reduce the routing failover time if the WX device fails. The port where the WX 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 56: Off-Path WX 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 WX device. To match this example, enter the following CLI command on the WX 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 57, 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 57: 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 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 WX device. To match this example, enter the following CLI command on the WX 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 58. 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 58: 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 WX device. However, entering the “show packet-interception” command on the WX 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 WX 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 WX device. Do not use the “set interface” command to redirect traffic to the interface where the WX device is connected.

## Alternatives to Packet Interception

In some environments, you can install a WX device off path 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 59, the two WX devices are connected in “two-legged” VLANs on two Layer 2 switches. All traffic is switched through the WX devices as it passes to and from the WAN routers.


**Figure 59: 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 each WX device is the HSRP address of the WAN routers. If one router fails, traffic is directed to the other router.
- Use a crossover cable to connect the Local interface to the switch so that traffic is blocked if one WX device fails. The Layer 3 switches can then route the traffic through the other WX device.

### Layer 3 Switch Sandwich

Figure 60 shows a single device connected across Layer 2 and Layer 3 VLANs on an L2/L3 switch. All traffic is switched through the WX device as it passes to and from the WAN routers.

**Figure 60: 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 the WX device is the HSRP address of the WAN routers. If one router fails, traffic is directed to the other router.

## 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.




**NOTE:** Each Multi-Path endpoint counts as two reduction tunnels.

To use Multi-Path, you configure both WX 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.

For example, in Figure 61, most traffic is normally 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 61: Multi-Path Deployment**



The following topics describe how to configure policy-based, multi-path tunnels:

- “Procedure for Configuring Multi-Path” in the next section
- “Enabling Multi-Path and Defining Marking Methods” on page 117
- “Defining Multi-Path Templates” on page 119

- “Defining Multi-Path Endpoints” on page 121
- “Configuring Routers to Support Multi-Path” on page 123

### ***Procedure for Configuring Multi-Path***

To configure Multi-Path for a pair of WX devices, do the following on BOTH devices:

1. Verify that the WX device that acts as the registration server is running WXOS 5.0 or later.
2. Verify that data reduction is enabled in both directions between the two devices (refer to “Configuring Endpoints for Reduction Tunnels” on page 129).
3. Verify that the appropriate traffic classes are defined (refer to “Assigning Applications to Traffic Classes” on page 95).
4. Enable the multi-path feature and specify a secondary IP address (refer to “Enabling Multi-Path and Defining Marking Methods” on page 117).
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 (refer to “Defining Multi-Path Templates” on page 119).
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 121).
7. If necessary, configure the WAN router to route traffic to the appropriate path (refer to “Configuring Routers to Support Multi-Path” on page 123).
8. Optionally, enable encryption for both paths or just the less-secure path (refer to “Configuring IP Security (IPSec)” on page 213).

## Enabling Multi-Path and Defining Marking Methods

To enable Multi-Path, you must specify a secondary IP address to be used as the source address on all packets to be routed to the secondary path. Optionally, packets sent on the primary and secondary paths can be marked with different gateway addresses or ToS/DSCP values.

To enable Multi-Path:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, click Multi-Path, and then click Start/Stop.

**Figure 62: Multi-Path Start/Stop Page**

2. Specify the following information:


Multi-Path	Select Enabled to activate the multiple-path feature on this device.
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 <b>Gateway IP</b> 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 123).</p> <p><b>NOTE:</b> If you enter an address assigned to another device, the path will remain inactive (refer to “Defining Multi-Path Endpoints” on page 121). If you must change the address, do the following:</p> <p>Enter the new address and click Submit.</p> <p>Disable Multi-Path and click Submit, and then enable Multi-Path. The configuration for the old secondary path is disabled, including settings for QoS, PFA, and IPSec.</p>

Optionally, you can mark packets sent on the primary and secondary paths with different ToS/DSCP values or gateway addresses. You can specify values for both marking methods, but only one method can be used for each remote endpoint that supports Multi-Path.

Gateway IP	<p>If the WAN routers for the primary and secondary paths are on the same subnet as the WX device, and the WX device is connected to a Layer 2 switch (see Figure 63), 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>
IP Precedence/DSCP	<p>Select <b>IP Precedence</b> or <b>DSCP</b> 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.</p> <p>These values override the IP precedence or DSCP settings for:</p> <ul style="list-style-type: none"> <li>■ Outbound QoS (refer to “Changing Outbound ToS/DSCP Values” on page 181)</li> <li>■ WX control packets (refer to “configure reduction” on page 343)</li> </ul> <p>The multi-path DSCP values also override ToS marking for router-based balancing (refer to “configure route” on page 356).</p>

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

**Figure 63: Multi-Path with Primary and Secondary Gateways**



## Defining Multi-Path Templates

To configure Multi-Path, at least one multi-path template must be defined 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 WX device that supports Multi-Path, refer to “Defining Multi-Path Endpoints” on page 121.

To define multi-path templates:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, click Multi-Path, and then click Templates.

**Figure 64: Defining Multi-Path Templates**

Peribit 1

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

ADVANCED

- Features/Topology
- Source/Destination Filter
- ARP
- Prime Time
- Packet Interception
- WAN Performance Monitor
- Multi-Path
  - Start/Stop
  - Templates**
  - Endpoints

**Multi-Path Templates** HELP

MP Template Name Delete

Primary cong/fail ☐

Submit Reset New Template...

Multi-Path templates are used to define the multi-path behavior for the various traffic classes.

Later, on the 'Endpoints' page, you can select the template to use for each Multi-Path enabled endpoint in your network.

From the Multi-Path 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 a template is applied to an endpoint, it cannot be deleted.

2. To add a new template, click New Template.

**Figure 65: Defining a New Multi-Path Template**

Peribit 1

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

- BASIC
- AAA
- APPLICATIONS
- IPSEC
- ADVANCED**

**Multi-Path Templates** HELP

Template Name

Traffic Class	Preferred Path	Divert
Default	Primary	Never
Business Critical	Primary	Never
Business Standard	Primary	Never
Low-Latency	Primary	Never

Submit Reset Cancel

For each class, traffic is initially directed to the path indicated in the table. If the Bypass Condition is met, traffic is diverted to the alternate path.

As a safeguard, traffic is not diverted unless the performance of the alternate path is better than the initial path.

If a link failure occurs, all traffic classes (except those with Bypass Condition set to 'Never') are immediately diverted to the Secondary path.

Specify the following information:

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 95).	
Preferred Path	Select <b>Primary</b> or <b>Secondary</b> to indicate the path used for each traffic class under normal network conditions.
Divert	<p>Select the conditions under which each traffic class can be switched to the alternate path:</p> <p>Never. The traffic class is never diverted from the preferred path.</p> <p>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.</p> <p><b>Congestion/Failure.</b> 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> <p>If <b>Congestion/Failure</b> 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 121. 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 “configure multi-path” on page 326).</p>



**NOTE:** Outbound QoS settings do not affect how traffic is diverted between alternate paths.

---

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.



## Defining Multi-Path Endpoints

After you specify a Multi-Path secondary IP address for one or more remote WX devices, you can assign a Multi-Path template to each remote endpoint, and specify the latency threshold and supplemental marking method (if any) for each path.

To define Multi-Path endpoints:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, click Multi-Path, and then click Endpoints.

**Figure 66: Defining Multi-Path Endpoints**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

AAA

APPLICATIONS

IPSEC

ADVANCED

Features/Topology

Source/Destination Filter

ARP

Prime Time

Packet Interception

WAN Performance Monitor

Multi-Path

Start/Stop

Templates

Endpoints

**Multi-Path Endpoints**

Find:  GO HELP

When Multi-Path is enabled, in the event of network congestion or link failure, outbound traffic destined for the checked Peribit devices is diverted to alternate paths according to conditions defined by the selected Multi-Path Template and Supplemental Marking Method.

Use the 'Start/Stop' page to enable or disable Multi-Path and to specify supplemental marking methods. Use the 'Templates' page to view and modify templates.

Only endpoints which have been selected for Reduction and which have a secondary IP address can be enabled for Multi-Path. If an endpoint has been disabled and you want to enable Multi-Path, go to Reduction/Endpoints page and enable Reduction for it and go to that device and set the secondary IP address on it.

Device Name	Status		Latency Threshold (msec)		Multi-Path Template	Supplemental Marking Method
	Pri	Sec	Primary	Secondary		
<input checked="" type="checkbox"/> SR-10.15.2.12	●	✗	50	75	Primary_cong/fail	None (Sec. IP only)
<input checked="" type="checkbox"/> SR-10.2.2.32	●	●	50	100	Primary_cong/fail	None (Sec. IP only)

Select All Clear

Submit Reset

Remote devices that are greyed out do not have a secondary IP address defined (refer to “Enabling Multi-Path and Defining Marking Methods” on page 117).



2. To enable Multi-Path between this device and a remote endpoint, select the check box next to the remote endpoint.

To view the list of Multi-Path endpoints starting with a specific device name, enter the first part of the name (or the entire name) in the Find box at the top of the page, and click GO. To select all the devices displayed on the page, click Select All. To deselect all displayed devices, click Clear. If you disable an endpoint, all subsequent traffic to that endpoint is sent on the primary path.

3. Specify the following for each selected endpoint:

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><b>NOTE:</b> 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. Traffic is switched 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 “configure multi-path” on page 326.</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>
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 119.</p>
Supplemental Marking Method	<p>Optionally, select one of the additional marking methods for the packets sent on each path (refer to “Enabling Multi-Path and Defining Marking Methods” on page 117). By default, all packets to be sent on the secondary path have the source address set to the secondary IP address.</p>

4. Click Refresh to update the icons in the **Status** column. The following icons are used to indicate the status of the primary and secondary paths of each multi-path endpoint:

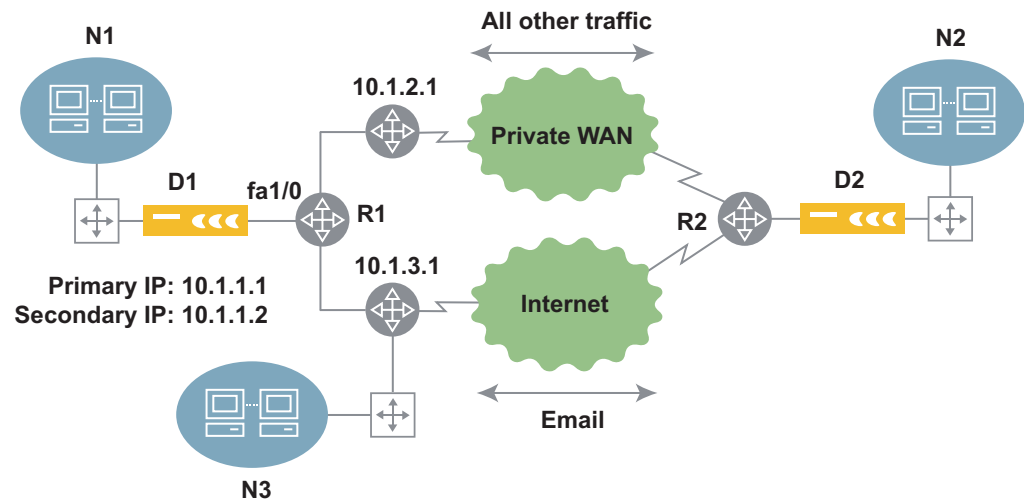
Icon	Description
	The reduction tunnel is up and the path's loss and latency are below the specified thresholds.
	<p>Connection or performance problem. Move the cursor over the icon to see which one of the following conditions applies.</p> <ul style="list-style-type: none"> <li>■ No secondary IP address for the remote endpoint (the address specified may belong to another device)</li> <li>■ Outbound reduction tunnel is down</li> <li>■ Loss threshold exceeded</li> <li>■ Latency threshold exceeded</li> </ul> <p>Note that when loss or latency thresholds are exceeded, traffic is switched to the alternate path only if the alternate reduction tunnel is up and the loss and latency are below the specified thresholds. If a reduction tunnel is down, traffic is switched to the alternate path regardless of the alternate's performance (provided the alternate reduction tunnel is up).</p>

5. Click Submit to activate the changes, or click Reset to discard them.
6. To retain your changes when the device is restarted, click SAVE in the menu frame.

## 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 67. A similar configuration is needed for R2.

**Figure 67: 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 D2.

To configure the WAN router R1 to use both the source address and the ToS IP precedence or DSCP values:

4. Define a route map for Multi-Path (see the previous example).
5. 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
```

6. 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 WX 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 67, you may not be able to access D1 from N3 because management responses have the primary IP address, and are routed to the private WAN.

## Configuring WAN Performance Monitoring

WAN performance monitoring lets you measure the latency and loss between the current device and 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 (refer to “WAN Performance Statistics” on page 231). 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 Policy-Based Multi-Path” on page 115).

To enable WAN performance monitoring:

1. In the Setup page, click ADVANCED in the left-hand navigation frame, and then click WAN Performance Monitor.

Figure 68: Configuring WAN Performance Monitoring

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

AAA

APPLICATIONS

IPSEC

ADVANCED

Features/Topology

Source/Destination Filter

ARP

Prime Time

Packet Interception

WAN Performance Monitor

Multi-Path

**WAN Performance Monitoring** Find:  GO

☐ Enable WAN Performance Monitoring for checked endpoints

Device Name	Endpoint	Latency Threshold (msec)
<input type="checkbox"/> 55/22-SR100	192.168.55.22	5000
<input type="checkbox"/> 55/32-SR55-BACKUP-SR	192.168.55.32	5000
<input type="checkbox"/> 55/42-SR55	192.168.55.42	5000
<input type="checkbox"/> 57/22-SR50	192.168.57.22	5000
<input type="checkbox"/> 58/22-SR50	192.168.58.22	5000
<input type="checkbox"/> 59/22-SR20	192.168.59.22	5000

Select All Clear

Submit Reset

2. Select the **Enable WAN Performance Monitoring for checked endpoints** check box.
3. To enable WAN performance monitoring between this device and a remote endpoint, select the check box next to the remote endpoint.

To view the list of endpoints starting with a specific device name, enter the first part of the name (or the entire name) in the Find box at the top of the page, and click GO. To select all devices displayed on the page, click Select All. To deselect all displayed devices, click Clear.

4. Specify the following for each selected endpoint:

**Latency Threshold** 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.

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 “configure wan-performance-monitor” on page 366.

Note that availability on the WAN Performance report is measured as the percentage of minutes for which at least one probe was acknowledged.

5. Click Submit to activate the changes, or click Reset to discard them.
6. To retain your changes when the device is restarted, click SAVE in the menu frame.









## Chapter 5

# Configuring Reduction Policies

This chapter describes how to configure basic and advanced reduction policy settings through the Web Console.

- “Configuring Basic Reduction Policies” in the next section
- “Configuring Advanced Reduction Policies” on page 137

### Configuring Basic Reduction Policies

---

The following topics describe how to configure basic reduction policies:

- “Configuring Endpoints for Reduction Tunnels” in the next section
- “Advertising Reduction Subnets” on page 131
- “Configuring Network Sequence Caching” on page 134
- “Reducing Applications” on page 135

### Configuring Endpoints for Reduction Tunnels

When you install a new WX device and specify a registration server, the device attempts to form a reduction tunnel with each registered device, or “endpoint,” in the same community. The existing devices also attempt to form tunnels with the new device, so that each device can have two types of tunnels—OUT tunnels that convey reduced data to remote devices, and IN 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 131). At any time, you can disable data assembly and/or reduce data only for specific WX devices in the community.

To configure the endpoints for reduction tunnels:

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

Figure 69: Configuring Endpoints for Reduction Tunnels

Peribit 1 Number of Active Clients -- 1

SETUP REDUCTION OOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

Endpoints

Reduction Subnets

Application Filter

ADVANCED

**Endpoints** Find:  GO

☒ Enable this device to ASSEMBLE traffic from all other Peribit devices

☒ Enable this device to REDUCE traffic destined for:

☒ ALL discovered Peribit devices

☐ ONLY Peribit devices designated as hubs

☐ ONLY checked Peribit devices below

Device name	IP address	Duties	Tunnel Status		Description
			OUT	IN	
<input checked="" type="checkbox"/> SM-192.168.70.11	192.168.70.11				
<input checked="" type="checkbox"/> SR-192.168.71.10	192.168.71.10				No request received
<input checked="" type="checkbox"/> SR-192.168.73.11	192.168.73.11				
<input checked="" type="checkbox"/> SR-192.168.74.11	192.168.74.11				
<input checked="" type="checkbox"/> SR-192.168.75.10	192.168.75.10				
<input checked="" type="checkbox"/> SR-192.168.76.10	192.168.76.10				
<input checked="" type="checkbox"/> SR-192.168.78.10	192.168.78.10				Maximum tunnel limit / No request received

Select All Clear Click on the IP address to login to another device

Submit Reset

**Legend**











- Hub
- Spoke
- Mesh
- Registration Server
- Secondary Reg. Server
- Backup
- Backup (Active)
- OUT Reduction tunnel from this device to remote device
- IN Reduction tunnel from remote device to this device
- Tunnel established by server
- Tunnel established by client on port 1
- No tunnel established
- Broken tunnel
- Temporarily Unavailable
- Unavailable

- To stop this device from assembling reduced data, 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 this device.
- To stop this device from 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 WX devices (default).
  - ONLY Peribit devices designated as hubs.** Data is reduced only for WX devices designated as a hub.
  - ONLY checked Peribit devices below.** Data is reduced only for the selected WX devices. Click the check box next to the appropriate devices. To view the list of endpoints starting with a specific device name, enter the first part of the name (or the entire name) in the Find box at the top of the page, and click GO. To select all devices displayed on the page, click Select All. To deselect all displayed devices, click Clear.

Note the following about the list of devices:

- To access another device, click the device name and enter the administrator user name and password for the device.

- The following icons are used in the **Duties** and **Tunnel Status** columns. The **IN** column indicates the status of the tunnel from the remote device; the **OUT** column indicates the status of the tunnel from this device to the remote device.

Icon	Description
	<b>Hub</b> — The device is designated as a hub in the community. Each device attempts to form a reduction tunnel with a hub before creating tunnels to other WX devices (refer to “Setting Community Feature and Topology Parameters” on page 99).
	<b>Spoke</b> — The device is designated as a spoke in a Hub and Spoke topology. By default, a spoke reduces and assembles data only for the hub device(s).
	<b>Mesh</b> — The device is designated as part of a mesh topology.
	<b>Registration Server</b> — The device is the primary registration server for the community.
	<b>Secondary Registration Server</b> — Indicates that this device is the secondary registration server for the community.
	<b>Backup and Backup (Active)</b> — The device is designated as backup for one or more primary devices. The icon flashes when the backup device is active. To configure a device as a backup, refer to “configure backup” on page 312.
	<b>Tunnel established</b> — A reduction tunnel exists between this device and the remote device at the specified IP address. On an WX 100 that has one or more clients, an “S” or a number (1 to 6) is enclosed in the circle to indicate whether the server (the WX 100) or a client is handling the tunnel. The number indicates the port on the WX 100 where the client is connected (also called the client ID). Note that remote devices see only the WX 100, not the clients.
	<b>No tunnel established</b> — No reduction tunnel exists between this device and the remote device due to a policy setting. For example, if you disable data reduction to a remote device by clearing the check box next to its IP address, this icon is displayed in the <b>OUT</b> column, and the message “Disallowed by policy” is displayed in the <b>Description</b> column.
	<b>Broken tunnel</b> — No reduction tunnel exists between this device and the remote device because of a policy setting or an error. If you manually disable data reduction to a remote device, the remote device displays this icon in the corresponding <b>IN</b> column and “No request received” in the <b>Description</b> column.
	<b>Temporarily unavailable</b> — The reduction tunnel is in a transitory state.

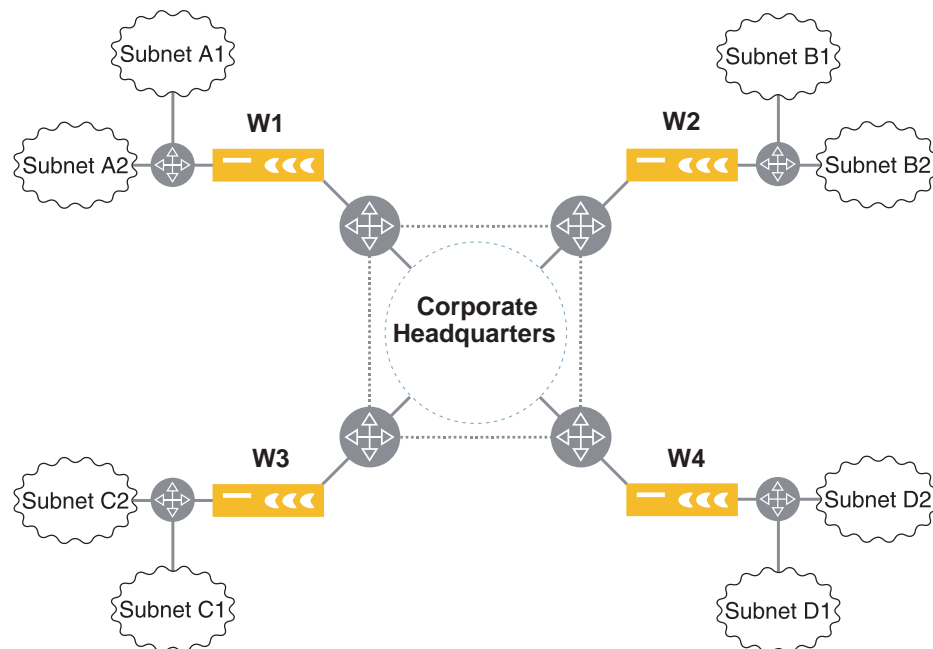
4. Click Submit to activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

### Advertising Reduction Subnets

Reduction subnets are the subnets on the LAN side of the WX 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 WX device is installed. To identify more LAN-side subnets, refer to “Configuring Local Routes” on page 66.

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. Figure 70 shows four WX devices, each with two subnets on its Local side.

**Figure 70: Selecting Specific Subnets for Data Reduction**



To disable data reduction for Subnet D1, log in to W4 and deselect Subnet D1 on the Reduction Subnet list. Data from other subnets that is destined for Subnet D1 passes through the community without reduction. Data that is destined for Subnet D2 is still reduced by the other WX devices and assembled by W4.

For further control of the traffic being reduced, you can specify application filters, as described in “Reducing Applications” on page 135, and source/destination filters, as described in “Filtering Data Reduction by Source and Destination” on page 101.

If a WX device has 4000 or more reduction subnets, it may take considerable time to load them into the Web console. In this case, you may want to use the CLI to view and configure the reduction subnets. For more information, refer to “configure reduction-subnet” on page 350.



**NOTE:** If a host or gateway in an advertised subnet becomes unreachable, the WX device dynamically adjusts the advertised subnets to exclude (“carve out”) the unreachable address. To view the most recent advertised subnets, refer to “Viewing and Fetching Remote Routes” on page 138. To disable the carve-out feature (refer to “configure reduction-subnet” on page 350).

To advertise reduction subnets:

1. Click REDUCTION in the menu frame, and then click Reduction Subnets in the left-hand navigation frame.

**Figure 71: Configuring Reduction Subnets**

The **Cost** column is not the standard routing cost, but an internal value used to calculate the cost of reduction tunnels. The **Interface** column indicates whether the route was discovered on the Local or Remote interface.



**NOTE:** Normally, reduction subnets include only subnets discovered on the Local interface (the LAN side). Subnets discovered on the WAN side are included if the device is installed off-path (refer to “Configuring Packet Interception” on page 106) or if the WAN reduction subnet option is enabled manually (refer to “configure reduction-subnet” on page 350). However, WAN-side subnets are excluded if their next hop is the default gateway.

The WAN option is useful when local routes are discovered on the Remote interface, such as in some VLAN environments. For an off-path device, the interface value is “N/A”, so be careful to advertise only the true LAN-side subnets.

2. Select one of the following parameters for the reduction subnet list:
  - **Advertise ALL subnets.** Advertises all subnets in the list to all the devices in the community. This option is not available when the WAN reduction subnet option is enabled.
  - **Advertise checked subnets ONLY.** Advertises only the selected subnets. Select the subnets in the list that you want to advertise.
  - **Advertise all subnets EXCEPT checked.** Advertises all subnets in the list, except those that are checked. Select the subnets that you do NOT want to advertise.



Note that changes to advertised subnets are propagated to the other devices immediately. However, if reduction is disabled, changes are propagated every hour unless the fetch interval is changed (refer to “Viewing and Fetching Remote Routes” on page 138).

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

### Configuring Network Sequence Caching

Network Sequence Caching (NSC) is an enhanced data reduction technique available on WXC devices. NSC 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). NSC is most effective where large files are often sent over the WAN, such as for database backups.

Disk icons displayed in the banner of a WXC device indicate the status of the hard disk(s):

Icon	Description
	The hard disk is operating normally.
	The hard disk has failed. On the WXC 250, NSC stops and only MSR is used for reduction; on the WXC 500, NSC continues operation unless the second disk also fails. Contact Technical Support about any disk failures.

To use NSC between two WXC devices, reduction tunnels must exist between them in both directions (refer to “Configuring Endpoints for Reduction Tunnels” on page 129), and Active Flow Pipelining (AFP) and outbound QoS must be enabled on both devices (refer to “Enabling Packet Flow Acceleration by Endpoint” on page 194). Applications that are enabled for AFP can then be enabled for NSC (refer to “Reducing Applications” on page 135).

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

To configure NSC for remote WXC devices:

1. Click REDUCTION in the menu frame, and then click Network Sequence Caching in the left-hand navigation frame.

**Figure 72: Configuring Reduction Subnets**

IP address	Device name	Circuit Speed (Kbps)
<input checked="" type="checkbox"/> 192.168.53.2	53/2-SR55	100000

2. To disable NSC on this device so that standard data reduction is used for all remote devices, clear the **Enable Network Sequence Caching...** check box. Otherwise, select one of the following options:
  - **All NSC-capable Peribit devices.** NSC is used for all remote WXC devices in the community (default).
  - **ONLY NSC-capable Peribit devices checked below.** NSC is used only for the selected WXC devices. Click the check box next to the IP address of the appropriate devices. To select all devices, click Select All. To deselect all devices, click Clear.



**NOTE:** NSC takes effect only if it is enabled in both directions.

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Reducing Applications

For each application, you can enable or disable data reduction and Network Sequence Caching (NSC). 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”. If the undefined applications are reduced, they are monitored automatically. To define additional applications, refer to “Managing Applications” on page 88.

To select applications to be reduced:

1. Click REDUCTION in the menu frame, and then click Application Filter in the left-hand navigation frame.

**Figure 73: Selecting Applications for Reduction**

**Peribit-SM**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Network Sequence Mirror
- Application Filter

ADVANCED

**Application Filter**

Application name	Reduce	Network Sequence Mirror
<a href="#">AOL</a>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
<a href="#">CIFS</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">Clearcase</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">CVS</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">DNS</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">Exchange</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">Filenet</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">FTP</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">Groupwise</a>	<input type="checkbox"/>	<input type="checkbox"/>
<a href="#">Hostname Resolution</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">HTTP</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">HTTPS</a>	<input type="checkbox"/>	<input type="checkbox"/>
<a href="#">ICA</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">Kerberos</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">LDAP</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<a href="#">XWindows</a>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Undefined applications	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Only applications checked in the 'Reduction' column will be reduced. Unchecked applications will be passed through without reduction.

If Network Sequence Mirroring (NSM) is enabled for traffic destined for NSM-capable Peribit devices (See the "Network Sequence Mirror" page.) then applications checked in the "Network Sequence Mirror" column will be reduced using NSM.

Applications for which Active Flow Pipelining (AFP) have not been enabled cannot be enabled for NSM.

Reduce All NSM All Clear

Submit Reset

2. To view or change an application's definition, click an application name, make any needed changes, and click Submit.



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, SMTP, 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 “Accelerating WAN Traffic” on page 189).</p>
NSC	<p>On a WXC device, you can enable Network Sequence Caching (NSC) for reduced applications (enabled by default for most applications). If NSC is enabled for one or more remote WXC devices (refer to “Configuring Network Sequence Caching” on page 134), then NSC is used to reduce the application traffic sent to those devices.</p> <p>NSC 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 WX devices or to WXC devices where NSC is disabled.</p> <p>To enable NSC for all reduced applications, click NSC All. To use NSC 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 197).</p>

4. Click Submit to activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring Advanced Reduction Policies

---

The following topics describe the advanced data reduction policies:

- “Viewing and Fetching Remote Routes” in the next section
- “Configuring Tunnel Load Balancing Policies” on page 139
- “Defining Default Assemblers” on page 141
- “Defining Preferred Assemblers” on page 144
- “Configuring Tunnel Mode Settings” on page 145
- “Configuring Pre-Synchronization for Network Sequence Caching” on page 146
- “Configuring Tunnel Switching” on page 148

## Viewing and Fetching Remote Routes

Remote routes are the reduction subnets advertised by the other WX devices in the community. Each device can reduce only the traffic that is destined for a remote route advertised by another WX device. You can view the remote routes to determine which routes are advertised by multiple devices. You can also specify how often remote routes are fetched from the other devices, and enable a test to validate each remote route. The set of subnets advertised by each device is called a netmap.



**NOTE:** The remote routes shown for a device may not match the list of advertised subnets shown on the device. Each device dynamically adjusts its advertised subnets to exclude (carve out) unreachable addresses. To exclude an address from an advertised subnet, multiple smaller subnets are generated, so that one advertised subnet may produce several remote routes. To disable the carve-out feature, refer to “configure reduction-subnet” on page 350.

To view the remote routes:

1. Click REDUCTION in the menu frame, click ADVANCED in the left-hand navigation frame, and then click Remote Routes.

**Figure 74: Displaying and Updating Remote Routes**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Application Filter

ADVANCED

- Remote Routes**
- Load Balancing
- Default Assemblers
- Preferred Assemblers
- Tunnel Mode

**Remote Routes**

8 subnets are advertised by remote Peribit devices in this community

IP Address	Subnet Mask	Assembler IP	Cost	Route Type	Last Update
192.168.5.0	255.255.255.240	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.16	255.255.255.248	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.24	255.255.255.252	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.28	255.255.255.254	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.30	255.255.255.255	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.32	255.255.255.224	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.64	255.255.255.192	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.128	255.255.255.128	192.168.5.131	1	Dynamic	09-23-2003, 10:16

It is recommended that the Route Validation feature be turned off when a load balancing policy is in effect.

Fetch advertised routes: Hourly  ☐ Validate advertised routes

The **Assembler IP** column shows the address of one or more remote devices that can assemble data for the specified subnet. The **Cost** column indicates the relative cost of the route for each device. Static routes have the highest cost (1000). The lowest cost device is used whenever possible. Load balancing can be used when multiple devices have equal cost paths, as described in “Configuring Tunnel Load Balancing Policies” in the next section.

2. To change how often the remote routes are fetched from the other WX devices in the community, select a frequency from the drop-down menu at the bottom of the page. To update the remote routes immediately, click Fetch Now.

Note that remote routes are advertised each time a device starts, and route changes are advertised as soon as they occur. However, if reduction is disabled, the advertisement of route changes depends on the fetch interval. 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 WX 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” on page 139).


4. Click Submit to activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring Tunnel Load Balancing Policies

If two or more WX devices in the same community have equal cost paths to the same subnet, you can use tunnel 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 144. If neither load balancing nor preferred assemblers are used, the path selection is arbitrary.

For example, in Figure 75, devices D2 and D3 advertise a local route to Subnet 2. On D1, the two routes to Subnet 2 have equal cost paths, and are grouped together in the Remote Routes page (Figure 76).

**Figure 75: Configuring Tunnel Load Balancing Policies**



**Figure 76: Remote Routes with Equal Cost Paths**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Application Filter

ADVANCED

- Remote Routes
- Load Balancing
- Default Assemblers
- Preferred Assemblers
- Tunnel Mode

**Remote Routes**

8 subnets are advertised by remote Peribit devices in this community

IP Address	Subnet Mask	Assembler IP	Cost	Route Type	Last Update
192.168.5.0	255.255.255.240	192.168.5.131	1	Dynamic	09-23-2003, 10:16
		192.168.52.22	1	Dynamic	09-23-2003, 10:16
192.168.5.16	255.255.255.248	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.24	255.255.255.252	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.28	255.255.255.254	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.30	255.255.255.255	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.32	255.255.255.224	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.64	255.255.255.192	192.168.5.131	1	Dynamic	09-23-2003, 10:16
192.168.5.128	255.255.255.128	192.168.5.131	1	Dynamic	09-23-2003, 10:16

Common destinations with equal cost paths

It is recommended that the Route Validation feature be turned off when a load balancing policy is in effect.

Fetch advertised routes: Hourly Fetch Now ☐ Validate advertised routes

Submit Reset

To configure tunnel load balancing policies:

1. Click REDUCTION in the menu frame, click ADVANCED in the left-hand navigation frame, and then click Load Balancing.

**Figure 77: Configuring Tunnel Load Balancing Policies**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Application Filter

ADVANCED

- Remote Routes
- Load Balancing
- Default Assemblers
- Preferred Assemblers
- Tunnel Mode

**Load Balancing**

The rule selected below determines how traffic is routed when more than one reduction tunnel exists for a given subnet. The selected rule also applies to Default Assemblers if more than one has been specified.

It is recommended that the Route Validation feature be turned off when a load balancing policy is in effect.

☒ Off All traffic is routed to one of the available tunnels.

☐ 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, i.e. round robin.

☐ Flow based Traffic is distributed over available tunnels based on source and destination IP addresses and ports.

Submit Reset

- 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.
2. Click Submit to activate the changes, or click Reset to discard them.
  3. To retain your changes when the device is restarted, click SAVE in the menu frame.

### **Defining Default Assemblers**


You can sometimes simplify route administration by designating a WX 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 (such as to another WX device). 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 (refer to "Filtering Data Reduction by Source and Destination" on page 101).

Figure 78 shows a simple example of a remote site with one outbound connection to the corporate network. If D1 is the default assembler for D2, all traffic that uses the default route on D2 is reduced and sent to D1.

**Figure 78: Setting a Default Assembler**

To disable data reduction for traffic sent to subnet S4, you can add S4 to the exclusion list on D2. You can specify up to six default assemblers. If you specify more than one default assembler, the current load balancing policies are applied (refer to “Configuring Tunnel Load Balancing Policies” on page 139).



**NOTE:** Default assemblers can be used only in live operation (not in Profile Mode). Be sure you understand which WX devices will use a default assembler, and which local routes the default assembler supports.

To define default assemblers:

1. Log in to the device where you want to specify default assemblers.
2. Click REDUCTION in the menu frame, click ADVANCED in the left-hand navigation frame, and then click Default Assemblers.

**Figure 79: Creating a Default Assembler List**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Application Filter

ADVANCED

- Remote Routes
- Load Balancing
- Default Assemblers**
- Preferred Assemblers
- Tunnel Mode

**Default Assemblers**

If traffic is destined for a subnet that is NOT included in the Remote Routes list, it is normally passed through without reduction. However, if a Default Assembler (a Peribit device) is entered below, this traffic will be reduced and routed to the Default Assembler.

If a Default Assembler is entered, certain traffic can be excluded from this feature by entering the destination IP address/subnet mask in the Exclude List below.

**Default Assemblers**

Enter IP addresses of Peribit devices, one per line. A maximum of 6 Default Assemblers may be entered. If more than one Default Assembler is entered, then the Load Balancing policy will be applied. If Load Balancing is set to "Off", then the precedence of the Default Assemblers will be based on their order in the list.

**Exclude List**

Enter addresses/subnets, one per line. For an individual host, enter the IP address only. For a subnet, enter the IP address and subnet mask separated by a slash (/). Examples:  
123.123.123.123  
123.123.123.1/255.255.255.0

Submit Reset

3. 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.
4. 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 another WX device, the exclusion is ignored.
5. Click Submit to activate the changes, or click Reset to discard them.
6. Log in to each default assembler you specified and, 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). To add a static route, refer to "Configuring Local Routes" on page 66.
7. To retain your changes when the device is restarted, click SAVE in the menu frame.

## 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 139. 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 80, data from Subnet 1 has two network paths to Subnet 2. If the WX 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 80: Designating a Preferred Assembler**



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

To create a list of preferred assemblers:

1. Click REDUCTION in the menu frame, click ADVANCED in the left-hand navigation frame, and then click Preferred Assemblers.

**Figure 81: Defining Preferred Assemblers**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Application Filter

ADVANCED

- Remote Routes
- Load Balancing
- Default Assemblers
- Preferred Assemblers**
- Tunnel Mode

**Preferred Assemblers**

If more than one reduction tunnel exists for a given subnet and no other policies apply, traffic will be routed to one of the tunnels on an arbitrary basis. However, you can give precedence to a particular tunnel by entering the associated assembler in the list of Preferred Assemblers below.

If the Load Balancing policy is set to anything other than "Off", then it will override this policy.

Preferred Assemblers

Enter IP addresses of Peribit devices, one per line. If more than one Preferred Assembler is contending for the same traffic, then precedence will be based on the order in which they appear in the list.

Submit Reset



2. Enter the IP address of a remote preferred assembler. You can specify up to 80 preferred assemblers (one per line).
3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Configuring Tunnel Mode Settings

The tunnel mode determines how a WX device sends reduced traffic to the remote WX 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. Click REDUCTION in the menu frame, click ADVANCED in the left-hand navigation frame, and then click Tunnel Mode.

**Figure 82: Configuring Tunnel Mode Settings**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Application Filter

ADVANCED

- Remote Routes
- Load Balancing
- Default Assemblers
- Preferred Assemblers
- Tunnel Mode**

**Tunnel Mode**

Select the tunnel mode below. IPComp is the default setting and provides optimum reduction in most environments.

In some networks, edge devices make use of packet information such as source and destination ports and IP addresses. This information is normally not available in the reduction tunnel. If this information is required in your network, it can be made available by selecting the appropriate option below.

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

☐ UDP All reduction tunnel traffic is transmitted on port 3577. Reduced traffic appears as a single flow between two Peribit devices.

☐ Multi-flow emulation Different traffic flows can be distinguished in the reduction tunnel.  
Number of flows supported:  (Range 256-1024)

☐ Application visibility Original source/destination port information is preserved in the reduction tunnel.

Submit Reset

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.   |
3. Click Submit to activate the changes, or click Reset to discard them.
  4. To retain your changes when the device is restarted, click SAVE in the menu frame.

### **Configuring Pre-Synchronization for Network Sequence Caching**

On WXC devices where Network Sequence Caching (NSC) is enabled, “pre-synchronization” can be used to improve user response times and reduction rates for large files, such as database files and software updates. During pre-synchronization, the repeated patterns in the files are added to the reduction dictionaries of the selected devices, so that reduction occurs when the first user requests the files.

With standard tools such as Perl and Cron, you can put the pre-sync CLI commands in a script and schedule pre-synchronization to occur automatically during off-peak hours (refer to the CLI commands in “Pre-Synchronization” on page 346).

Pre-synchronization has the following requirements:

- The files must be loaded on an FTP server.
- The WXC where you configure pre-synchronization must have the FTP server on its LAN side so that it can reduce all the traffic sent from the FTP server to the remote WXC devices.
- NSC must be enabled between WXC devices, refer to “Configuring Network Sequence Caching” on page 134.
- NSC must be enabled for the application that users will access to retrieve the files (refer to “Reducing Applications” on page 135).

To configure pre-synchronization for NSC:

1. Log in to a WXC that has the FTP server on its LAN side.



**NOTE:** Do not configure pre-synchronization on an off-path device that uses RIP for packet interception. Traffic from the FTP server will be routed directly to the remote WXCs without being reduced.

---

2. Click REDUCTION in the menu frame, click ADVANCED in the left-hand navigation frame, and then click Pre-Sync.

**Figure 83: Configuring Pre-Synchronization for Network Sequence Caching**

**Peribit-SM**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Reduction**

BASIC

- Endpoints
- Reduction Subnets
- Network Sequence Mirror
- Application Filter

ADVANCED

- Remote Routes
- Load Balancing
- Default Assemblers
- Preferred Assemblers
- Tunnel Mode
- Pre-Sync**

**Pre-Sync**

This page allows you to increase the efficiency with which the data dictionary is created in NSM-capable Peribit devices. This is done by sending large files to the checked Peribit devices before they are requested by a client.

Enter file pathnames for FTP server's below (one per line). Then check the Peribit devices from the list below. When you click "Go", the file(s) will be sent to each of the Peribit devices in turn.

FTP file pathnames are of the form `ftp://host:port:user:password/path` where some or all of the parts `:user:password`, `:password`, `:port` and `/path` may be excluded.

File Pathname

IP address Device name

☐ 192.168.53.2 53/2-SR55

Select All Clear

Go

- Specify the FTP server location of each file (one file per line). The general format is:

`ftp://host:port:user:password/path`

Where:

- **host**. FTP server name or IP address.
  - **:port**. FTP port number. Omit if port 21 is used.
  - **:user:password**. FTP user name and password. Omit if server allows anonymous access.
  - **/path**. File location on the server.
- Select the check box next to the IP address of each WXC device where you want to send the specified file(s). To select all devices, click Select All.



**NOTE:** The FTP server must be reachable from each remote device.

- Click Go to send the files to the selected devices.

To view the results of the last 50 pre-synchronization tasks, enter the following CLI command:

```
show reduction network-sequence-mirroring
```

## Configuring Tunnel Switching

Each WX device can perform data reduction (form reduction tunnels) for a varying number of remote WX devices, depending on the device type. For example, an WX 20 can support five reduction tunnels. Tunnel switching allows each device to reduce data for every other WX device in the network, without having to form a reduction tunnel with each remote device.

To optimize performance, devices can be deployed hierarchically in either of the following ways:

- Assign devices to separate communities (only devices in the same community can reduce and assemble data for each other)
- Apply hub and spoke topology designations to selected devices in the same community (by default, spoke devices reduce and assemble data only for hub devices)

Tunnel switching can be used to send compressed traffic between devices in separate communities or between spoke devices associated with the same hub or different hubs.

If outbound QoS is enabled, bandwidth management is applied to tunnel-switched packets in the normal manner.




**NOTE:** Inbound QoS applies only to traffic received on the Remote interface. When 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.

---

### Tunnel Switching Between Communities

You can organize devices into hierarchical communities by assigning selected devices to multiple communities. On these common devices, which act as default assemblers for traffic leaving the community, you must enable tunnel switching to reduce the traffic for the next device in the path.

In Figure 84, tunnel switching is enabled on D2 and D3 in Community 3. These are the common devices that convey reduced traffic between Community 1 and Community 2.

**Figure 84: Example of Tunnel Switching Between Communities**

When D1 in Figure 84 encounters traffic destined for a subnet advertised by D4, the following processing occurs:

1. D1 cannot match the destination to an advertised reduction subnet (D4 is in a separate community), so D1 reduces the traffic and sends it to the default assembler (D2).
2. D2 assembles the traffic from D1, matches the destination to a reduction subnet advertised by D3, recompresses the traffic, and sends it to D3. If tunnel switching is disabled on D2, the traffic is sent to D3 uncompressed.
3. D3 assembles the reduced traffic from D2, matches the destination to a reduction subnet advertised by D4, recompresses the traffic, and sends it to D4. If tunnel switching is disabled on D3, the traffic is sent to D4 uncompressed. If the traffic from D2 is uncompressed, the traffic is compressed and sent to D4 (no recompression).
4. Traffic from D4 to D1 follows the reverse path, with D3 serving as the default assembler for D4 (and the other devices in Community 2).

Note that Figure 84 can be simplified by omitting Community 3 and assigning D2 and D3 to both Community 1 and 2. In that way, D2 can send recompressed traffic directly to D4, without requiring assembly and recompression on D3.

### Procedure for Configuring Tunnel Switching Between Communities

To configure tunnel switching between communities:

1. Identify the common devices in each community that convey reduced traffic between communities.

2. On each of the other devices in a community, designate the common device as the default assembler (refer to “Defining Default Assemblers” on page 141). Alternatively, on each common device, you can manually define static routes for the external subnets that you want to advertise to the other devices in the community.
3. On each common device, do the following:
  - Enable tunnel switching and disable LAN/WAN checking (refer to “configure reduction” on page 343). For off-path devices, LAN/WAN checking is disabled by default.
  - Add static routes for the reduction subnets to be advertised to the other common devices (refer to “Adding Static Routes” on page 68). For example, on D2 in Figure 84, define static routes for the reduction subnets in Community 1. Note that one or two static routes may be sufficient, depending on the subnet addressing scheme.



**NOTE:** Define static routes carefully to avoid the creation of routing loops.

---

4. If necessary, enable WAN reduction subnets on the common devices (refer to “configure reduction-subnet” on page 350). For example, in Figure 84, if the other devices in Community 1 are on the Remote (WAN) side of D2, WAN reduction subnets must be enabled on D2 so that the appropriate subnets can be advertised to D3. In this case, since D2 is an off-path device, WAN reduction subnets are enabled by default.

### Tunnel Switching Between Hub and Spoke Devices

In the same community, you can organize devices into multiple sets of hubs and spokes, and then use tunnel switching to send reduced traffic between any two spoke devices. By default, spoke devices reduce and assemble data only for the hubs. You can further restrict each spoke to work only with a specific hub.


Tunnel switching for hub and spoke topologies works in the same manner as for hierarchical communities. Figure 85 is very similar to Figure 84, except that all the devices are in the same community, and tunnel switching is enabled on the hub devices (D2 and D3).

Note that tunnel switching can be used between spokes on the same hub, such as when you add new spokes to a hub and some existing spokes cannot support any additional direct tunnels.



**NOTE:** 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.

---

**Figure 85: Example of Tunnel Switching Between Spoke Devices**

To configure tunnel switching between spoke devices:

1. Identify the hub devices that convey reduced traffic between spokes.
2. On each of the spoke devices, designate the hub device as the default assembler (refer to “Defining Default Assemblers” on page 141). Alternatively, on each hub device, you can manually define static routes for the external subnets that you want to advertise to the other devices in the community.
3. On the hub devices, enable tunnel switching and disable LAN/WAN checking (refer to “configure reduction” on page 343). In some routing environments, on each hub you may need to add static routes for the reduction subnets associated with the spoke devices. Note that one or two static routes may be sufficient, depending on the subnet addressing scheme.



**NOTE:** Define static routes carefully to avoid the creation of routing loops.

4. If necessary, enable WAN reduction subnets on the hub devices (refer to “configure reduction-subnet” on page 350). For example, in Figure 84, if the spoke devices for Hub 1 are on the Remote (WAN) side of D2, WAN reduction subnets must be enabled on D2 so that the appropriate subnets can be advertised to D3.
5. If necessary, define static routes between the hub devices (refer to “Adding Static Routes” on page 68).





## Chapter 6

# Applying Quality of Service (QoS) Policies

The following sections describe how WX devices use Quality of Service (QoS) policies to allocate WAN bandwidth to your network applications:

- “Using Outbound QoS to Enhance Performance” in the next section
- “Understanding Outbound Bandwidth Management” on page 154
- “Configuring Outbound QoS Policies” on page 162
- “Configuring Inbound QoS Policies” on page 184
- “Summary of Key Terms” on page 187

## Using Outbound QoS to Enhance Performance

---

Outbound QoS provides two key benefits:

- **Basic bandwidth allocation.** Data reduction performance is automatically optimized based on the WAN speeds, 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 WX 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 devices upstream from the WX device. 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 WAN circuit speeds, as described in “Defining Outbound QoS Endpoints” on page 176. For guidance on adjusting the WAN speeds to account for router overhead, refer to “WAN Circuit Speeds and Router Overhead” on page 156.
2. Start outbound QoS using Weighted Fair Queuing (WFQ) or Weighted Strict Priority (WSP), as described in “Starting and Stopping Outbound QoS” on page 183. Unless you need strict priority treatment for traffic classes, WFQ is recommended.

## Understanding Outbound Bandwidth Management

---

If all WAN traffic goes through the WX device, then outbound QoS policies can control how the entire WAN bandwidth is allocated to all contending applications, regardless of whether traffic is being reduced. 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 WX 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 155
- “QoS Templates and Endpoints” on page 155
- “WAN Circuit Speeds and Router Overhead” on page 156
- “Dedicated, Oversubscribed, and Variable Rate WANs” on page 157
- “Direct Setup Versus Wizard Configuration Results” on page 159
- “Class Priorities and Excess Bandwidth Allocation” on page 161
- “ToS/DSCP Values” on page 162

- “Unadvertised Subnets” on page 162
- “Summary of Key Terms” on page 187

## 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 86 shows the default settings for the standard traffic classes created by the Setup Wizard. You can have up to 16 traffic classes.

**Figure 86: Predefined 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 %

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 161).

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 “Managing Applications” on page 88).

## 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 WX 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 105.
- **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. You can view the bandwidth reports for prime time or non-prime time hours (refer to “Outbound Bandwidth Statistics” on page 244).

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

### **WAN Circuit Speeds and Router Overhead**

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

- **Outbound speed.** The sum of the WAN circuit speeds on the adjacent router that conduct traffic from the WX device. You must specify the outbound speed only if it is less than the sum of the remote WAN speeds—that is, if the WAN is “oversubscribed” (refer to “Dedicated, Oversubscribed, and Variable Rate WANs” on page 157).
- **Endpoint circuit speeds.** The maximum WAN circuit speed associated with each remote WX device or virtual endpoint for which you want to manage the outbound bandwidth. You can use the Ethernet speed for a remote WX device if you enable congestion control for that 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, the endpoint WAN circuit speeds 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.

For an oversubscribed WAN, always set the outbound speed as accurately as possible, even if congestion control is enabled.

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.

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 WX 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 by the WX device, 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 87, the outbound speed for WX device D1 is 1.5 Mbps, which equals the total speed of the remote endpoints for D2 and 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 and D3 (.6 x 500 Kbps). However, in dedicated mode, Other Traffic is unconstrained by QoS.

**Figure 87: Dedicated WAN**




You can specify a dedicated WAN as “oversubscribed” if you want the Other Traffic to be managed by QoS.

- **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 88, the WAN is oversubscribed from the perspective of 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 ( $.8 \times 1500 = 1200$  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 88: 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 WAN speed specified for each remote endpoint is not critical. For example, you can enable congestion control for all remote endpoints, and then specify the WX Ethernet speed for each remote device. For oversubscribed WANs, the outbound speed of the adjacent WAN router must be specified accurately (refer to “WAN Circuit Speeds and Router Overhead” on page 156).



**NOTE:** Congestion control manages only traffic sent to other WX endpoints. In oversubscribed mode, if you have substantial passthrough traffic for non-WX 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 176).

## 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 3 shows the Wizard and direct setup results when D1 in Figure 88 is configured with two traffic classes and the same guaranteed bandwidths for each endpoint.

**Table 3: 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 Business	15 % 40 %	15 % 40 %	75 Kbps 200 Kbps	10.92 % 29.12 %	54 Kbps 145 Kbps
E2	1500 Kbps	Default Business	15 % 40 %	15 % 40 %	225 Kbps 600 Kbps	10.92 % 29.12 %	163 Kbps 436 Kbps
E3	60 Kbps	Default Business	15 % 40 %	15 % 40 %	9 Kbps 24 Kbps	10.92 % 29.12 %	6 Kbps 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 3 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 E1 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 89 shows the “Oversubscribed” template specifying the 15 % and 40 % guarantees; Figure 90 shows the guaranteed bandwidths in Kbps displayed on the Outbound QoS Overview page when the template is applied to each endpoint.

**Figure 89: Oversubscribed Template for Device D1**

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 90: Direct Setup Results on the Outbound QoS Overview Page for Device D1**

Endpoint		Template	Circuit Speed (Kbps)	Traffic Classes		Total Guaranteed Bandwidth
				Default	Business	
Other traffic	<a href="#">EDIT</a>	Oversubscribed	60	9	24	33
192.168.53.5	<a href="#">EDIT</a>	Oversubscribed	500	75	200	275
192.168.52.22	<a href="#">EDIT</a>	Oversubscribed	1500	225	600	825
<b>Total</b>				<b>309</b>	<b>824</b>	<b>1133</b>

### 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 3. 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 E1 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 D1. Figure 91 shows the guaranteed bandwidths in Kbps generated by the Setup Wizard and displayed on the Outbound QoS Overview page.

**Figure 91: Wizard Results on the Outbound QoS Overview Page for Device D1**

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

The Wizard adjusts the bandwidths for oversubscribed WANs only when there are multiple remote endpoints. For example, in Figure 88 on page 158, 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. The allocation depends on whether the WAN is dedicated or oversubscribed.

**Table 4: WFQ Allocation of Excess Bandwidth**

WAN Type	Excess Bandwidth Allocation
Dedicated	<p>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> $(\text{Class Priority} + 1) / (\text{Sum of active class priorities} + 1 \text{ for each class})$ <p>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>Note that if only one class has traffic, then that class receives 100 % of the bandwidth.</p> <p>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>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>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>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>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>Note that you must calculate the sum of the guaranteed bandwidths for each class on each remote endpoint. For the example in Table 3 on page 159, 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 WX “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 181.

If necessary, queue processing can be determined by the ToS/DSCP values of the incoming traffic (refer to “Processing Queues Based on Incoming ToS/DSCP Values” on page 184).

## **Unadvertised Subnets**

For an oversubscribed WAN, traffic to 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 WX 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 “Filtering Data Reduction by Source and Destination” on page 101).

By default, each WX 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.

## **Configuring Outbound QoS Policies**

---

This section describes how to configure outbound QoS policies for bandwidth management, and covers the following topics:

- “Procedure for Configuring Outbound QoS Policies” on page 163
- “Using the Outbound QoS Setup Wizard” on page 164
- “Defining Outbound QoS Settings by Endpoint” on page 171
- “Defining Traffic Classes” on page 173
- “Defining Outbound QoS Templates” on page 174
- “Defining Outbound QoS Endpoints” on page 176
- “Changing Outbound ToS/DSCP Values” on page 181
- “Starting and Stopping Outbound QoS” on page 183

## Procedure for Configuring Outbound QoS Policies

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

1. For best results, verify that each WX device advertises all the subnets it can access. In oversubscribed mode, traffic to unadvertised subnets is 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 “Filtering Data Reduction by Source and Destination” on page 101).
2. Run the Setup Wizard or specify the outbound QoS policies directly:
  - To run the Setup Wizard, 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 “Defining Traffic Classes” on page 173).
  - b. Define one or more templates to specify the priorities and bandwidths for each traffic class (refer to “Defining Outbound QoS Templates” on page 174).
  - c. Specify the local outbound speed (if the WAN is oversubscribed) and the maximum circuit speeds for each remote endpoint (refer to “WAN Circuit Speeds and Router Overhead” on page 156 and “Defining Outbound QoS Endpoints” on page 176).
  - d. Assign a prime-time and nonprime-time template to each endpoint (refer to “Defining Outbound QoS Settings by Endpoint” on page 171).
  - e. Enable QoS and select a queuing model (refer to “Starting and Stopping Outbound QoS” on page 183).
- 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 171).
  - Change traffic class names (refer to “Defining Traffic Classes” on page 173).
  - Add new templates, change a template name, or change just one of the Wizard templates (refer to “Defining Outbound QoS Templates” on page 174)

- Define virtual endpoints or exclude address or subnet pairs from bandwidth management (refer to “Defining Outbound QoS Endpoints” on page 176).
- Change the ToS/DSCP values for one or more traffic classes (refer to “Changing Outbound ToS/DSCP Values” on page 181).

### 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 105.
- **Wizard-NonPrimeTime.** Applies to nonprime time hours. To view QoS reports for prime time or nonprime time hours, refer to “Outbound Bandwidth Statistics” on page 244.

Each time you run the Setup Wizard, both of the Wizard templates and all customized settings are overwritten. To change just one of the templates, refer to “Defining Outbound QoS Templates” on page 174.

To run the outbound QoS Setup Wizard:

1. Click QOS in the menu frame, and click Setup Wizard in the left-hand navigation frame.
2. Click Enable Outbound QoS and click Next.



3. Select one of the following WAN modes and click Next:

- |                    |   |
|--------------------|---|
| Dedicated Circuits | Indicates that the local outbound WAN speed equals or exceeds the sum of the WAN speeds for the remote endpoints whose bandwidths you want to manage (the default). In dedicated mode, traffic sent to non-WX endpoints (“Other traffic”) is unconstrained by QoS.<br><br>If the WAN is dedicated, but you want Other Traffic to be managed by QoS, select Oversubscribed and use the default outbound speed. |
| Over-subscribed    | Indicates that the local outbound WAN speed is less than the sum of the remote WAN speeds. Add up the speeds of all the WAN interfaces on the adjacent router, and enter the total in the <b>Outbound Speed</b> field. Be sure to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 156).  |

**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
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.

4. Select the check box next to the IP address of each remote WX device (endpoint) for which you want to manage the outbound bandwidth (or click Select All), and enter the maximum remote WAN circuit speed (in Kbps) for each selected endpoint. If Packet Flow Acceleration is enabled for an endpoint, outbound QoS cannot be disabled.



**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 156). Exceeding the actual WAN speed effectively shifts bandwidth management to the router, and may cause the router to drop traffic.

Note the following:

- For WX devices that support Multi-Path, a “\_Pri” and “\_Sec” are appended to the device name to indicate the primary and secondary path. You can enable QoS for one or both paths. To configure Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 115.

- In oversubscribed mode, the two generated templates are also applied to the “Other traffic” endpoint, which is used to manage the bandwidth for all traffic that is not sent to one of the selected WX devices. The circuit speed for “Other traffic” defaults to the outbound speed.
- If any “No Remote Peribit” endpoints have been defined to manage the traffic sent to remote subnets that do not have a WX device (refer to “Defining Outbound QoS Endpoints” on page 176), you can change their circuit speeds or disable them. You can change the settings for “Other traffic” and non-WX endpoints in the same manner as other endpoints (refer to “Defining Outbound QoS Settings by Endpoint” on page 171).

5. 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 38: 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"/>

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.

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 WX 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.



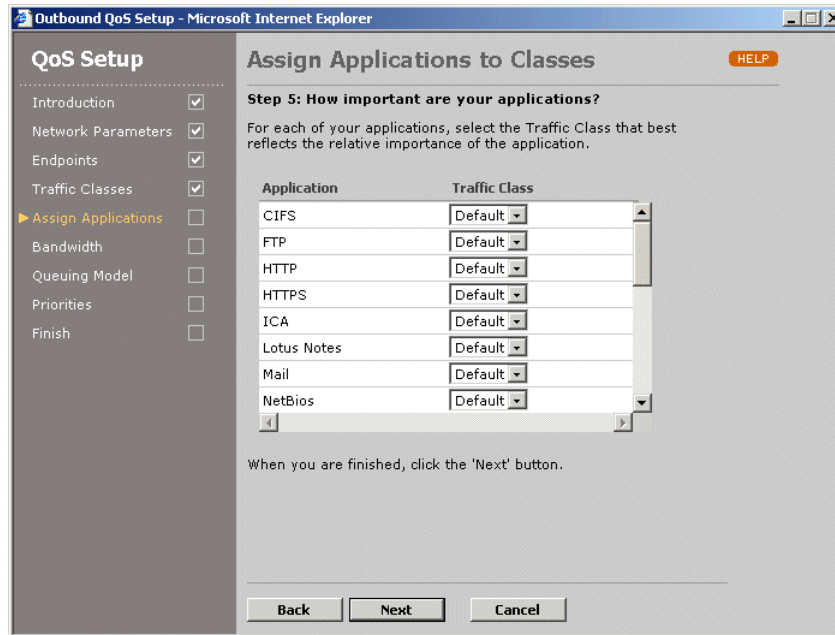
**NOTE:** Congestion control manages only traffic sent to other WX endpoints for which reduction tunnels are enabled. In oversubscribed mode, if you have substantial passthrough traffic for non-WX destinations, you may want to reduce the maximum speed for the “Other traffic” endpoint to limit the bandwidth allocated to passthrough traffic. After you complete the Wizard configuration, refer to “Defining Outbound QoS Endpoints” on page 176.

7. Click Next. To define custom traffic classes, click Custom, and then click Next.

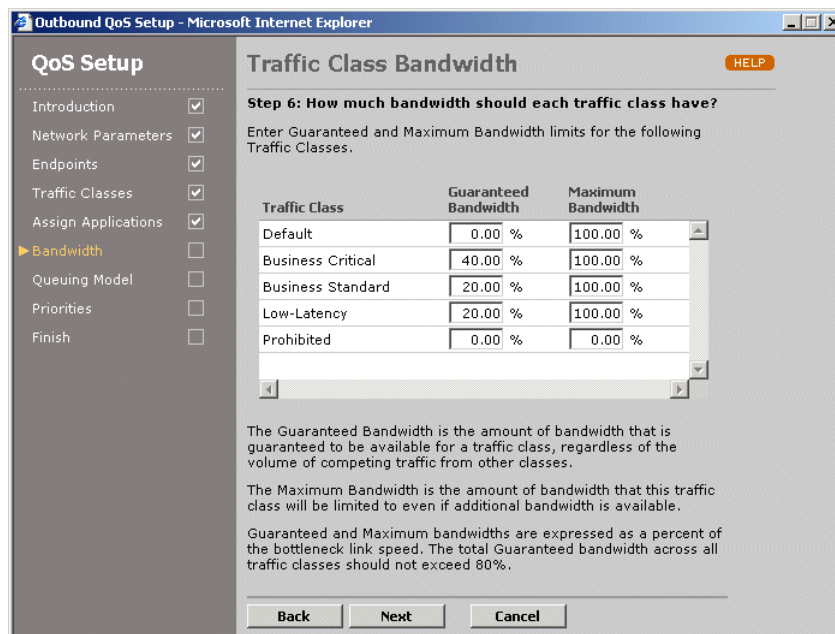
The screenshot shows the 'Outbound QoS Setup - Microsoft Internet Explorer' window. On the left is a 'QoS Setup' sidebar with a list of steps: Introduction (checked), Network Parameters (checked), Endpoints (checked), Traffic Classes (selected), Assign Applications, Bandwidth, Queuing Model, Priorities, and Finish. The main area is titled 'Traffic Classes' and contains 'Step 4: Create Traffic Classes'. It instructs the user to create up to 15 traffic classes by entering a name and clicking 'Add', or to delete checked classes by clicking 'Delete'. A list box shows 'Default' as the only traffic class. At the bottom are 'Back', 'Next', and 'Cancel' buttons. A 'HELP' button is in the top right corner.

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. Note that the Default class is reserved for undefined application traffic and cannot be deleted. Click Next.





9. Select the appropriate traffic class for each of your defined applications. If one of your network applications is not shown, you must create a application definition for it, as described in “Managing Applications” on page 88. Click Next.






10. Enter the bandwidth information for each traffic class, and click Next.

Guaranteed Bandwidth	<p>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 local outbound WAN speed.</p>
Maximum Bandwidth	<p>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.</p>



**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 excess bandwidth as load conditions change, and click Next. The excess 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 excess 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 excess bandwidth is allocated by processing the queues based on their priority. Processing is weighted equally for traffic classes that have the same priority.

**Outbound QoS Setup - Microsoft Internet Explorer**

**QoS Setup**

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

**Priorities** HELP

**Step 8: Set priorities for each traffic class**

Select a priority for each of the traffic classes. Choose a higher priority number for important traffic classes -- a lower priority number for unimportant traffic classes.

Traffic Class	Priority
Default	0 (Lowest)
Business Critical	0 (Lowest)
Business Standard	0 (Lowest)
Low-Latency	7 (Highest)
Prohibited	0 (Lowest)

Back Next Cancel

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 excess (unguaranteed) bandwidth to the competing traffic classes.



**NOTE:** 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.
14. Click QOS in the menu frame to refresh the Outbound QoS Overview page, which now shows the template name, circuit speed, and guaranteed bandwidths for each endpoint.
15. To retain your changes when the device is restarted, click SAVE in the menu frame.

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 176.

To view or change the outbound QoS settings by endpoint:

1. Click QOS in the menu frame to open the Outbound QoS Overview page.

**Figure 92: Outbound QoS Overview**

**Peribit 1**

SETUP REDUCTION **QOS** ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**QoS**

OUTBOUND  
Setup Wizard...  
**Overview**  
Direct Setup

INBOUND

**Outbound QoS Overview** HELP

Display: Guaranteed Bandwidth

Show bandwidth as: % of circuit speed

Time Frame: Prime Time

Outbound Speed: 100000 Kbps

Outbound QoS Setting: Weighted Fair Queueing Bandwidth Allocation

Update

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

The Outbound QoS Overview page shows the outbound speed for the WX device, the selected queuing model, and the template name, circuit speed, and guaranteed bandwidths for each remote endpoint. In oversubscribed mode, 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 specify bandwidth percentages as circuit speeds (Kbps). Percentages must be used if congestion control is enabled.



**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 remote speeds are the true WAN speeds (not the WX 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 105).
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 93: Changing Endpoint Templates or Template Settings**

The screenshot shows the Peribit 1 QoS configuration interface. The left sidebar has a 'QoS' section with 'OUTBOUND' and 'INBOUND' sub-sections. The 'OUTBOUND' section is expanded, showing 'Setup Wizard...', 'Overview', and 'Direct Setup'. The 'Overview' link is selected. The main content area is titled 'Outbound QoS Overview > Other traffic'. It contains a description: 'This page determines bandwidth limits for Outbound QoS to the selected endpoint.' Below this, the 'Endpoint' is 'Other traffic', 'Circuit Speed' is '1500 Kbps', and 'Time Frame' is 'Prime Time'. There are two radio buttons: 'Use QoS template' (selected) and 'Use custom setting'. A dropdown menu shows 'Wizard-PrimeTime'. Below this, a 'Show bandwidth as' dropdown is set to '% of circuit speed'. A table shows bandwidth limits for various traffic classes. A note on the right states: 'Bandwidth limits are stored as a percent of circuit speed. If circuit speed is modified, the bandwidth Kbps values will also change.' At the bottom are 'Submit' and 'Cancel' buttons.

Traffic Class	Priority	Guaranteed Bandwidth	Maximum Bandwidth
Default	0 (Lowest)	0.00 %	100.00 %
Business Critical	0 (Lowest)	13.33 %	100.00 %
Business Standard	0 (Lowest)	6.67 %	100.00 %
Low-Latency	7 (Highest)	6.67 %	100.00 %
Prohibited	0 (Lowest)	0.00 %	0.00 %
<b>Total</b>		<b>26.67 %</b>	

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 174.
  - 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. Do not include leading zeros on bandwidths (they indicate octal values).

To increase the guaranteed bandwidth for a traffic class on an oversubscribed WAN, you may have to decrease the bandwidth on another class (on the same endpoint or a different endpoint), reduce the remote 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 164).

5. To retain your changes when the device is restarted, click SAVE in the menu frame.

If you override the template settings for an endpoint, the template name is changed to None on the Outbound QoS Overview page. To restore the original settings, reapply the template.

If you customize the settings for specific endpoints, rather than change the template, new templates are created whose names include the IP address of the endpoint. To view these templates, use the “show -run qos outbound” CLI command on the device.

- PTO- <IP\_address> for customized prime-time templates
- NTO- <IP\_address> for customized nonprime-time templates

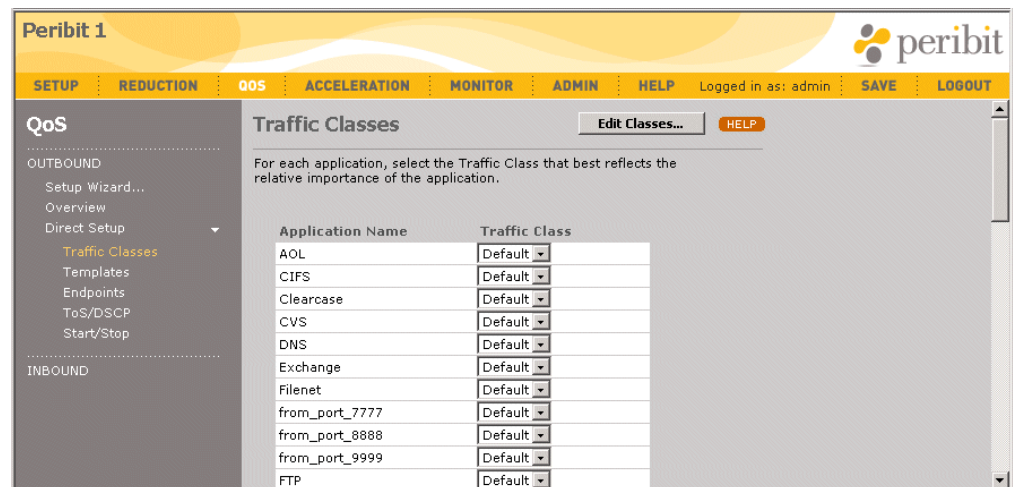
## Defining Traffic Classes

Outbound application traffic is managed by traffic class. You can assign one or more applications to each of the predefined traffic classes provided by the Setup Wizard or create your own classes. Initially, all applications belong to the Default class. Note that an application can belong to only one traffic class, but it can belong to different classes on different WX devices. You can have up to 16 traffic classes.

To define traffic classes and add applications to a class:

1. Click QOS in the menu frame, click Direct Setup in the left-hand navigation frame, and then click Traffic Classes.

**Figure 94: Defining Outbound QoS Traffic Classes**



2. To change the applications assigned to each traffic class, select the appropriate traffic class for each application, and click Submit.
3. To add or change the current traffic classes, click Edit Classes.

From the Traffic Classes > Edit Classes page, you can:

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

- 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. All 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.



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

4. To retain your changes when the device is restarted, click SAVE in the menu frame.


## Defining Outbound QoS Templates

Outbound QoS templates specify the priority, guaranteed bandwidth, and maximum bandwidth for each traffic class. You can change the templates created by the Setup Wizard or create new templates. To apply a QoS template to an endpoint, refer to “Defining Outbound QoS Settings by Endpoint” on page 171.

To define outbound QoS templates:

1. Click QOS in the menu frame, click Direct Setup in the left-hand navigation frame, and then click Templates.

**Figure 95: Defining Outbound QoS Templates**



From the Outbound QoS Templates page, you can:

- Add a new template, as described in Step 2 through Step 4.
- 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.

Figure 96: Defining a New QoS Template

Peribit 1

SETUP REDUCTION QoS ACCELERATION MONITOR ADMIN HELP Logged in as: admin LOGOUT

**QoS**

OUTBOUND

- Setup Wizard...
- Overview
- Direct Setup
- Traffic Classes
- Templates
- ToS/DSCP
- Start/Stop

INBOUND

**Outbound QoS Templates > New Template**

HELP

Template Name

Traffic Class	Priority	Bandwidth Limit (%)	
		Guaranteed	Maximum
Default	0 (Lowest)	0	100
Business Critical	0 (Lowest)	0	100
Business Standard	0 (Lowest)	0	100
Low-Latency	0 (Lowest)	0	100
Prohibited	0 (Lowest)	0	100

For each traffic class, select a priority and then set the guaranteed and maximum bandwidth limits.

'Guaranteed bandwidth' is the bandwidth reserved for a given traffic class regardless of the amount of traffic from other classes.

'Maximum bandwidth' is an upper limit on bandwidth allocated for a given traffic class, even if there is no other traffic from other classes.

These values are percentages of an endpoint's circuit speed.

Submit Reset Cancel

3. 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.

4. Click Submit to activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Defining Outbound QoS Endpoints

Each WX device can manage the outbound bandwidth for multiple remote endpoints. An endpoint is a WX device or a virtual endpoint that specifies one or more remote subnets that are not reachable through a WX device. After you run the Setup Wizard, you can:

- Specify the local WAN as dedicated or oversubscribed
- Enable congestion control for one or more endpoints.
- Enable or disable bandwidth management for any endpoint.
- Define virtual endpoints.
- Change the remote WAN circuit speeds.
- Specify LAN/WAN address or subnet pairs to be excluded from bandwidth management.



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

---

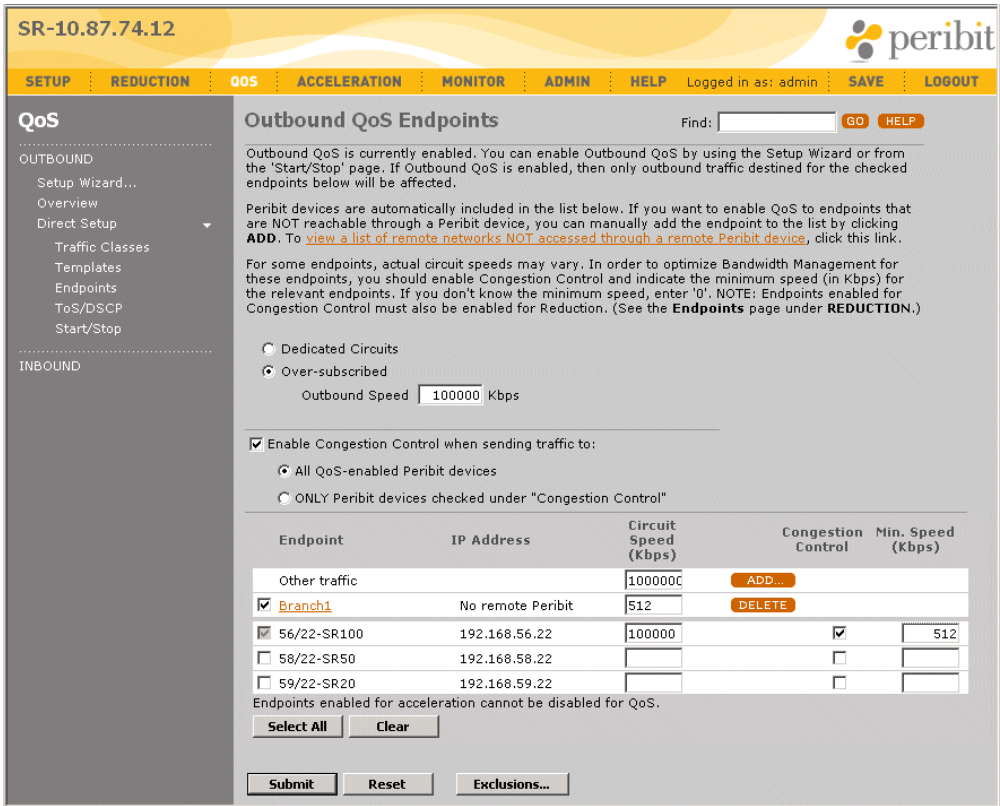
For oversubscribed WANs, you may have to decrease some speeds or guaranteed percentages before increasing others. If you use the Setup Wizard to change QoS settings, all percentages are adjusted automatically (refer to “Using the Outbound QoS Setup Wizard” on page 164).

To define the outbound QoS endpoints:

1. Click QOS in the menu frame, click Direct Setup in the left-hand navigation frame, and then click Endpoints.



Figure 97: Enabling Bandwidth Management by Endpoint



From the Outbound QoS Endpoints page, you can:

- Change the WAN mode:

- Dedicated Circuits** Indicates that the local outbound WAN speed equals or exceeds the sum of the WAN speeds for the remote endpoints whose bandwidths you want to manage (the default). In dedicated mode, traffic sent to non-WX endpoints (“Other traffic”) is unconstrained by QoS.
- If the WAN is dedicated, but you want “Other traffic” to be managed by QoS, you can define virtual endpoints (see Step 3) or select Oversubscribed and use the default outbound speed.
- Over-subscribed** Indicates that the local outbound WAN speed is less than the sum of the remote WAN speeds. Add up the speeds of all the WAN interfaces on the adjacent router, and enter the total in the **Outbound Speed** field (in Kbps). Be sure to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 156).
- Unlike the Setup Wizard, selecting oversubscribed mode here does not assign a template to the “Other traffic” endpoint. Unless you assign a template manually, “Other traffic” will have no guaranteed bandwidth (refer to “Defining Outbound QoS Settings by Endpoint” on page 171).

- Enable or disable outbound QoS for specific remote endpoints, as described in Step 2.

- Add or delete a virtual endpoint for remote subnets that do not have a WX device, as described in Step 3. To view the subnets associated with the current virtual endpoints, click **view a list of remote networks...**
- Change a virtual endpoint's name or subnets. Click the endpoint name, make the changes, and click **Submit**.
- Enable congestion control for one or more endpoints, as described in Step 4.
- Exclude specific address or subnet pairs from bandwidth management, as described in Step 5 through Step 7.

2. To enable bandwidth management for a remote endpoint:

- a. Select the check box next to the endpoint. To view the list of endpoints starting with a specific device name, enter the first part of the name (or the entire name) in the Find box at the top of the page, and click GO. To select all devices displayed on the page, click Select All. To deselect all displayed devices, click Clear.

Virtual endpoints have a DELETE button next to them, and are listed first by default. If Packet Flow Acceleration is enabled for an endpoint, outbound QoS cannot be disabled (the endpoint is greyed out).

- b. Enter the maximum remote WAN circuit speed (in Kbps) for each selected endpoint, and click Submit.



**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 156). Exceeding the actual WAN speed effectively shifts bandwidth management to the router, and may cause the router to drop traffic.

---

Note the following:

- For WX devices that support Multi-Path (refer to “Configuring Policy-Based Multi-Path” on page 115), a “\_Pri” or “\_Sec” is appended to the device name to indicate the primary or secondary path. You can enable QoS for one or both paths.
- The “Other traffic” endpoint is always enabled, and in oversubscribed mode is used to manage the bandwidth for all traffic that is not sent to one of the selected endpoints. The circuit speed for “Other traffic” defaults to the outbound speed.
- When you select 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 171.

- 3. To add a virtual endpoint, click ADD. Virtual endpoints let you manage the traffic to specific remote subnets that do not have a WX device (in dedicated or oversubscribed mode). By default, all such traffic is managed by the “Other traffic” endpoint, which is unconstrained by QoS in dedicated mode.

Figure 98: Adding Virtual Endpoints

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QoS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The main content area is titled 'Outbound QoS Endpoints > Add Endpoint'. It contains a text box for 'Name', a text box for 'Circuit Speed' with a unit dropdown set to 'kbps', and a text area for 'Subnets'. A 'Submit' button is at the bottom. A sidebar on the left shows the 'QoS' section expanded, with 'OUTBOUND' sub-sections like 'Setup Wizard...', 'Overview', 'Direct Setup', 'Traffic Classes', 'Templates', 'Endpoints', 'ToS/DSCP', and 'Start/Stop'.

Specify the following information, and click **Submit**. The maximum number of virtual endpoints (up to 320) depends on the device type (2 for the WX 15, 5 for the WX 20 and WXC 250, and 60 for the WXC 500).

Name	Enter the endpoint name (up to 20 characters).
Circuit Speed	Enter the maximum WAN circuit speed associated with this endpoint (in Kbps).
Subnets	Enter the IP addresses or subnets associated with this endpoint (one per line). The subnet format is: <IP address>/<subnet mask> Subnets specified here are ignored if they are also advertised by a WX device.

To delete a virtual endpoint, click DELETE next to the endpoint. Traffic to deleted virtual endpoints is managed by the “Other-traffic” endpoint.



**NOTE:** Traffic to a virtual endpoint has the lowest priority unless you specify QoS policies for the endpoint (refer to “Defining Outbound QoS Settings by Endpoint” on page 171).

- 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.

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 usually 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 WX 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). In oversubscribed mode, if you have substantial passthrough traffic for non-WX 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. To exclude one or more LAN/WAN pairs of addresses or subnets from bandwidth management, click Exclusions.

**Figure 99: Excluding Subnets or Hosts from Bandwidth Management**

Traffic that does not traverse the WAN should be excluded from outbound QoS. For example, if the WAN router has several LAN interfaces, traffic sent to those LANs should be excluded. To avoid managing traffic addressed to the router on the WAN side of the WX device, all LAN traffic sent to the device's local subnet is excluded by default.

6. Enter a local IP address or subnet in the **Between LAN side network** field, and enter a remote IP address or subnet 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.

7. Click Submit to activate the changes, or click Reset to discard them.
8. To retain your changes when the device is restarted, click SAVE in the menu frame.

### 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 WX “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. 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 5 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 5: 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	–

Name	DSCP	IP Precedence
AF31	26	–
AF32	28	–
AF33	30	–
AF41	34	–
AF42	36	–
AF43	38	–
EF	46	–

To set ToS/DSCP values by traffic class:

1. Click QoS in the menu frame, click Direct Setup in the left-hand navigation frame, and then click ToS/DSCP.

**Figure 100: Setting ToS/DSCP Values**

**Peribit 1**

SETUP REDUCTION QoS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**QoS**

OUTBOUND

- Setup Wizard...
- Overview
- Direct Setup
  - Traffic Classes
  - Templates
  - Endpoints
  - ToS/DSCP**
  - Start/Stop

INBOUND

**ToS/DSCP**

☐ Do not alter ToS/DSCP bits  
☒ Set IP Precedence bits for checked traffic classes (preserve incoming ToS byte if no classes set)  
☐ Set DSCP bits for checked traffic classes (preserve incoming DSCP if no classes set)

Options are disabled if IP Precedence or DSCP bits are being modified through Multi-Path or Router-based Route Load Balancing.

☒ Restore original ToS/DSCP bits after assembly

Traffic class	ToS/DSCP value
<input type="checkbox"/> Default	0

Select All Clear

Submit Reset

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 DSCP option is disabled if DSCP values are set by Multi-Path (refer to “Enabling Multi-Path and Defining Marking Methods” on page 117) or if ToS marking for router-based balancing is in use (refer to “configure route” on page 356).

The default selection, **Do not alter ToS/DSCP bits**, indicates that WX 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 “Defining Traffic Classes” on page 173.



**NOTE:** Changes to the ToS/DSCP values do not affect the outbound QoS reports. Also, these values are overridden by the ToS/DSCP settings defined for Multi-Path (refer to “Configuring Policy-Based Multi-Path” on page 115).

5. After reduced traffic from remote WX devices is assembled, the **Restore original ToS/DSCP bits after assembly** option resets the ToS/DSCP value to its original value (if the remote WX device changed it).
6. To retain your changes when the device is restarted, click SAVE in the menu frame.

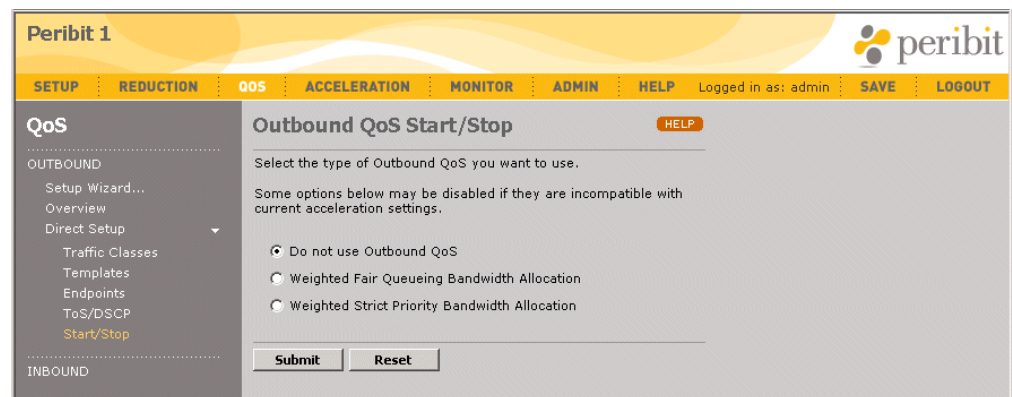
## Starting and Stopping Outbound QoS

You can start or stop outbound bandwidth management at any time, as well as change the prioritization method used to allocate the excess (unguaranteed) bandwidth among the contending applications. The selected prioritization model applies to all the managed endpoints.

To stop the outbound QoS service or change the prioritization:

1. Click QOS in the menu frame, click Direct Setup in the left-hand navigation frame, and then click Start/Stop.

**Figure 101: Starting and Stopping Outbound QoS**



2. To stop the outbound QoS service, click Do not use Outbound QoS. If Packet Flow Acceleration (PFA) is enabled, you cannot disable outbound QoS.
3. To restart the service or change the prioritization method used for each endpoint, select one of the following.
  - **Weighted Fair Queueing Bandwidth Allocation.** Queues are created for each traffic class, and the excess 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 excess bandwidth is allocated by processing the queues based only on priority.
4. Click Submit to activate the changes, or click Reset to discard them.
  5. To retain your changes when the device is restarted, click SAVE in the menu frame.

### ***Processing Queues Based on Incoming ToS/DSCP Values***

Optionally, queues can be processed based on the incoming ToS/DSCP values, as follows:

1. Create application definitions based solely on the ToS or DSCP settings. For example, an application named "DSCP-5" could be defined with a DSCP value of 5 (no other settings) that would apply to all traffic with a matching DSCP value.
2. Define a traffic class for each application definition (maximum of 15 classes available). For example, assign the "DSCP-5" application to a traffic class named "DSCP-5-class".
3. Specify the QoS policies for each traffic class, as normal, and enable QoS with Weighted Fair Queuing or Weighted Strict Priority.

Traffic coming in to the WX device from the LAN will now be queued for processing based on the ToS/DSCP values.

## **Configuring Inbound QoS Policies**

Inbound bandwidth management lets you specify 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. Thus, inbound QoS does not apply to off-path WX devices (which use only the Local interface). Also, on devices configured for tunnel switching, 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 6: 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 Traffic report to help create the list of subnets (refer to “Traffic Statistics” on page 256).
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 102, 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 102: Configuring Inbound Bandwidth Management**

To configure the inbound QoS service:

1. Click QOS in the menu frame, and then click INBOUND in the left-hand navigation frame.

**Figure 103: Configuring Maximum Inbound QoS Bandwidths**

**Peribit 1**

SETUP REDUCTION **QoS** ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**QoS**

OUTBOUND  
Setup Wizard...  
Overview  
Direct Setup

INBOUND

**Inbound QoS**

If 'Enable Inbound QoS' is checked, traffic from the following four predefined traffic classes will be limited to the specified maximum bandwidths.

☐ Enable Inbound QoS

Inbound Speed  Kbps

Traffic Class	Maximum Bandwidth	Description
Reduced	<input type="text"/> 100 %	Any traffic that has been reduced by a Peribit device.
Intranet	<input type="text"/> 100 %	TCP traffic originating from the corporate network that has NOT been reduced.
TCP	<input type="text"/> 100 %	TCP traffic NOT originating from the corporate network.
Default	<input type="text"/> 100 %	All other protocols (e.g. UDP, streaming)

Submit Reset

- To start the inbound QoS service, click **Enable Inbound QoS**.
- Add up the speeds of all the WAN interfaces on the adjacent router that conduct traffic to the WX device, and enter the value (in Kbps) in the **Inbound Speed** field.
- Enter the maximum bandwidth of each traffic class as a percentage of the inbound speed.
- Click Submit to activate the changes, or click Reset to discard them.
- Click Intranet to specify the remote subnets whose traffic belongs to the Intranet class.

**Figure 104: Configuring Subnets for the Inbound QoS Intranet Class**

**Peribit 1**

SETUP REDUCTION **QoS** ACCELERATION MONITOR ADMIN HELP Logged in as: admin LOGOUT

**QoS**

OUTBOUND  
Setup Wizard...  
Overview  
Direct Setup

INBOUND

**Inbound QoS > Intranet**

Traffic that originates from within the corporate network (and has NOT been reduced) is categorized as belonging to the 'Intranet TCP' class.

Enter the IP address/subnet mask for all subnets belonging to the corporate network, one per line.  
For example: 123.123.123.0/255.255.255.0

Submit Reset Cancel

7. In the list box, enter the remote subnets (one per line) whose traffic belongs to the Intranet traffic class. The subnet format is:

<IP address>/<subnet mask>

8. Click Submit to activate the changes, or click Reset to discard them.
9. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Summary of Key Terms

---

The following terms and concepts are key to understanding outbound QoS:

- **Dedicated WAN.** The local outbound speed equals or exceeds the sum of the remote endpoint circuit speeds.
- **Endpoint circuit speed.** Maximum WAN circuit speed associated with a remote WX device or a “virtual” endpoint (no WX device).
- **Excess bandwidth.** Difference between the total available bandwidth and the guaranteed bandwidth currently being used.
- **Guaranteed bandwidth.** Amount of bandwidth guaranteed to a given traffic class.
- **Maximum bandwidth.** Maximum amount of bandwidth a traffic class can consume.
- **Outbound speed.** Sum of the WAN circuit speeds on the adjacent router that conduct traffic to the WX device.
- **Oversubscribed WAN.** The local outbound speed is less than the sum of the remote endpoint circuit speeds.
- **QoS template.** Specifies a priority, guaranteed bandwidth, and maximum bandwidth for each traffic class. You can apply a different QoS template to the traffic sent to each remote WX device.
- **Queuing model.** Scheduling algorithm that allocates bandwidth to each traffic class. The queuing models are Weighted Fair Queuing (WFQ) and Weighted Strict Priority (WSP).
- **Setup Wizard.** Prompts you to specify the QoS settings for your network, including WAN link speeds, traffic classes, the priority and bandwidths for each class, and the queuing model. For an oversubscribed WAN, the Setup Wizard automatically adjusts the guaranteed bandwidths to ensure fair traffic delivery to each endpoint.
- **Traffic class.** A group of one or more applications.



## Chapter 7

# Accelerating WAN Traffic

The following sections describe how to configure traffic acceleration:

- “Packet Flow Acceleration” in the next section
- “Application Flow Acceleration” on page 200

## Packet Flow Acceleration

---

The following sections describe how to configure Packet Flow Acceleration:

- “Overview of Packet Flow Acceleration” in the next section
- “Requirements for Using Packet Flow Acceleration” on page 193
- “Enabling Packet Flow Acceleration by Endpoint” on page 194
- “Enabling Active Flow Pipelining by Application” on page 197
- “Enabling Fast Connection Setup by Application” on page 198

## 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) provides four methods to improve the throughput of reduced TCP application flows across high-speed, high-latency WAN links. For WX devices that support Multi-Path, you can enable PFA for the primary and/or secondary paths.

The following topics describe each PFA method:

- “Active Flow Pipelining” in the next section
- “Forward Error Correction” on page 192
- “Fast Connection Setup” on page 192

## 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.



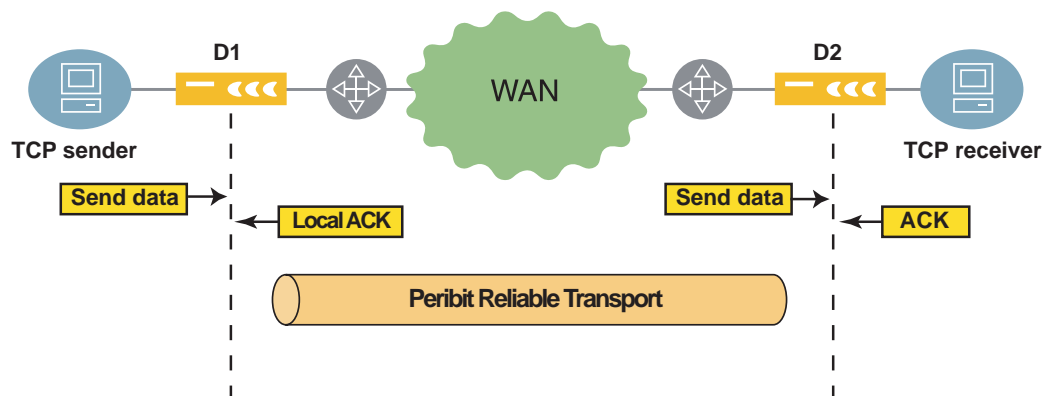
**NOTE:** AFP is required to use Network Sequence Caching (NSC) on WXC devices, or to accelerate Microsoft CIFS, Microsoft Exchange, and HTTP traffic using Application Flow Acceleration.

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 WX device, between the two WX devices, and between the receiving WX device and the destination.

Since the WX devices acknowledge all transmissions locally, more data can be put “in flight” at once. The WX device returns ACKs 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 105: 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 traffic 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 WX 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 WX device. Each cluster can have four devices.

In the following example, if D3 receives a SYN packet from D1, the SYN and its source are advertised to D4. If D4 receives the SYN ACK, it can establish the flow with D1.

**Figure 106: AFP Clusters for Asymmetric Routing Support**



Note the following about asymmetric routing support:

- Load balancing on the router or switch must be flow- or destination based (not packet-based).
- If you have a cluster on both sides of the WAN, reduction tunnels must be enabled in both directions between all the WX devices in the two clusters. Note that the routers need not be fully meshed. For example, the physical path between D1 and D4 can be R1 to R3 to R4.
- If a device is in a cluster, it can accelerate traffic only to remote devices that are running WXOS 5.1 or later.
- If Multi-Path is enabled on one peer, it must be enabled for all devices in the cluster. Also, traffic is accelerated only if the same path is used in both directions (primary or secondary).
- Asymmetric routing support takes precedence over preferred assemblers and tunnel load balancing settings (if any) defined on the WX device.

### AFP Statistics

Figure 107. shows an example of the statistics provided for Active Flow Pipelining. The acceleration factor is the actual average throughput divided by the estimated throughput without acceleration. Note that performance improvements will be more noticeable to users as the accelerated session count and traffic load increases.

**Figure 107: Sample Active Flow Pipelining Statistics**

Application	Total TCP Sessions (count)	Accelerated Sessions (count)	Traffic (MB)	Average Session Throughput (Mbps)		Acceleration Factor
				Actual	w/o Accel.*	
FTP	56	56	369	123	61	2.1 X
CIFS	12	12	1235	76	24	3.1 X
HTTP	78	78	698	28	7	4.2 X
Others	17495	17495	76	8	5	1.7 X

On a given path between two WX devices, AFP may also benefit from Forward Error Correction, but AFP cannot be used simultaneously with Fast Connection Setup.

### Forward Error Correction

Forward Error Correction (FEC) enables the sending WX 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 WX 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 extra overhead required to process recovery packets.

### Fast Connection Setup

With Fast Connection Setup (FCS), the sending WX 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 108: 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 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.

Figure 109. shows an example of the FCS statistics. 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.

In Figure 109, the FTP gains apply to a small number of sessions that probably affect only the traffic on the control port.

**Figure 109: Sample Fast Connection Setup Statistics**

Application	Total TCP Sessions (count)	Short Sessions*		Average Short Session Time (msec)		Average Short Session Acceleration (percent)				
		(count)	(percent)	with Accel.	w/o Accel.					
HTTP	329	121	36.8%	772.30	1020.51	24.3%				
NetBios	714	68	9.5%	873.91	1088.90	19.7%				

## Requirements for Using Packet Flow Acceleration

To use PFA to accelerate application traffic between two WX devices, the following conditions must be met:

- The applications must be reduced (refer to “Reducing Applications” on page 135).
- A reduction tunnel must exist in both directions between the WX devices (refer to “Configuring Endpoints for Reduction Tunnels” on page 129).

- Outbound QoS must be enabled, and the WAN circuit speed must be specified for each remote WX device for which you want to accelerate traffic (refer to “Using Outbound QoS to Enhance Performance” on page 153).

Note that if the circuit speeds are specified incorrectly, too much data may be sent to the router, and the acceleration reports may show performance gains that cannot be realized due to router congestion.

- To use Active Flow Pipelining, you must:
  - Enable AFP on the Community Topology page (refer to “Setting Community Feature and Topology Parameters” on page 99).
  - Enable asymmetric routing support (clustering) if the outbound and return traffic does not always traverse the same two WX devices (refer to “Asymmetric Routing for AFP” on page 191 and “configure acceleration” on page 302). For AFP to accelerate a traffic flow, all the traffic must traverse the same two WX devices in both directions.

Check the sent and received packet counts in the Top Flows traffic report to verify that traffic is traversing the same devices in both directions, (refer to “Traffic Statistics” on page 256).



**NOTE:** PFA is most effective in networks with high-speed connections and high latency, and/or very high reduction rates. 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.

---

### **Enabling Packet Flow Acceleration by Endpoint**

You can enable each method of Packet Flow Acceleration for all remote WX 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 WX devices (refer to “Configuring Endpoints for Reduction Tunnels” on page 129).
- Enable reduction for the applications you want to accelerate (refer to “Reducing Applications” on page 135).
- Enable outbound QoS using Weighted Fair Queuing or Weighted Strict Priority, and specify the WAN circuit speed for the remote endpoint (refer to “Using Outbound QoS to Enhance Performance” on page 153).

If you enable Active Flow Pipelining or Fast Connection Setup, you must select the applications that each method is applied to (refer to “Enabling Fast Connection Setup by Application” on page 198 and “Enabling Active Flow Pipelining by Application” on page 197).

To enable Packet Flow Acceleration by endpoint:

1. Verify that reduction tunnels exist in both directions between the WX devices that you want to support Packet Flow Acceleration (refer to “Configuring Endpoints for Reduction Tunnels” on page 129).
2. Click ACCELERATION in the menu frame.

**Figure 110: Enabling Packet Flow Acceleration**

**Peribit 1**

SETUP REDUCTION QoS **ACCELERATION** MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Acceleration**

PACKET FLOW ACCELERATION

Overview

Active Flow Pipelining

Fast Connection Setup

APPLICATION FLOW ACCELERATION

CIFS

Exchange

**Acceleration Overview**

Find:  GO 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	Recovery Packets	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"/>	1	9
<input checked="" type="checkbox"/> SR-192.168.72.10	192.168.72.10	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input checked="" type="checkbox"/> SR-192.168.73.11	192.168.73.11	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input checked="" type="checkbox"/> SR-192.168.74.11	192.168.74.11	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input type="checkbox"/> SR-192.168.75.10	192.168.75.10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input type="checkbox"/> SR-192.168.76.10	192.168.76.10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input type="checkbox"/> SR-192.168.78.10	192.168.78.10	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9

Note: QoS must be enable endpoint before it can be accelerated

\* Should only be used for connections with applicato generate many very short TCP connections (e.g. HTT across high latency links.

† Should only be used for connections that are subje high loss and do not have enabled on the satellite m CSU/DSU.

(Note that QoS must be enabled on an endpoint before it can be accelerated.)

(Note that NSM must be disabled on an endpoint before AFP can be turned off.)

Select All Clear ☒ Show Advanced Settings

Submit Reset

3. 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. To enable Active Flow Pipelining, you must enable AFP on the Community Topology page (refer to “Setting Community Feature and Topology Parameters” on page 99).
4. Select one of the following options:
  - **Accelerate all QoS enabled endpoints using default settings.** Traffic is accelerated to all remote WX 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 WX devices, and different PFA methods can be used for each endpoint. Click the check box next to the appropriate devices.

To view the list of endpoints starting with a specific device name, enter the first part of the name (or the entire name) in the Find box at the top of the page, and click GO. To select all devices displayed on the page, click Select All. To deselect all displayed devices, click Clear. An endpoint is greyed out if no reduction tunnel exists or outbound QoS is not configured for the endpoint.

For WX devices that support Multi-Path, a “\_Pri” or “\_Sec” is appended to the device name to indicate the primary or secondary path. You can enable PFA for one or both paths. To configure Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 115.

5. 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 Caching and Application Flow Acceleration.</p> <p><b>NOTE:</b> In some cases, you may need to do one or more of the following (refer to “configure acceleration” on page 302):</p> <ul style="list-style-type: none"> <li>Adjust the buffer size for optimum performance.</li> <li>Increase the number of lost heartbeat packets allowed on high-loss links (reduction may stop when consecutive heartbeat packets are lost).</li> <li>Enable clustering if the outbound and return traffic does not always traverse the same two WX devices.</li> <li>If tunnel load balancing is enabled, verify that it is “Flow based” or “Per-destination” (refer to “Configuring Tunnel Load Balancing Policies” on page 139)</li> <li>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” command “Fast reduction tunnels” on page 348).</li> </ul>
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 <b>Show Advanced Settings</b> 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 7.

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 7: 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 %

- Click Submit to activate the changes, or click Reset to discard them.
- To retain your changes when the device is restarted, click SAVE in the menu frame.

You can now enable PFA for specific applications, as described in the following sections.

### **Enabling Active Flow Pipelining by Application**

After you enable Active Flow Pipelining, as described in “Enabling Packet Flow Acceleration by Endpoint” on page 194, you can select the applications whose outgoing traffic you want to accelerate. Active Flow Pipelining is intended for applications that transfer large amounts of data, such as FTP and CIFS, over high-latency links, such as satellite connections, and long-haul high-bandwidth links, such as E3 and T3.

To enable Active Flow Pipelining for one or more applications:

- Click ACCELERATION in the menu frame, and then click Active Flow Pipelining in the left-hand navigation frame.

**Figure 111: Enabling Active Flow Pipelining by Application**

**Peribit 1**

SETUP REDUCTION QOS **ACCELERATION** MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Acceleration**

PACKET FLOW ACCELERATION

- Overview
- Active Flow Pipelining**
- Fast Connection Setup

APPLICATION FLOW ACCELERATION

- CIFS
- Exchange

**Active Flow Pipelining** HELP

When Active Flow Pipelining is enabled, checked applications will be accelerated, provided they have also been checked for Reduction (See the **Application Filter** page under REDUCTION.)

Checked application traffic will be included in Acceleration report statistics only if the applications are also checked for Monitoring. (See the **Monitoring** page under SETUP > APPLICATIONS.)

Application

<input checked="" type="checkbox"/>	AOL
<input checked="" type="checkbox"/>	CIFS
<input checked="" type="checkbox"/>	Clearcase
<input checked="" type="checkbox"/>	CVS
<input checked="" type="checkbox"/>	DNS
<input checked="" type="checkbox"/>	epmap
<input checked="" type="checkbox"/>	Exchange
<input checked="" type="checkbox"/>	Filenet
<input checked="" type="checkbox"/>	FTP
<input checked="" type="checkbox"/>	Groupwise
<input checked="" type="checkbox"/>	Hostname Resolution
<input checked="" type="checkbox"/>	HTTP
<input checked="" type="checkbox"/>	HTTPS
<input checked="" type="checkbox"/>	ICA
<input checked="" type="checkbox"/>	Kerberos
<input checked="" type="checkbox"/>	LDAP
<input checked="" type="checkbox"/>	Undefined Applications

Select All Clear

Submit Reset

2. Select the check box next to each application that you want to accelerate using Active Flow Pipelining, or click Select All. The selected applications are accelerated only if they are also being reduced (refer to “Reducing Applications” on page 135).



**NOTE:** Active Flow Pipelining must be enabled on both the sending and receiving WX devices.

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

### Enabling Fast Connection Setup by Application

After you enable Fast Connection Setup, as described in “Enabling Packet Flow Acceleration by Endpoint” on page 194, 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 and NetBios.

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 are accelerated.

To enable Fast Connection Setup for one or more applications:

1. Click ACCELERATION in the menu frame, and then click Fast Connection Setup in the left-hand navigation frame.

**Figure 112: Enabling Fast Connection Setup by Application**

**Peribit 1**

SETUP REDUCTION QOS **ACCELERATION** MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Acceleration**

PACKET FLOW ACCELERATION

- Overview
- Active Flow Pipelining
- Fast Connection Setup**

APPLICATION FLOW ACCELERATION

- CIFS
- Exchange

**Fast Connection Setup** HELP

When Fast Connection Setup is enabled, checked applications will be accelerated, provided they have also been checked for Reduction (See the **Application Filter** page under **REDUCTION**.)

Checked application traffic will be included in Acceleration report statistics only if the applications are also checked for Monitoring. (See the **Monitoring** page under **SETUP > APPLICATIONS**.)

**Application**

- ☐ AOL
- ☐ CIFS
- ☐ Clearcase
- ☐ CVS
- ☐ DNS
- ☐ epmap
- ☐ Exchange
- ☐ Filenet
- ☐ FTP
- ☐ Groupwise
- ☐ Hostname Resolution
- ☒ HTTP
- ☐ HTTPS
- ☐ ICA
- ☐ Kerberos
- ☐ Telnet

Clear

Submit Reset

2. Select the check box next to each application that you want to accelerate using Fast Connection Setup. The selected applications are accelerated only if they are also being reduced (refer to “Reducing Applications” on page 135).



**NOTE:** To accelerate application traffic in both directions between two WX devices, you must enable Fast Connection Setup for the application on both devices.

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Application Flow Acceleration

---

The following topics describe Application Flow Acceleration for Microsoft CIFS, Microsoft Exchange, and HTTP traffic:

- Overview of Application Flow Acceleration on page 200
- Enabling Microsoft CIFS Acceleration on page 204
- Enabling Microsoft Exchange Acceleration on page 207
- Enabling HTTP Acceleration on page 209



**NOTE:** Application Flow Acceleration must be enabled on the Features/ Topology page (refer to “Setting Community Feature and Topology Parameters” on page 99).

---

### Overview of Application Flow Acceleration

Though technologies such as compression (MSR and NSC) 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 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 optimized: 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 WX 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 WX devices closest to the clients.



**NOTE:** Application Flow Acceleration and tunnel switching cannot be enabled on the same WX device.

---




### 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 113), the client requests one block of data at a time. The WX device closest to the client (D1) requests the next  $N$  blocks. The WX 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 113: Microsoft CIFS/Exchange Read Operations**



In write operations (Figure 114), the client writes one block at a time. The WX device closest to the client (D1) acknowledges each block locally, and discards the acknowledgements from the server.

**Figure 114: 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 204). 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). WX devices can cache only HTTP response headers, but WXC 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 WXC locally responds to both GET and GET IF-MODIFIED-SINCE requests for the static objects in the WX 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 WXC devices, only objects that are considered “stale” by the cache are pre-fetched.

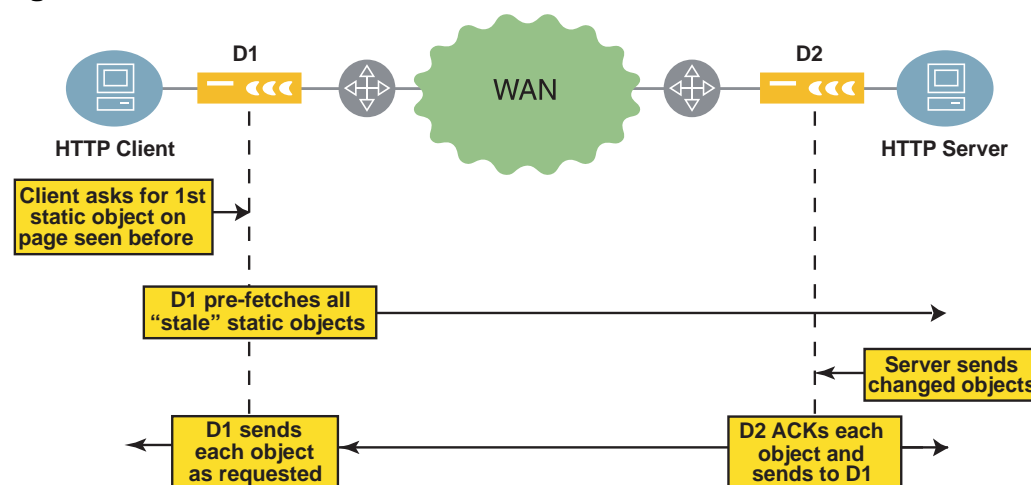
In HTTP cache header-only mode (Figure 115), the client sends HTTP GET IF-MODIFIED-SINCE queries before reloading an object from the browser cache. Based on its own caching timer, the WX device closest to the client (D1) indicates the object has not changed or forwards the query to the server. Even if the query is not forwarded, D1 sends its own query to verify the object's last-modified date.

**Figure 115: HTTP Caching—Header-Only Mode**

In HTTP cache header-and-body mode (Figure 116), the client sends HTTP GET requests for static objects on a page that has been visited before. The WXC 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 116: HTTP Caching—Header-and-Body Mode (WXC Devices)**

If pre-fetch is enabled (Figure 117), 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 WX device (D1) requests all the static objects that are considered stale. The objects returned by the server are acknowledged locally by D2.

**Figure 117: HTTP Pre-Fetch**

To view the current cache usage, refer to "configure acceleration" on page 302.



**NOTE:** HTTP traffic is not accelerated if a proxy server exists between the server-side WX device and the actual HTTP server. However, if the proxy server is between the Web client and the client-side WX device, HTTP traffic will be accelerated.

## 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 WX 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 92)
- Active Flow Pipelining must be enabled on both the client- and server-side WX devices (refer to "Enabling Active Flow Pipelining by Application" on page 197).
- Application Flow Acceleration must be enabled on the client-side WX device (select **All features** on the Features/Topology page, as described in "Setting Community Feature and Topology Parameters" on page 99). On the server-side WX 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. Click SETUP in the menu frame, click APPLICATIONS and Definitions in the left-hand navigation frame, 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.

**Figure 118: Adding New CIFS Application Definitions**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC  
Addresses  
Interfaces  
Time  
License Key  
SNMP  
Syslog Server  
Local Routes  
Registration Server

AAA

APPLICATIONS  
Overview  
Definitions  
Traffic Classes  
Monitoring

IPSEC

ADVANCED

**Application Definitions > New**

Application Name: CIFS\_Server\_1  
Application Type: CIFS

Application traffic will be identified using the following rules

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
10.10.20.25	139,445			Any	Advanced CLEAR
		10.10.20.25	139,445	Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR

Enter IP Address or subnet. Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0  
Use commas to enter multiple ports. Use hyphen (-) to specify a range. Example: 25, 27, 125-135  
To match any value, leave the field blank. Do not use asterisk (\*).

Submit Cancel

- 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, click ACCELERATION in the menu frame, and then click CIFS in the left-hand navigation frame.

**Figure 119: 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.

The **Disable SMB signing** check box allows CIFS transactions to be accelerated for servers that have SMB signing enabled, but not required. To accelerate CIFS transactions for Windows 2000 and Windows 2003 domain controllers that require SMB signing, you must change the Group Policy on the domain controller (see Step 6).

4. Click Submit to activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.
6. Verify that SMB signing is disabled on Windows 2000 and Windows 2003 domain controllers that are also file servers. For more information, refer to the Microsoft Web site:  
<http://support.microsoft.com/default.aspx?scid=kb;en-us;887429>.

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.
7. 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 WXC devices using NSC disk-based compression provide substantial benefits for such traffic without acceleration. Also, Exchange 2003/Outlook 2003 use compression by default. Since WX reduction is more effective, Microsoft Exchange compression should be disabled (refer to <http://support.microsoft.com/?kbid=825371>).

Enable Exchange acceleration on the WX 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 92)
- Active Flow Pipelining must be enabled on both the client- and server-side WX devices (refer to “Enabling Active Flow Pipelining by Application” on page 197).
- Application Flow Acceleration must be enabled on the client-side WX device (select **All features** on the Features/Topology page, as described in “Setting Community Feature and Topology Parameters” on page 99). On the server-side WX 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. Click SETUP in the menu frame, click APPLICATIONS and Definitions in the left-hand navigation frame, 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.

**Figure 120: Adding New Exchange Application Definitions**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

- Addresses
- Interfaces
- Time
- License Key
- SNMP
- Syslog Server
- Local Routes
- Registration Server

AAA

APPLICATIONS

- Overview
- Definitions
- Traffic Classes
- Monitoring

IPSEC

ADVANCED

**Application Definitions > New**

Application Name: Exchange\_Server\_1

Application Type: Exchange

Application traffic will be identified using the following rules

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
10.10.20.45	135			Any	Advanced CLEAR
		10.10.20.45	135	Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR

Enter IP Address or subnet.  
Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0

Use commas to enter multiple ports.  
Use hyphen (-) to specify a range.  
Example: 25, 27, 125-135

To match any value, leave the field blank. Do not use asterisk (\*).

Submit Cancel

- 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, click **ACCELERATION** in the menu frame, and then click **Exchange** in the left-hand navigation frame.

**Figure 121: 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 activate the changes, or click **Reset** to discard them.
5. To retain your changes when the device is restarted, click **SAVE** in the menu frame.

## 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 WX 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 92)
- Active Flow Pipelining must be enabled on both the client- and server-side WX devices (refer to “Enabling Active Flow Pipelining by Application” on page 197).
- Application Flow Acceleration must be enabled on the client-side WX device (select **All features** on the Features/Topology page, as described in “Setting Community Feature and Topology Parameters” on page 99). On the server-side WX 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 WX device and the actual HTTP server. However, if the proxy server is between the Web client and the client-side WX 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. Click **SETUP** in the menu frame, click **APPLICATIONS** and **Definitions** in the left-hand navigation frame, 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**.

**Figure 122: Adding New HTTP Application Definitions**

Peribit 1

peribit

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

- Addresses
- Interfaces
- Time
- License Key
- SNMP
- Syslog Server
- Local Routes
- Registration Server

AAA

APPLICATIONS

- Overview
- Definitions
- Traffic Classes
- Monitoring

IPSEC

ADVANCED

**Application Definitions > New**

Application Name: HTTP\_Server\_1

Application Type: HTTP

Application traffic will be identified using the following rules

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
10.10.20.55	80	10.10.20.55	80	Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR
				Any	Advanced CLEAR

Enter IP Address or subnet.  
Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0

Use commas to enter multiple ports.  
Use hyphen (-) to specify a range.  
Example: 25, 27, 125-135

To match any value, leave the field blank. Do not use asterisk (\*).

Submit Cancel

- 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, click **ACCELERATION** in the menu frame, and then click **HTTP** in the left-hand navigation frame.

**Figure 123: Enabling HTTP Acceleration**

**Peribit 1**

peribit

SETUP REDUCTION QOS **ACCELERATION** MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Acceleration**

PACKET FLOW ACCELERATION

- Overview
- Active Flow Pipelining
- Fast Connection Setup

APPLICATION FLOW ACCELERATION

- CIFS
- Exchange
- HTTP

**HTTP Acceleration** [HELP](#)

Active Flow Pipelining must be enabled before HTTP acceleration can be used. It is only necessary to enable HTTP acceleration on the Peribit device closest to the client. HTTP acceleration does not work for HTTPS traffic.

☐ Disable HTTP Acceleration  
☒ Enable HTTP Acceleration for checked applications

Application

☐ HTTP  
☒ HTTP\_Server\_1

Select All Clear

Submit Reset

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 activate the changes, or click Reset to discard them.
5. To retain your changes when the device is restarted, click SAVE in the menu frame.

On WXC devices, static objects are cached by default (header-and-body mode). To change the cache setting, refer to “configure acceleration” on page 302.



## Chapter 8

# Configuring IP Security (IPSec)

The following sections describe how to configure IP security (IPSec) to authenticate and encrypt traffic between any pair of WX devices:

- “Overview of IPSec” in the next section
- “Procedure for Configuring IPSec Policies” on page 215
- “Using the IPSec Setup Wizard” on page 215
- “Defining IPSec Settings by Endpoint” on page 219
- “Defining IPSec Templates” on page 222
- “Defining the Default IPSec Policy” on page 224

### Overview of IPSec

---

IPSec can be used to authenticate and encrypt traffic between any pair of WX devices (endpoints) 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 WX 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 WX devices (the WX 15 and WX 20 are limited to 2 and 5 devices, respectively).

### Default IPSec Policy

When two WX 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 WX 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 WX device fails.



**NOTE:** If an inline WX device fails, all traffic is passed through without encryption. To block all traffic during a hardware failure, use a crossover cable (rather than a straight-through cable) to connect the WX device to the WAN router. This works only if Ethernet auto-MDI negotiation is disabled on the router.

## IPSec Implementation Details

The WX implementation of 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 has not been approved, but 3DES is both 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 Peer1 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.

Supporting IPSec allows WX devices 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 the WX implementation because they cannot decompress traffic. However, uncompressed WX 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 configure a pair of WX devices to support IPSec, do the following on both devices:

1. Verify that reduction and assembly are enabled (refer to “Configuring Endpoints for Reduction Tunnels” on page 129). Specific endpoints need not be enabled for reduction.
2. Run the Setup Wizard to create and apply the IPSec Wizard template to selected WX devices (endpoints), as described in “Using the IPSec Setup Wizard” on page 215.



**NOTE:** Each time you run the Setup Wizard the existing Wizard template is overwritten.

---

3. To change the template settings, run the Setup Wizard again (overwriting the template) or make the changes manually. The following changes must be made manually:
  - Change a template for a specific endpoint, or enable encryption for the WX management traffic (refer to “Defining IPSec Settings by Endpoint” on page 219).
  - Add new templates, or change a template name or key lifetimes (refer to “Defining IPSec Templates” on page 222).

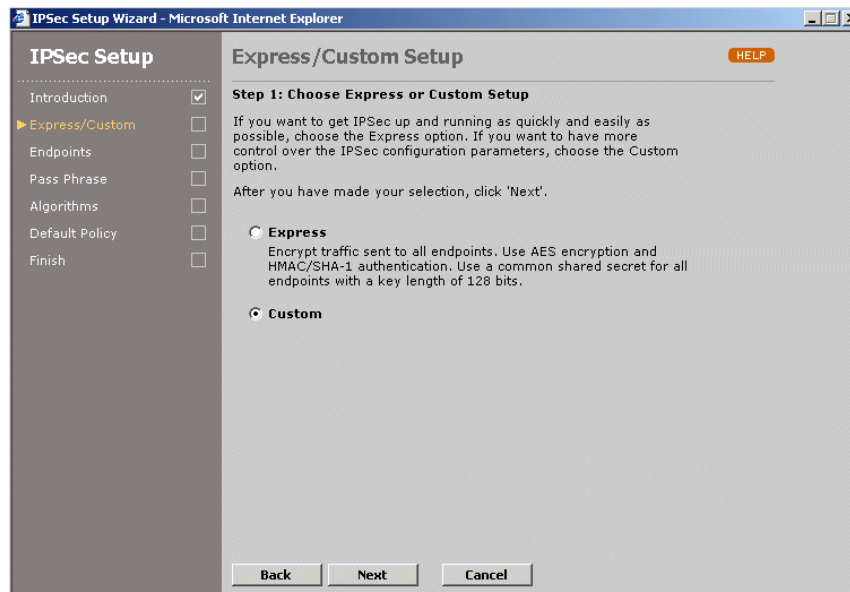
## Using the IPSec Setup Wizard

---

Always use the Setup Wizard the first time you define the IPSec policies. The Setup Wizard creates a template called Wizard and applies it to the selected endpoints. Each time you run the Setup Wizard, the Wizard template is overwritten. To define other templates, refer to “Defining IPSec Templates” on page 222.

To run the IPSec Setup Wizard:

1. In the Setup page, click IPSEC in the left-hand navigation frame, and then click Setup Wizard.
2. Click Enable IPSec and click Next.



3. Select one of the following, and click Next.

Express	<p>Applies the default Wizard template to all endpoints, and prompts you to enter a single pass phrase that applies to all endpoints (up to 64 characters). Note that assigning the same pass phrase to all endpoints is a poor security practice, and is not recommended.</p> <p>The default template uses the following algorithms:</p> <ul style="list-style-type: none"> <li>■ <b>Encryption.</b> Advanced Encryption Standard with a 128-bit key (AES-128)</li> <li>■ <b>Authentication.</b> Secure Hash Algorithm (HMAC/SHA-1)</li> </ul> <p>If you select this option, enter the pass phrase for all endpoints on the next page, and go to Step 9.</p>
Custom	<p>Allows you to select specific endpoints that support IPsec, specify a separate pass phrase for each endpoint, and select the template algorithms.</p>

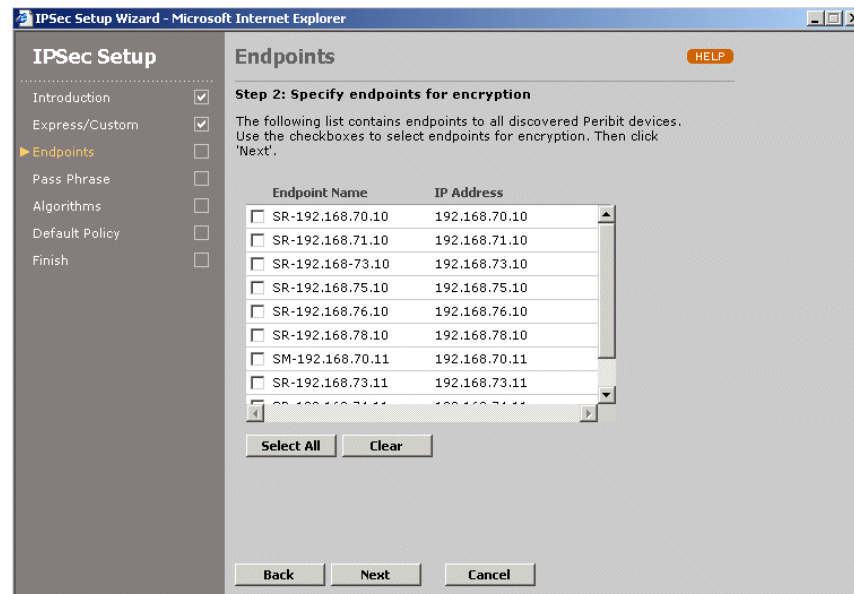
4. If you select the Custom setup, select one of the following, and click Next.

All endpoints	Enables IPsec for all remote endpoints.
Selected endpoints	Allows you to select specific endpoints that support IPsec.

5. To enable IPsec for specific endpoints, select the check box next to the appropriate remote devices, or click Select All, and then click Next.

For WX devices that support Multi-Path, a “\_Pri” or “\_Sec” is appended to the device name to indicate the primary or secondary path. You can enable IPsec for one or both paths. To configure Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 115.






6. Select one of the following, and click Next.

- |                         |   |
|-------------------------|---|
| Common Pass Phrase      | Prompts you to specify one pass phrase for all endpoints. (This is a poor security practice, and is not recommended.) |
| Individual Pass Phrases | Prompts you to specify a separate pass phrase for each endpoint.  |

The pass phrase is used to generate a preshared key of the appropriate length. The pass phrase can be from 4 to 64 four characters, but 8 characters is the recommended minimum.


7. Enter and verify a pass phrase for each endpoint or all endpoints, and click Next. The same pass phrase must be specified on the remote devices.
8. Select the encryption and authentication algorithms, and click Next.

- |                          |  |
|--------------------------|--|
| Encryption Algorithm     | <p>Select the algorithm used to encrypt outbound traffic:</p> <ul style="list-style-type: none"> <li>■ <b>Any.</b> The algorithm selected for the peer endpoint is used. If both endpoints specify Any, AES-128 is used.</li> <li>■ <b>AES-128.</b> Advanced Encryption Standard with a 128-bit key.</li> <li>■ <b>AES-192.</b> AES with a 192-bit key.</li> <li>■ <b>AES-256.</b> AES with a 256-bit key.</li> <li>■ <b>3DES.</b> Triple Digital Encryption Standard with a 168-bit key.</li> </ul> |
| Authentication Algorithm | <p>Select the algorithm used to authenticate outbound traffic:</p> <ul style="list-style-type: none"> <li>■ <b>Any.</b> The algorithm selected for the peer endpoint is used. If both endpoints specify Any, HMAC/SHA-1 is used.</li> <li>■ <b>HMAC/SHA-1.</b> Secure Hash Algorithm.</li> <li>■ <b>HMAC/MD5.</b> Message Digest 5.</li> </ul>   |



9. The default IPsec policy is applied to traffic sent to unadvertised subnets (no WX device) and to subnets advertised by devices where encryption is enabled, but the key negotiation has failed. By default, all such traffic is unencrypted.

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 WX device fails.



Specify destination addresses and subnets where encryption is required, and click Next.

- |                     |  |
|---------------------|--|
| Encryption Required | Enter destination addresses or subnets (one per line) for which traffic must be dropped and logged. The subnet format is:<br><b>&lt;IP address&gt;/&lt;subnet mask&gt;</b>   |
| Encryption Optional | Enter destination addresses or subnets (one per line) for which traffic can be sent unencrypted.<br><br>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, the traffic is sent unencrypted. |

10. Click Next, click Submit, and then click Close.
11. Under IPSEC in the left-hand navigation frame, click Overview to refresh the IPSec Overview page, which now shows the template name (Wizard) assigned to each of the selected endpoints.
12. To retain your changes when the device is restarted, click SAVE in the menu frame.

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

## Defining IPSec Settings by Endpoint

After you run the Setup Wizard to create the initial IPSec settings, 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. You can also view the status of each secure connection. To add or change IPSec templates, refer to “Defining IPSec Templates” on page 222.

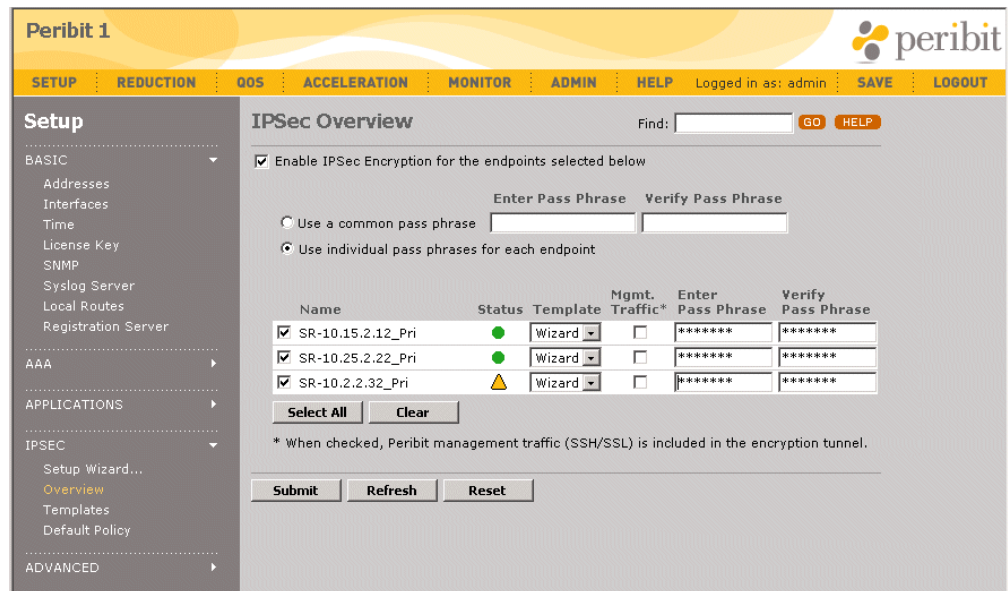


**NOTE:** When the Status column indicates that IPSec is operating normally with a remote WX device, it is highly recommended that you enable encryption of management traffic for that device. Also, remember to save the configuration so that encryption is not lost when the device is restarted.

To view or change the IPSec settings by endpoint:

1. In the Setup page, click IPSEC in the left-hand navigation frame, and then click Overview.

Figure 124: IPSec Overview



Peribit 1

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Setup**

BASIC

- Addresses
- Interfaces
- Time
- License Key
- SNMP
- Syslog Server
- Local Routes
- Registration Server

AAA

APPLICATIONS

IPSEC

- Setup Wizard...
- Overview
- Templates
- Default Policy

ADVANCED

**IPSec Overview** Find:  GO HELP

☒ 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

Name	Status	Template	Mgmt. Traffic*	Enter Pass Phrase	Verify Pass Phrase
<input checked="" type="checkbox"/> SR-10.15.2.12_Pri	●	Wizard	<input type="checkbox"/>	*****	*****
<input checked="" type="checkbox"/> SR-10.25.2.22_Pri	●	Wizard	<input type="checkbox"/>	*****	*****
<input type="checkbox"/> SR-10.2.2.32_Pri	⚠	Wizard	<input type="checkbox"/>	*****	*****

Select All Clear

\* When checked, Peribit management traffic (SSH/SSL) is included in the encryption tunnel.

Submit Refresh Reset

- Do one or more of the following and click Submit to activate the changes, or click Reset to discard them.

- To enable IPSec, click the check box next to **Enable IPSec Encryption for the endpoints selected below**. You can then select the check box next to the remote endpoints where you want to send encrypted traffic.

To view the list of endpoints starting with a specific device name, enter the first part of the name (or the entire name) in the Find box at the top of the page, and click GO. To select all devices displayed on the page, click Select All. To deselect all displayed devices, click Clear. If you disable an endpoint, all subsequent traffic to that endpoint is sent unencrypted.

For WX devices that support Multi-Path, a “\_Pri” or “\_Sec” is appended to the device name to indicate the primary or secondary path. You can enable IPSec for one or both paths. To configure Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 115.




- To change the common pass phrase or an individual pass phrase, enter and verify the new pass phrase (4 to 64 characters, eight is recommended) in the appropriate boxes. The pass phrase is used to generate a preshared key of the appropriate length.

To switch between common and individual pass phrases, select the appropriate radio button. If you select **Use a common pass phrase**, the individual pass phrases (if any) are retained for future use (click **Use individual pass phrases** to reactivate them).



**NOTE:** The pass phrase specified here must match the pass phrase specified on the remote device.

- To change the template for an endpoint, select a template from the Template drop-down menu. Note that two endpoints can establish a secure connection only if their IPSec templates specify the same authentication and encryption algorithms. To create new templates, refer to “Defining IPSec Templates” on page 222.
  - 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.
3. Click Refresh to update the icons in the **Status** column. The following icons are used to indicate the status of each IPSec connection:

Icon	Description
	<b>Normal operation</b> — A secure connection is established between this device and the remote device.
	<b>Configuration change</b> — New inbound and outbound security associations (SAs) are being negotiated due to a configuration change. Each SA specifies the algorithms and generated keys used to protect traffic in one direction. If this icon is displayed for more than a minute or two, the negotiation has failed and the old security association will eventually expire.
	<b>No security association</b> — A security association has not been negotiated. The default IPSec policy is applied to all traffic sent to this endpoint (refer to “Defining the Default IPSec Policy” on page 224).

4. To retain your changes when the device is restarted, click SAVE in the menu frame.


## 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 template created by the Setup Wizard or create new templates. To apply a template to an endpoint, refer to “Defining IPSec Settings by Endpoint” on page 219.

To define IPSec templates:

1. In the Setup page, click IPSEC in the left-hand navigation frame, and then click Templates.

**Figure 125: Defining IPSec Templates**



From the IPSec Templates page, you can:

- Add a new template, as described in Step 2 through Step 4.
- Change a template name or settings. Click the template name, change the template name and/or the settings, and click Submit.
- Delete a template. Click the check box next to the template name, and click Submit. If the deleted template is applied to an endpoint, the endpoint reverts to the Wizard template. The Wizard template can be changed, but not deleted.

2. To add a new template, click New Template.

**Figure 126: Defining a New IPSec Template**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, HELP, and a user status bar showing 'Logged in as: admin' with SAVE and LOGOUT buttons. The left sidebar is titled 'Setup' and contains a tree view with categories: BASIC (Addresses, Interfaces, Time, License Key, SNMP, Syslog Server, Local Routes, Registration Server), AAA, APPLICATIONS, IPSEC (Setup Wizard..., Overview, Templates, Default Policy), and ADVANCED. The main content area is titled 'IPSec Templates' and contains the following fields:

- Template Name:** A text input field.
- Encryption Algorithm:** A dropdown menu with 'AES-128' selected.
- Authentication Algorithm:** A dropdown menu with 'HMAC/SHA-1' selected.
- Key Lifetime:** A section with two sub-sections:
  - Time:** Radio buttons for 'Never expires' and 'Expires in' (with a value of 24 and the unit 'hours').
  - Data:** Radio buttons for 'Never expires' and 'Expires in' (with a value of 100 and the unit 'MB').

At the bottom of the form are three buttons: Submit, Reset, and Cancel.

3. Enter the following information:

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"> <li>■ <b>Any.</b> The algorithm selected for the peer endpoint is used. If both endpoints specify Any, AES-128 is used.</li> <li>■ <b>AES-128.</b> Advanced Encryption Standard with a 128-bit key.</li> <li>■ <b>AES-192.</b> AES with a 192-bit key.</li> <li>■ <b>AES-256.</b> AES with a 256-bit key.</li> <li>■ <b>3DES.</b> Triple Digital Encryption Standard with a 168-bit key.</li> </ul>
Authentication Algorithm	<p>Select the algorithm used to authenticate outbound traffic:</p> <ul style="list-style-type: none"> <li>■ <b>Any.</b> The algorithm selected for the peer endpoint is used. If both endpoints specify Any, HMAC/SHA-1 is used.</li> <li>■ <b>HMAC/SHA-1.</b> Secure Hash Algorithm.</li> <li>■ <b>HMAC/MD5.</b> Message Digest 5.</li> </ul>
Key Lifetime	<p>Specify the time and data limits for generated keys:</p> <ul style="list-style-type: none"> <li>■ <b>Time.</b> Enter the number of hours before a generated key expires (up to 2160), or select <b>Never expires</b>.</li> <li>■ <b>Data.</b> Enter the number of megabytes of traffic allowed before a generated key expires (up to 4000), or select <b>Never expires</b>.</li> </ul> <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>

- Click Submit to activate the changes, or click Reset to discard them.
- To retain your changes when the device is restarted, click SAVE in the menu frame.



## Defining the Default IPSec Policy

The default IPSec policy is applied to the following types of traffic:

- Passthrough traffic sent to unadvertised subnets (no remote WX device)
- Traffic between WX 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. To view the system log, refer to “Viewing and Saving System Logs” on page 277.

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 WX 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 Setup page, click IPSEC in the left-hand navigation frame, and then click Default Policy.

**Figure 127: Defining the IPSec Default Policy**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left-hand navigation menu is expanded to show the 'Setup' section, with 'IPSEC' selected. Under 'IPSEC', 'Default Policy' is highlighted. The main content area is titled 'IPSec Default Policy' and contains the following text:

For subnets not announced by a remote Peribit device, the following lists determine how packets destined for that subnet will be handled.

The lists also apply to subnets advertised by a Peribit device, which has been configured for encryption, but which has not successfully negotiated an IPSec security association.

Packets destined for subnets listed under 'Encryption Required' will be dropped and logged. Packets destined for subnets listed under 'Encryption Optional' will be passed-through unencrypted.

If both lists are blank, then all traffic will be considered 'Encryption' Optional and will be passed-through unencrypted.

Enter subnets, one per line, using the format: 10.123.0.0/255.255.0.0

Below the text are two text input fields: 'Encryption Required' and 'Encryption Optional'. At the bottom of the page are 'Submit' and 'Reset' buttons.



2. In the two text boxes, specify the destination addresses and subnets where encryption is required or optional, as follows:

Encryption Required	<p>Enter destination addresses or subnets (one per line) for which traffic must be dropped and logged. The subnet format is:</p> <p><b>&lt;IP address&gt;/&lt;subnet mask&gt;</b></p>
Encryption Optional	<p>Enter destination addresses or subnets (one per line) for which traffic can be sent unencrypted.</p> <p>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.</p>

3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.



## Chapter 9

# Monitoring and Reporting

This chapter describes how to view statistics for data reduction, bandwidth utilization, application acceleration, and overall traffic statistics. It covers the following topics:


- “Viewing and Printing Reports” in the next section
- “WAN Statistics” on page 228
- “Reduction Statistics” on page 235
- “Outbound Bandwidth Statistics” on page 244
- “Inbound Bandwidth Statistics” on page 246
- “Acceleration Statistics” on page 248
- “Traffic Statistics” on page 256
- “Endpoints Summary” on page 258
- “Executive Summary” on page 260

## Viewing and Printing Reports

---

Use the following methods to view additional details about report charts and graphs:

- Move the cursor over a bar chart, pie chart, or line graph to view the numerical values associated with each point on the chart or graph. Moving the cursor over the legend next to a pie chart has the same effect (clicking the legend highlights the corresponding wedge).
- On line graphs, move the cursor over a legend below the graph, such as **Data Out**, to show the y-axis values for all the associated points on the graph (Figure 128). Also, selecting the legend highlights the line on the graph.
- Click and drag the Zoom Scroller to the right of each line graph to zoom in on a portion of the graph. For example, to focus on the lower part of the graph, click the top of the Zoom Scroller and drag to the bottom (Figure 128).

**Figure 128: Viewing Graph Details**

- To view statistics for only prime time hours (if any), select Show Prime Time Only in the left-hand navigation frame, and click Submit. The selected report time period must be a day or longer. To set the prime time hours, refer to “Defining the Prime Time” on page 105.
- To print a report, select Printer Friendly Format in the left-hand navigation frame and click Submit. The report opens in a new browser window, and you can use the browser’s Print function to print the report.

## WAN Statistics

This section describes the WAN statistics displayed in the WXOS Web console.

- “WAN Throughput Statistics” in the next section
- “WAN Application Summary” on page 230
- “WAN Performance Statistics” on page 231

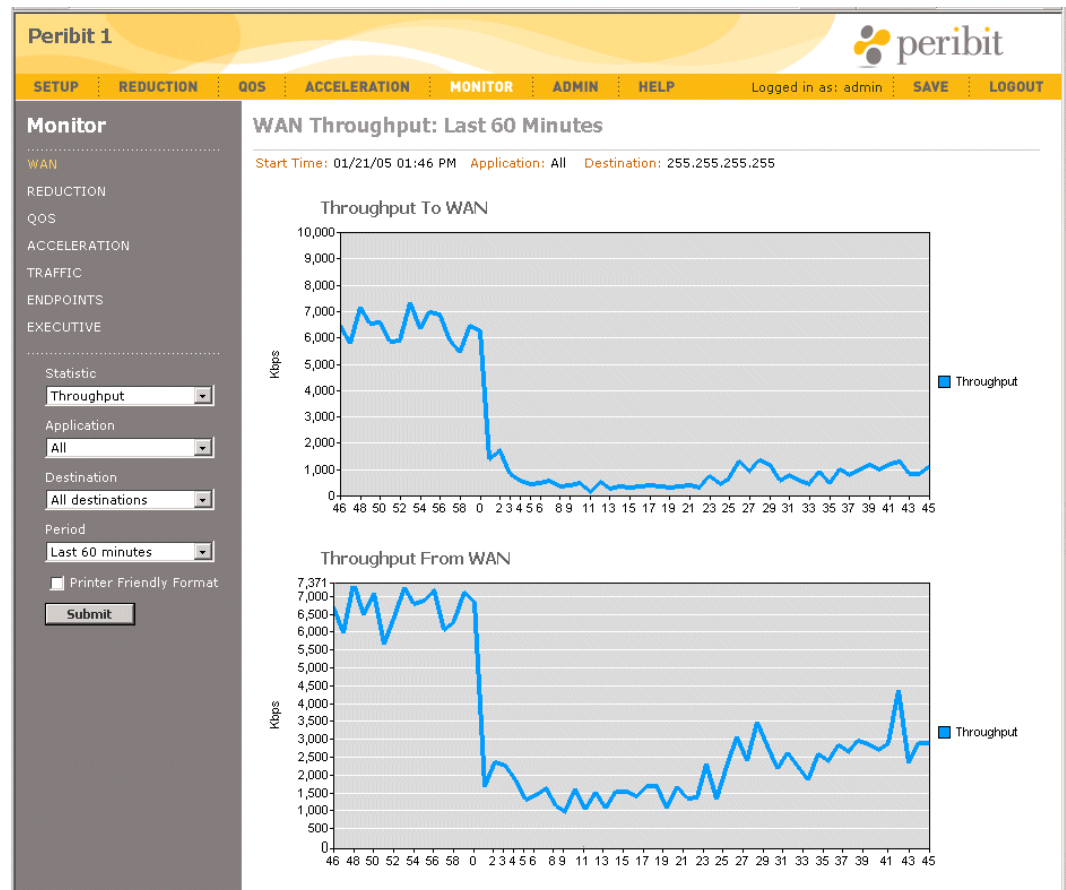
### WAN Throughput Statistics

The WAN throughput report shows separate graphs of the throughput to and from the WAN for all remote destinations, or for a specific WX device or virtual endpoint. To define virtual endpoints, refer to “Defining Outbound QoS Endpoints” on page 176. These statistics help you gauge the speed of the traffic to and from the WAN.

To view WAN throughput:

1. Click MONITOR in the menu frame, and WAN in the navigation frame.
2. Select Throughput from the Statistic menu, change one or more of the following report parameters, and click Submit.
  - Select a monitored application from the Application menu. Select Others to view statistics for applications that are undefined or unmonitored. The default is All. To specify the monitored applications, refer to “Reducing Applications” on page 135.

- Select a specific WX device or a virtual endpoint from the Destination menu to view the throughput to and from the WAN for the selected device.
- Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

**Figure 129: WAN Throughput Report**

3. 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.

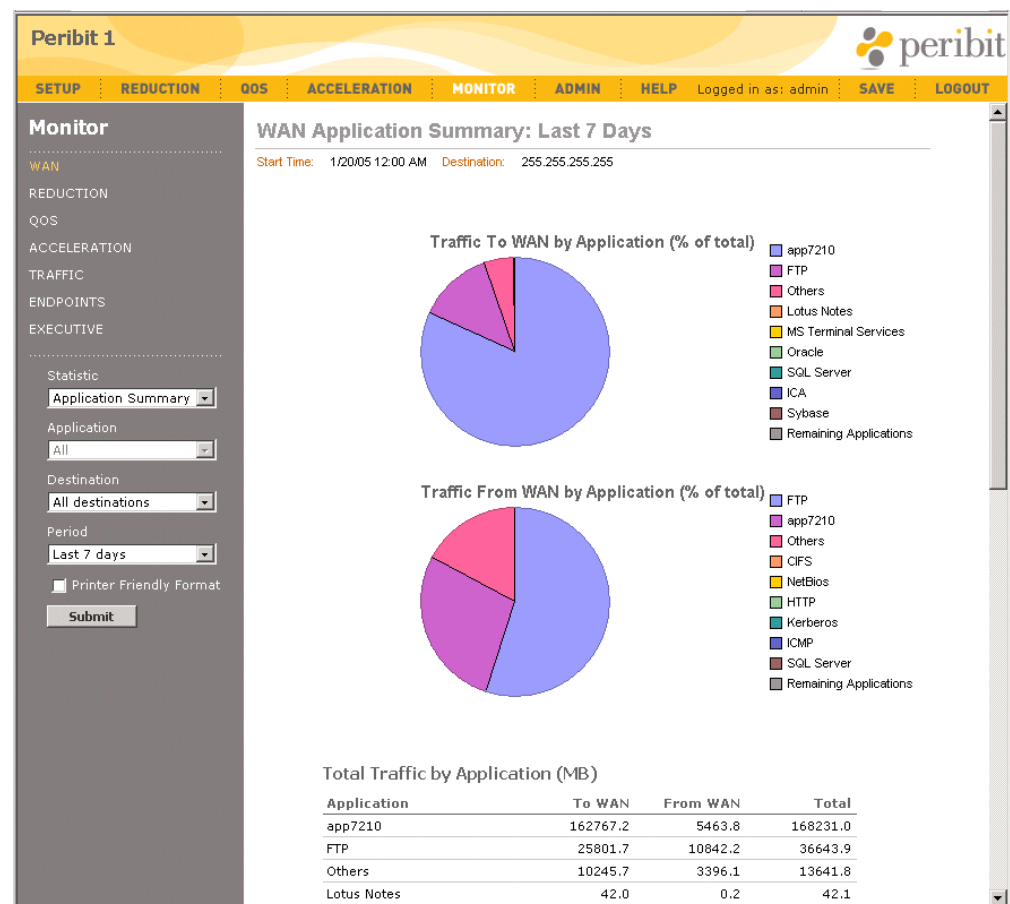
## WAN Application Summary

The WAN Application Summary shows the application traffic to and from the WAN for all remote destinations, or for a specific WX device or virtual endpoint. To define virtual endpoints, refer to “Defining Outbound QoS Endpoints” on page 176. The traffic to and from the WAN is shown for up to 40 monitored applications. To specify the monitored applications, refer to “Reducing Applications” on page 135.

To view the WAN Application Summary:

1. Click MONITOR in the menu frame, and WAN in the navigation frame.
2. Select Application Summary from the Statistic menu, change one or more of the following report parameters, and click Submit.
  - Select a specific device or a virtual endpoint from the Destination menu to view the application traffic to and from the WAN for the selected device.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

**Figure 130: WAN Application Summary**



3. 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 destination. 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. You can use the Traffic report to create definitions for the undefined applications that have the most traffic (refer to “Traffic Statistics” on page 256).

## WAN Performance Statistics

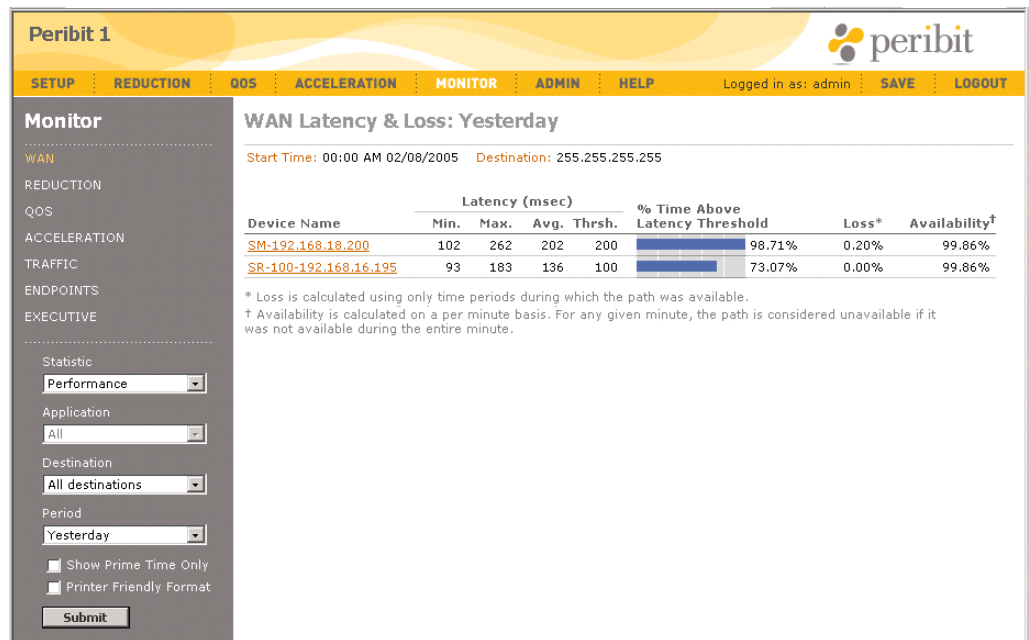
The WAN Performance report provides WAN loss and latency statistics, and performance events, between the current device and the remote WX devices that are enabled for either of the following:

- WAN performance monitoring (refer to “Configuring WAN Performance Monitoring” on page 124)
- Policy-Based Multi-Path, provided the local device is configured for Multi-Path and allows traffic for one or more traffic classes to change paths (refer to “Configuring Policy-Based Multi-Path” on page 115)

To view WAN performance:

1. Click MONITOR in the menu frame, and WAN in the navigation frame.
2. Select Performance from the Statistic menu, change one or more of the following report parameters, and click Submit.
  - Select a specific device from the Destination menu to view the performance graphs and events for the selected device. The default is All, which shows a table of performance statistics for all monitored devices.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.


Figure 131: WAN Performance Statistics



3. If the selected destination is All, the following information is shown for all monitored remote devices.
  - **Device Name.** Name of the remote WX 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 current device and the remote device (in milliseconds). The latency threshold for the remote device 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 WX probes that were lost.
  - **Availability.** Percentage of the minutes in the selected time period for which at least one probe was acknowledged. By default, 12 probes are sent per minute.
4. To view the performance graphs and events for a specific device, click the device name or select the device from the Destination menu. The information on the performance graphs depends on whether the device is enabled for Multi-Path (Figure 132) or WAN performance monitoring (Figure 133).





Figure 132: Multi-Path WAN Performance Charts



For a Multi-Path device, the following information is shown on the Loss and Latency charts (Figure 132):

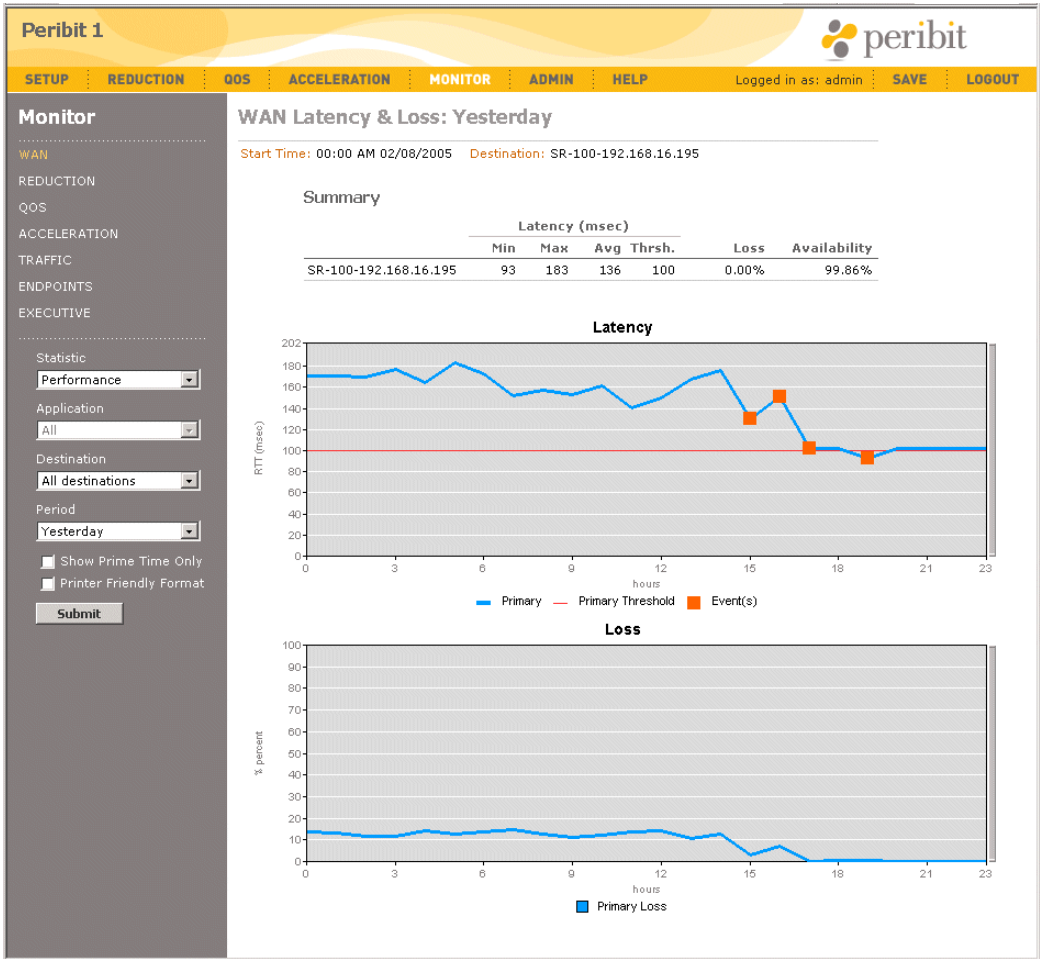
- The Latency chart shows the average round-trip time for the primary path (blue) and secondary path (black), and indicates the configured latency threshold for each path. The following icons are used to indicate performance events. An informational SNMP trap and a Syslog entry are generated for each event. Move the cursor over the icon to view the number of events in the time period.

Icon	Description
	Indicates that traffic was switched to the alternate path due to one of the following conditions: <ul style="list-style-type: none"> <li>■ <b>Loss or latency threshold exceeded.</b> Eligible traffic is diverted only if the alternate path's reduction tunnel is up and the loss and latency are below the specified thresholds.</li> <li>■ <b>Reduction tunnel is down.</b> 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).</li> </ul>

Icon	Description
	To view the status of the reduction tunnels, check the Multi-Path Endpoints page (refer to “Defining Multi-Path Endpoints” on page 121) or the Endpoints Summary report (refer to “Endpoints Summary” on page 258).
	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 WX 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 133: Single-Path WAN Performance Charts



For WAN performance monitoring endpoints (Figure 133), 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 WX 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.

## Reduction Statistics

This section describes the reduction statistics displayed in the WXOS Web console. There are four types of reduction statistics:

- “Device Throughput Statistics” in the next section
- “Data Reduction Statistics” on page 237
- “Application Summary Statistics” on page 240
- “Passthrough Statistics” on page 242
- “Packet Size Distribution Statistics” on page 243

### Device Throughput Statistics

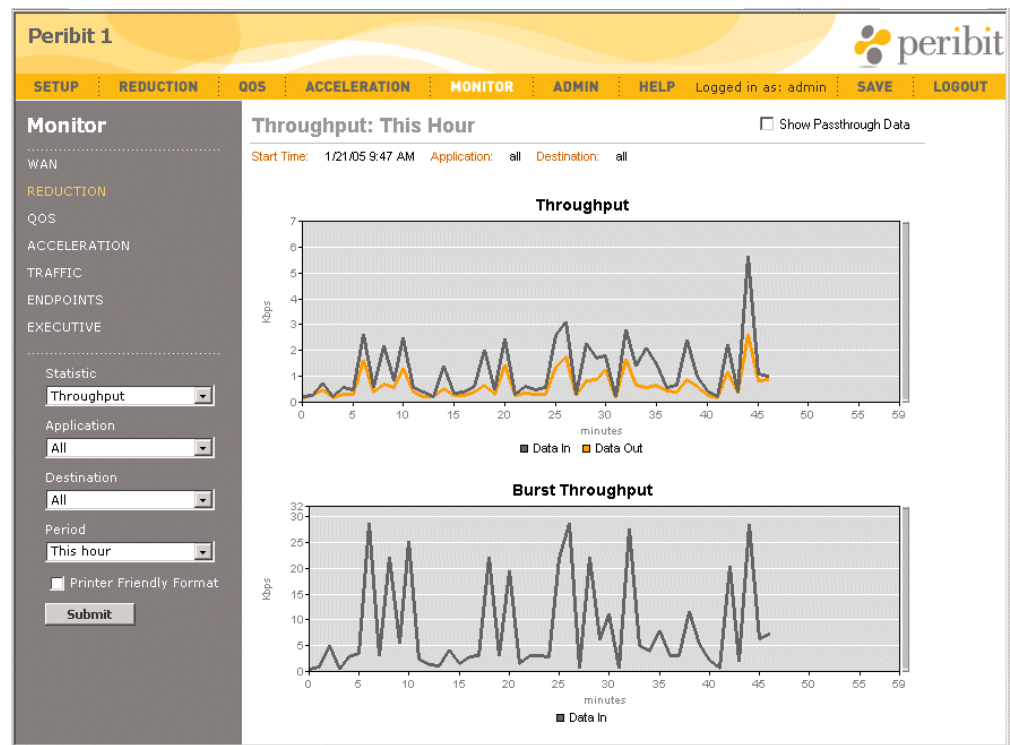
The device throughput statistics include a Throughput line graph and a Burst Throughput line graph. The Burst Throughput graph is shown only when you view the data for all applications and destinations. These statistics help you gauge the speed of traffic in and out of the device.

To view throughput statistics:

1. Click MONITOR in the menu frame.
2. Select Throughput from the Statistic menu, change one or more of the following report parameters, and click Submit.
  - Select a monitored application from the Application menu. Select Others to view statistics for applications that are undefined or unmonitored. The default is All. To specify the monitored applications, refer to “Reducing Applications” on page 135.
  - Select a specific WX device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

The Throughput page opens (Figure 134).

Figure 134: Device Throughput Statistics



3. 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 graph shows the following:
  - **Data In** (grey line). Average data throughput into the reduction engine.
  - **Data Out** (orange line). Average data throughput out of the reduction engine.
  - **Data In + Passthrough** (blue line). If All is selected from the Application and Destination menus, click the **Show Passthrough Data** check box at the top of the page to view the total average throughput into the device, including data that is passed through without being reduced.
- If All is selected from the Application and Destination menus, the Burst Throughput graph is displayed with the following:
  - **Data In** (grey line). Peak data throughput into the reduction engine. Based on five-second intervals for hourly reports, one-minute-intervals for daily reports, and one-hour intervals for weekly and monthly reports.
  - **Data In + Passthrough** (blue line). Click the **Show Passthrough Data** check box at the top of the page to view the peak throughput into the device, including data that is passed through without being reduced.



**NOTE:** The passthrough data shown here does not include the L2 multicast traffic. To view a breakdown of the passthrough traffic, including the amount of L2 multicast traffic, refer to “Passthrough Statistics” on page 242.

## Data Reduction Statistics

The data reduction statistics include a Summary table, a Percent Reduction graph, a Bytes graph, and a Packets graph. The Packets graph is shown only when you view the data for all applications. You can also view a details page that shows the percentage of data reduction achieved for the traffic sent to each of the other WX devices.

Note that the percentage of data reduction is not an average, but is based on the total number of bytes in and out of each device, as follows:

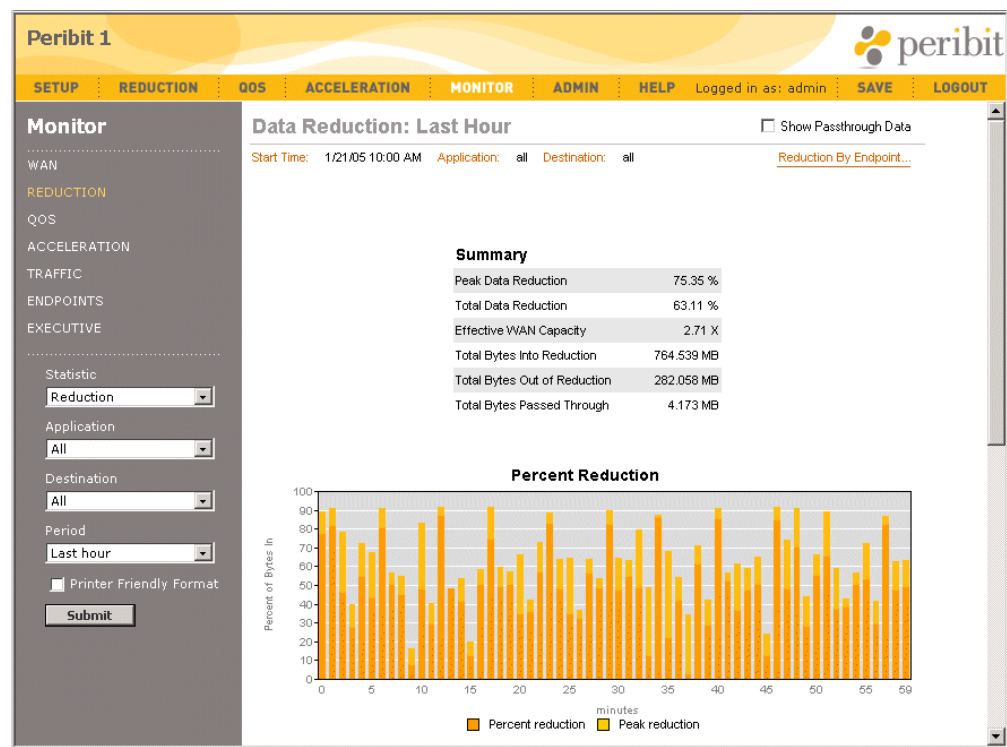
$$\% \text{ of Reduction} = [(\text{Bytes In} - \text{Bytes Out}) / \text{Bytes In}] \times 100$$

To view data reduction statistics:

1. Click MONITOR in the menu frame.
2. Select Reduction from the Statistic drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select a monitored application from the Application menu. Select Others to view statistics for applications that are undefined or unmonitored. The default is All. To specify the monitored applications, refer to “Reducing Applications” on page 135.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

The Data Reduction page opens (Figure 135).

Figure 135: Data Reduction Statistics



3. Review the following information on the data reduction graphs. Keep in mind that all values are for the selected application, destination, and time period.
  - The Data Reduction Summary table shows the following if All is selected from the Destination menu.
    - **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.
    - **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 %.
    - **Total Bytes Into Reduction.** Number of bytes into the data reduction engine.
    - **Total Bytes Out of Reduction.** Number of bytes after data reduction.
    - **Total Bytes Passed Through.** Number of bytes passed through without reduction. To view the different types of passthrough traffic, refer to “Passthrough Statistics” on page 242.




**NOTE:** If a specific device is selected from the Destination menu, the Summary table shows the total data reduction and the number of bytes in and out of the selected device, and for all devices in the community.

- If **All** is selected from the Application and Destination menus, click Reduction By Endpoint at the top of the page to view the data reduction for the traffic sent to each remote WX device. Note that historical data is maintained for at least two months, so devices may be listed that have no data for the selected time period.

Click a device name to view the data reduction by application for the device (to view the application statistics for all endpoints, refer to “Application Summary Statistics” on page 240).

- The Percent Reduction graph shows how the average and peak percentage of data reduction varied over the selected time period. Peak reduction is shown only for all applications.
- The Bytes graph shows the number of megabytes in and out of the device (Figure 136).

**Figure 136: Data Reduction Bytes Graph**



The Bytes graph includes the following:

- **Data In** (grey line). Number of bytes into the reduction engine.
- **Data Out** (orange line). Number of bytes out of the reduction engine.
- If **All** is selected from the Application and Destination menus, click the **Show Passthrough Data** check box at the top of the page to add a blue **Passthrough** line that shows the number of megabytes that are passed through the device without being reduced. The passthrough values are also added to the **Data In** and **Data Out** lines.




**NOTE:** The passthrough data shown here does not include the L2 multicast traffic. To view a breakdown of the passthrough traffic, including the amount of L2 multicast traffic, refer to “Passthrough Statistics” on page 242.



- If All is selected from the Application menu, the Packets graph is displayed. The Packets graph is similar to the Bytes graph, except that it shows the number of packets in and out of the device (Figure 137).

**Figure 137: Data Reduction Packets Graph**



### **Application Summary Statistics**


The Application Summary shows a pie chart of the nine monitored applications that have the highest percentage of the traffic into the WX device. A table is also included that shows the traffic statistics and percentage of data reduction for each monitored application (up to 40). To specify the monitored applications, refer to “Reducing Applications” on page 135.

To view application summary statistics:

1. Click MONITOR in the menu frame.
2. Select Application Summary from the Statistic drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.



Figure 138: Application Summary Statistics



3. Review the following information on the Application Summary. Keep in mind that all values are for the selected destination and time period.
  - 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. You can use the Traffic report to create definitions for the undefined applications that have the most traffic (refer to “Traffic Statistics” on page 256).
    - **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 135).

## Passthrough Statistics

Traffic that falls into one of several categories is passed through the WX device 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.

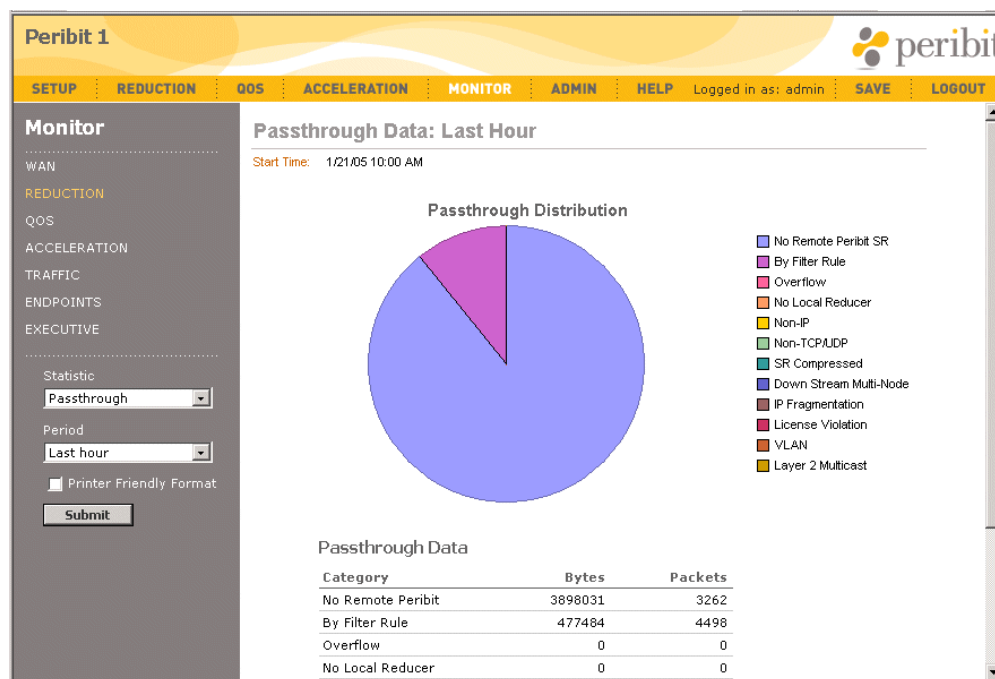


**NOTE:** There are no passthrough statistics for an off-path device where RIP is used to route traffic to the device. All traffic is sent through the reduction tunnel.

To view passthrough statistics:

1. Click MONITOR in the menu frame.
2. Select Passthrough from the Statistic drop-down menu.
3. Select a time period from the Period menu, and click Submit.

**Figure 139: Passthrough Statistics**



The following table describes the passthrough categories.

Category	Description
No Remote Peribit	No WX 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 (refer to “Reducing Applications” on page 135 and “Filtering Data Reduction by Source and Destination” on page 101).
Overflow	Traffic volume exceeded the device capacity.

Category	Description
No Local Reducer	Reduction is disabled on this device (refer to “Configuring Endpoints for Reduction Tunnels” on page 129).
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 WX device.
Down Stream Multi-Node	Traffic will be reduced by the next WX device (refer to “Multi-Node Configurations” on page 433).
IP Fragmentation	Always zero unless reduction of IP fragments is disabled (refer to “configure filter” on page 316).
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

The Packet Size Distribution report shows the number of packets in and out of the reduction engine, the percentage reduction in the number of packets, and the number of packets in each of six packet-size ranges.

To view packet size distribution statistics:

1. Click MONITOR in the menu frame.
2. Select Packet Distribution from the Statistic drop-down menu.
3. Select a time period from the Period menu and click Submit. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

**Figure 140: Packet Size Distribution Statistics**

## Outbound Bandwidth Statistics

If outbound QoS is enabled, the Outbound Bandwidth report shows the throughput of outbound traffic to the Remote interface and the amount of traffic dropped when one or more of the traffic classes exceeds its maximum allocated bandwidth. To configure outbound QoS settings, refer to “Configuring Outbound QoS Policies” on page 162.



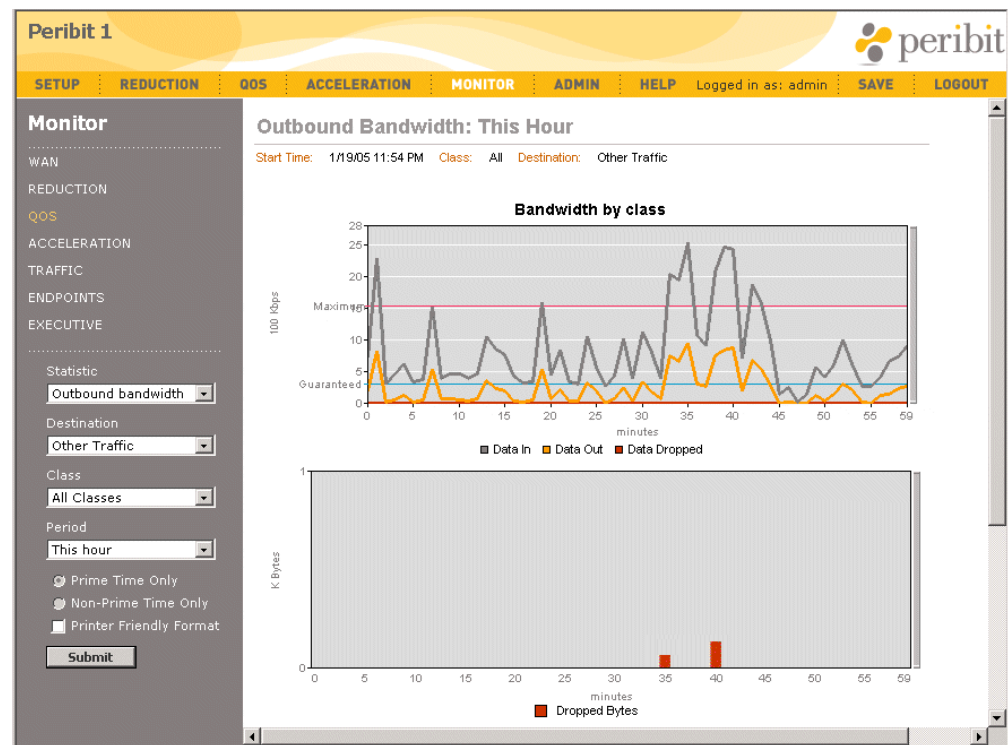
**NOTE:** Outbound bandwidth management is not effective for an off-path WX device unless all outbound WAN traffic is routed through the device.

To view outbound bandwidth statistics:

1. Click MONITOR in the menu frame, and Outbound Bandwidth in the navigation frame.
2. Select Outbound Bandwidth from the Statistic drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is Other traffic (all traffic that is not sent to a remote WX device).
  - Select a traffic class from the Class drop-down menu. The default is All Classes.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

If the selected time period is for a day or longer, you can select Non-Prime Time Only to view statistics based on the QoS policies defined for non-critical hours. By default, prime-time and non-prime time QoS policies are the same.

Figure 141: Outbound Bandwidth Statistics



3. Review the following information on the three line graphs. Keep in mind that all values are for the selected destination, traffic class, and time period.

- The Bandwidth by class graph shows the following:
  - **Data In** (grey line). Average data throughput into the Local interface from the LAN side of the WX device.
  - **Data Out** (orange line). Average throughput to the WAN side of the device. Indicates the data reduction achieved for the selected destination. The **Guaranteed** line shows the minimum bandwidth that is always available to the selected traffic class. If All Classes is selected, the guaranteed bandwidth is zero.



**NOTE:** The Data Out will be less than the remote circuit speed due to the overhead data produced by the WX device, but excluded from statistical reports.

- **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 (the **Maximum** line on the graph). If All Classes is selected, the maximum bandwidth is the circuit speed.

Note that brief bursts of traffic can cause data to be dropped, even when the average throughput is well below the maximum bandwidth.

- The Dropped Bytes and Dropped Packets graphs show the number of bytes and packets that were dropped when the maximum bandwidth for a traffic class (or the entire circuit) was exceeded.

## Inbound Bandwidth Statistics

---

If inbound QoS is enabled, the Inbound Bandwidth report shows the throughput of inbound traffic from the WAN and the amount of traffic dropped when one of the predefined traffic classes (Reduced, Intranet, TCP, and Default) exceeds its maximum allocated bandwidth. To configure inbound QoS settings, refer to “Configuring Inbound QoS Policies” on page 184.



**NOTE:** Inbound bandwidth management is not supported for off-path WX devices.


---

To view inbound bandwidth statistics:

1. Click MONITOR in the menu frame, click QOS, and select Inbound Bandwidth in the left-hand navigation frame.
2. Select Inbound Bandwidth from the Statistic drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select a traffic class from the Class drop-down menu. The default is All Classes.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

If the selected time period is for a day or longer, you can select Non-Prime Time Only to view inbound QoS statistics for non-critical hours.

Figure 142: Inbound Bandwidth Utilization Statistics



3. Review the following information on the three line graphs. Keep in mind that all values are for the selected traffic class and time period.
  - The Bandwidth by class graph shows the following:
    - **Data In** (grey line). Average data throughput from the WAN side of the device.
    - **Data Out** (orange line). Average data throughput out to the LAN side of the device.
    - **Data Dropped** (red line). Average rate that inbound data is dropped when the traffic for the selected class exceeds the maximum allocated bandwidth (the **Maximum** line on the graph). When All Classes is selected, the **Maximum** line is the inbound speed, which may be well above the maximum for the class(es) whose traffic is being dropped.
  - The Dropped Bytes and Dropped Packets graphs show the number of bytes and packets that were dropped when a traffic class exceeded its maximum allocated bandwidth.

## Acceleration Statistics

---

This section describes the statistics displayed for two types of Packet Flow Acceleration reports:

- “Active Flow Pipelining Statistics” in the next section
- “Fast Connection Setup Statistics” on page 250
- “Forward Error Correction Statistics” on page 252
- “CIFS and Exchange Acceleration Statistics” on page 253
- “HTTP Acceleration Statistics” on page 254

### Active Flow Pipelining Statistics

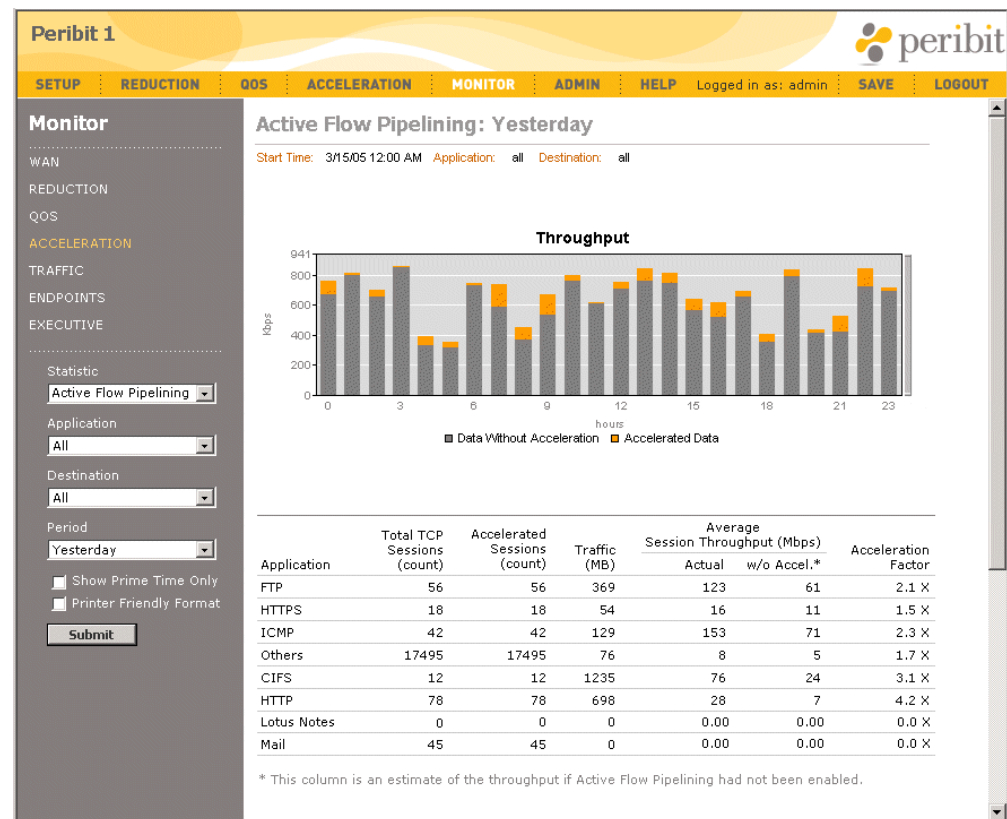
If Active Flow Pipelining is enabled for one or more endpoints and applications, the Active Flow Pipelining report shows the session statistics and the average throughput improvements due to AFP. To configure AFP for specific endpoints and applications, refer to “Accelerating WAN Traffic” on page 189.

To view Active Flow Pipelining statistics:

1. Click MONITOR in the menu frame, and then click ACCELERATION in the left-hand navigation frame.
2. Select Active Flow Pipelining from the Statistic drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select an application from the Application menu to view the acceleration statistics to each remote WX device. Select Others to view statistics for applications that are undefined or unmonitored. The default is All, which shows the average acceleration for all applications to all devices.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.



Figure 143: Active 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 was 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.

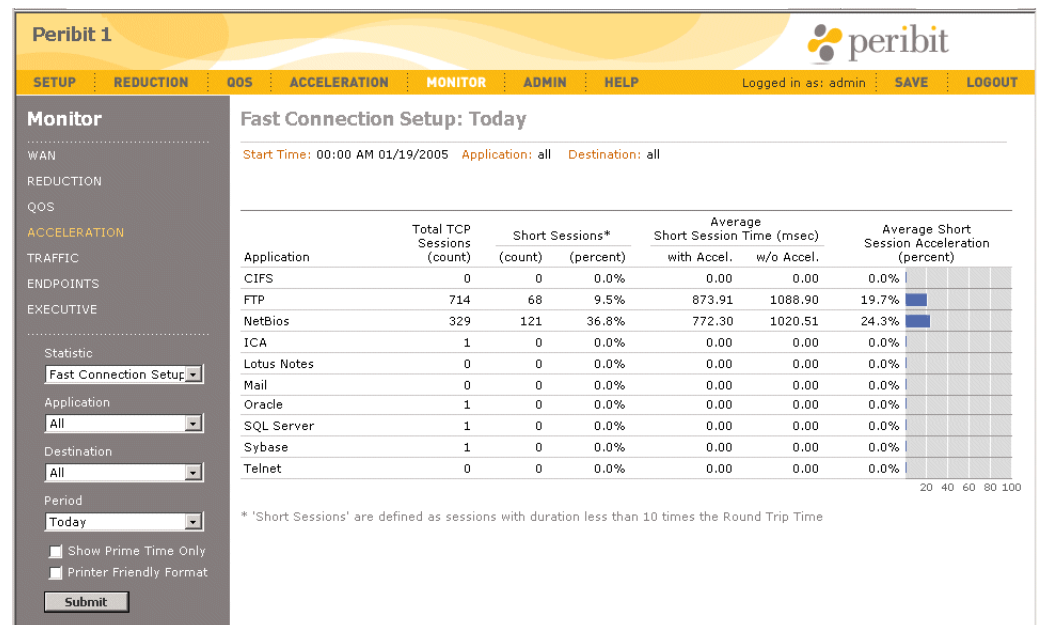
### **Fast Connection Setup Statistics**

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 configure Fast Connection Setup for specific endpoints and applications, refer to “Accelerating WAN Traffic” on page 189.

To view Fast Connection Setup statistics:

1. Click MONITOR in the menu frame, and then click ACCELERATION in the left-hand navigation frame.
2. Select Fast Connection from the Statistic drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select an application from the Application menu to view the acceleration statistics to each remote device. The default is All, which shows the average acceleration for all applications to all devices.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

Figure 144: Fast Connection Setup Statistics



- 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 WX 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.

## Forward Error Correction Statistics

The Forward Error Correction report shows the number of packets received from each remote endpoint that has forward error correction enabled. The report also shows the percentage of received packets that were lost, recovered, and retransmitted. The statistics are cumulative since the last time the counters were reset to zero. To reset the counters, use the CLI command:

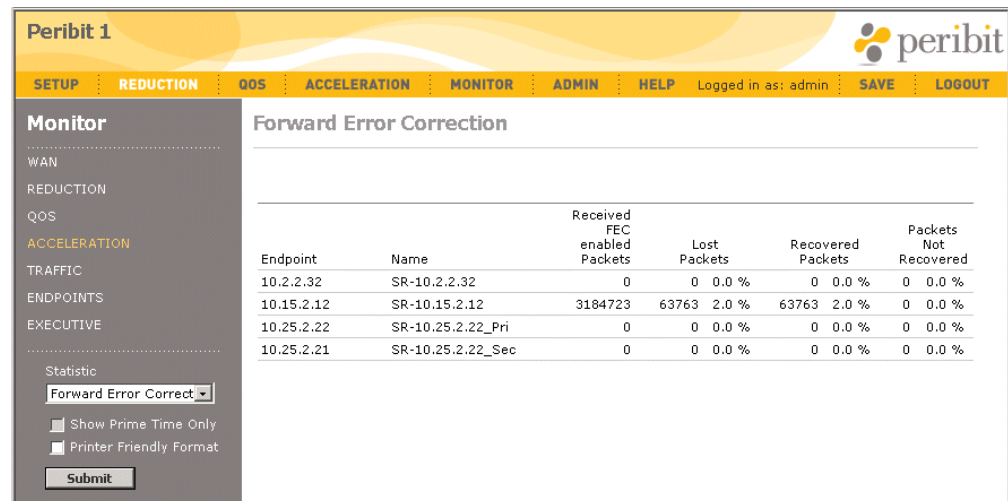
```
config acceleration forward-error-correction clear counters
```

To configure forward error correction for outgoing traffic to specific endpoints, refer to “Accelerating WAN Traffic” on page 189. Note that forward error correction is accepted on incoming traffic regardless of whether it is used for outgoing traffic.

To view forward error correction statistics:

1. Click MONITOR in the menu frame, and then click ACCELERATION in the left-hand navigation frame.
2. Select Forward Error Correction from the Statistic drop-down menu, and click Submit.

**Figure 145: Forward Error Correction Statistics**



3. Review the following information.
  - **Received Packets.** Number of error correction packets (data and recovery packets) received from the specified endpoint.
  - **Lost Packets.** Number and percentage of the received packets that were lost.
  - **Recovered Packets.** Number and percentage of the lost packets that were recovered using the recovery packets.
  - **Packets Not Recovered.** Number and percentage of the lost packets that had to be retransmitted.

## CIFS and Exchange Acceleration Statistics

If CIFS or Exchange application acceleration is enabled for one or more application definitions, the CIFS and Exchange acceleration reports shows the time saved due to CIFS and Exchange acceleration. To enable CIFS or Exchange acceleration, refer to “Application Flow Acceleration” on page 200.

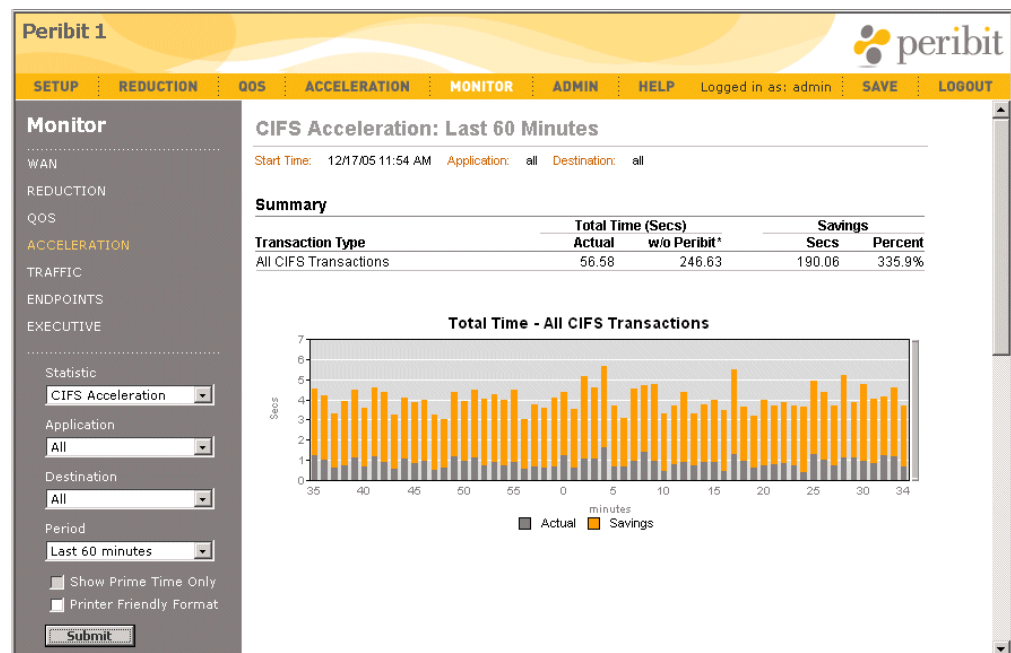


**NOTE:** View CIFS and Exchange acceleration reports on the client-side WX device. The acceleration statistics apply to the traffic in both directions. However, reduction statistics should probably be viewed on the server-side device.

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 drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select an application from the Application menu to view the acceleration statistics for a specific CIFS or Exchange application definition. The default is All.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

**Figure 146: 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 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 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.

### 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. To enable HTTP acceleration, refer to “Enabling HTTP Acceleration” on page 209.




**NOTE:** View HTTP acceleration reports on the client-side WX 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 drop-down menu, change one or more of the following report parameters, and click Submit.
  - Select an application from the Application menu to view the acceleration statistics for a specific HTTP application definition. The default is All.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

Figure 147: 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.

## Traffic Statistics

---

Traffic statistics are continuously collected for the most active traffic flows. The collected data for each flow includes the application name and protocol, the source and destination addresses and ports, and the number of bytes and packets sent and received. The collected statistics can be sent to a Cisco NetFlow server and displayed in the Web console. Undefined application flows displayed in the Web console are flagged so that you can quickly populate application definitions with the correct addresses and ports.



**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.

---

You can view the traffic statistics for the past hour, the past 24 hours, or all available hours (the length of time depends on the traffic volume). Up to 65,000 traffic 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 (for a description of the exported statistics, refer to “Top Traffic Export” on page 414).


To view the Traffic report:

1. Click MONITOR in the menu frame, and click TRAFFIC in the left-hand navigation frame.
2. To export the traffic statistics to a file in CSV format, click Export, enter the number of traffic flows you want to save, and click Export. To erase the current traffic statistics, click Clear.
3. To view the top 50 traffic flows that started in the past hour, click Submit.



Figure 148: Traffic Statistics



4. To create a new application definition with the addresses, ports, and protocol shown for a specific traffic flow, click **NEW APP** next to the traffic flow. For more information about defining applications, refer to “Managing Applications” on page 88. Note that an  is shown next to the flows for undefined applications.

5. To filter the traffic statistics, specify the following information and click **Submit**.

#### Statistic

Select a view of the traffic statistics. Each is displayed in descending order by traffic volume.

- **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

If you select the top flows, sending addresses, or receiving addresses, you can enter a subnet to view just the traffic from that subnet. The format is:

< IP address > / < subnet mask >

Where < subnet mask > is the number of bits used for the network portion of the address (such as “10.10.20.0/24”).

Traffic	<p>Select a view of the traffic for the selected statistic.</p> <ul style="list-style-type: none"> <li>■ <b>All.</b> All traffic for the selected statistic.</li> <li>■ <b>All Reduced.</b> Reduced traffic only.</li> <li>■ <b>Reduced Undefined Apps.</b> Reduced traffic for undefined applications only.</li> <li>■ <b>Passthrough Only.</b> Traffic sent from the WAN to the LAN that was not reduced. Does not apply to off-path WX devices or to in-line devices that use tunnel switching.</li> </ul>
Application	Select an application to limit the traffic to a specific application.
Port	<p>If you select the top flows, you can select a view of the port information.</p> <ul style="list-style-type: none"> <li>■ <b>Ignore Port.</b> Traffic is consolidated across all ports for each pair of source and destination addresses.</li> <li>■ <b>Source Only.</b> Traffic is consolidated across the same source ports for each pair of source and destination addresses.</li> <li>■ <b>Destination Only.</b> Traffic is consolidated across the same destination ports for each pair of source and destination addresses.</li> <li>■ <b>Source + Destination.</b> Traffic is shown for each combination of source and destination port.</li> </ul>
Aggregation mask	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 is "255.255.255.255", which shows a separate flow for each host. (This was the "Subnet Mask" field in previous versions of WXOS.)
Period	<p>Select the time period (last 60 minutes, last 24 hours, or all).</p> <p>Note that if you select <b>Last 60 minutes</b> or <b>Last 24 hours</b>, only the traffic flows that started in the selected time period are shown.</p>
Show reg. port names	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.
Show domain names	If you select the top flows, click the check box to view the domain names for each IP address. To specify the DNS servers to be queried, refer to "Configuring Device Address and Contact Information" on page 57. The IP address is displayed if its domain name cannot be resolved (the DNS queries may take a few seconds).

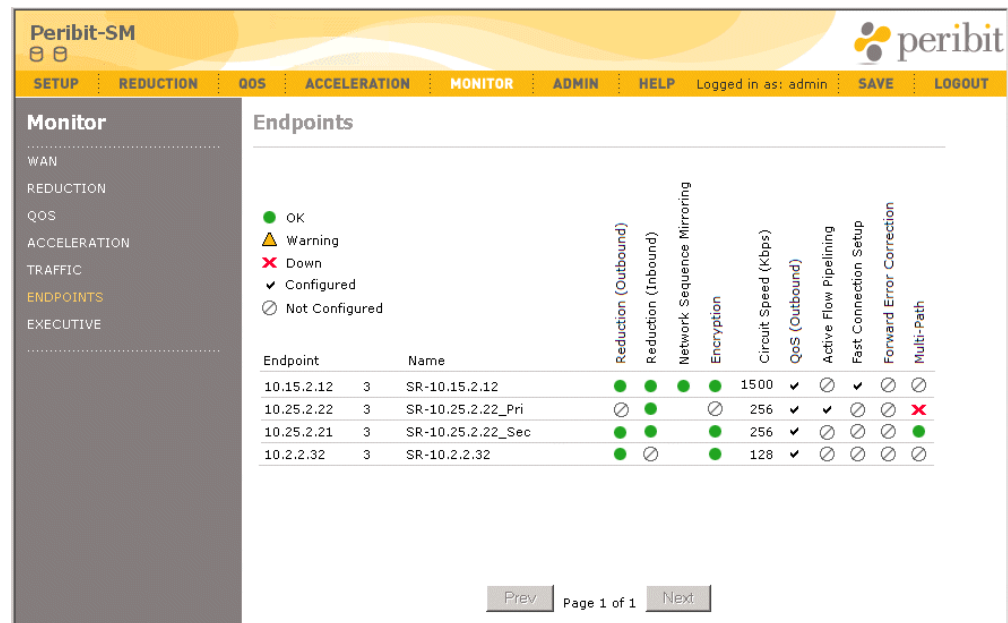
## Endpoints Summary

The Endpoints Summary report shows the status of reduction tunnels, Network Sequence Monitoring, IPSec encryption, Multi-Path, and Packet Flow Acceleration between the current WX device and each of the other devices in the community. The Endpoints Summary also indicates whether outbound QoS is enabled and, if so, the speed of the remote WAN circuit.

To view the Endpoints Summary:



1. Click MONITOR in the menu frame, and click ENDPOINTS in the left-hand navigation frame.

Figure 149: Endpoints Summary



2. The following icons are used to indicate the status of each connection:

Icon	Description
	<b>OK</b> — Indicates a connection between this device and the remote device for the following features: <ul style="list-style-type: none"> <li>■ Reduction (outbound or inbound)</li> <li>■ Network Sequence Caching (shown on WXC devices only)</li> <li>■ Encryption</li> </ul>
	<b>Warning</b> — Indicates that new IPsec security associations (SAs) are being negotiated due to an encryption configuration change. If this icon is displayed for more than a minute or two, the negotiation has failed and the old security association will eventually expire.
	<b>Down</b> — Indicates no connection between this device and the remote device for the following features: <ul style="list-style-type: none"> <li>■ <b>Reduction.</b> The outbound or inbound reduction tunnel is down (the remote device may be inaccessible).</li> <li>■ <b>Encryption.</b> A security association has not been negotiated, and the default IPsec policy is applied to all traffic sent to this endpoint (refer to “Defining the Default IPsec Policy” on page 224).</li> <li>■ <b>Network Sequence Caching.</b> A problem exists or NSC is enabled on the local device, but disabled on the remote device.</li> </ul>
	<b>Configured</b> — Indicates which of the following features are fully configured between the local device and each remote device (the feature must be enabled globally and for the remote device): <ul style="list-style-type: none"> <li>■ Outbound QoS</li> <li>■ Active Flow Pipelining (must also be configured on the remote device for AFP to take effect)</li> <li>■ Fast Connection Setup</li> <li>■ Forward Error Correction</li> </ul>

Icon	Description
	<b>Not configured</b> — The feature is not fully configured between this device and the remote device. However, in the <b>Reduction (Inbound)</b> column, this icon indicates that the remote device does not have a reduction tunnel enabled to the local device.
	<b>Unknown</b> — The connection is in a transitory state.
Blank	A blank in the Network Sequence Caching column indicates that the remote device is not a WXC or the reduction tunnel to the device is down. Also, the circuit speed is blank if outbound QoS is not configured.

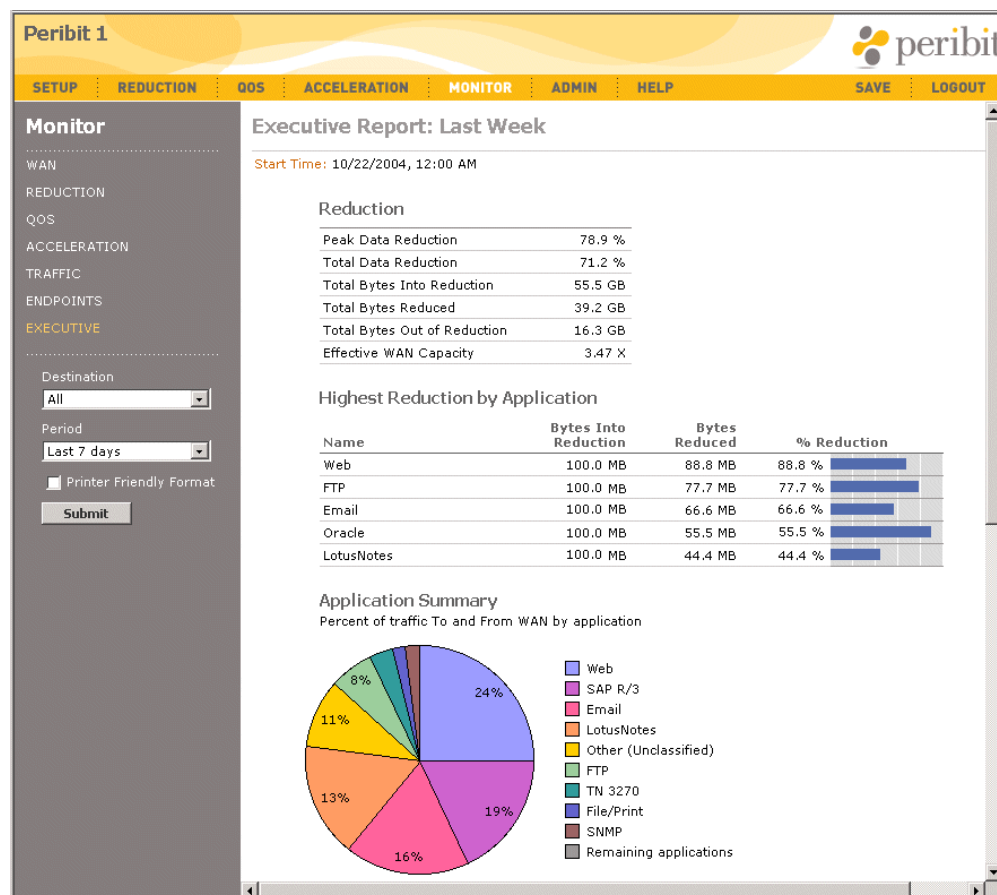
## Executive Summary

The Executive report summarizes reduction results, traffic volume by application, and average WAN performance (latency and loss) for one or all remote WX devices.

To view the Executive statistics:

1. Click MONITOR in the menu frame, and click EXECUTIVE in the left-hand navigation frame.
2. Optionally, change the following report parameters, and click Submit.
  - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All.
  - Select a time period from the Period menu. You can select the current or previous hour, day, or week. The default is Last 60 minutes.

Figure 150: Executive Summary



3. Review the following information. Keep in mind that all values are for the selected destination, and time period.
  - 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 Highest Reduction by Application table has the following columns.
  - **Application Name.** Names of the top five monitored applications with the highest reduction percentage. The **Others** category indicates the traffic for reduced applications that are undefined or unmonitored.
  - **Bytes Into Reduction.** Number of bytes into the device's reduction engine for each application.
  - **Bytes Reduced.** Number of bytes reduced for each application.
  - **Percent Reduction.** Percentage of data reduction achieved for each application.
- 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 the selected destination. 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 Application Volume by Application graph shows the traffic volume over the selected time period for the top nine monitored applications, plus the **Remaining applications** category.
- If WAN performance monitoring is enabled for the selected destination, the Average WAN Performance graph shows the average WAN latency and loss over the selected time period. If the selected destination is All, the graph averages the WAN latency and loss for all monitored WX endpoints (refer to "Configuring WAN Performance Monitoring" on page 124).

## Chapter 10

# Maintaining WX Devices

This chapter describes how to maintain the WX device through the Web console.

- “Maintaining Configurations and Software” in the next section.
- “Using Maintenance Tools” on page 271.



**NOTE:** If you have the Central Management System (CMS), you can schedule software and configuration updates for all WX devices in a community.

---

## Maintaining Configurations and Software

---

The following topics describe how to maintain the device’s configuration and software through the Web console:

- “Saving the Device Configuration” in the next section.
- “Displaying the Running Configuration” on page 265.
- “Loading a Device Configuration File” on page 266.
- “Loading a Boot Image” on page 267.
- “Clearing Application Monitoring Statistics” on page 268.
- “Setting the Device to the Factory Default Configuration” on page 268.
- “Rebooting the Device” on page 270.

## Saving the Device Configuration

When you change a device’s configuration, you must save the configuration file to Flash memory to preserve the settings the next time the device is restarted. You can also save the configuration file to another location for backup, such as an FTP or TFTP server. If a problem occurs where you must restore the factory default settings, you can load a saved configuration file to restore your network settings.



**NOTE:** A configuration file contains device-specific information, such as IP network addresses. Therefore, do not load the configuration file from one WX device to another.

---

1. Click ADMIN in the menu frame, and then click Save Configuration in the left-hand navigation frame.

**Figure 151: Saving the Configuration**

Peribit 1

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Admin**

MAINTENANCE

- Save Configuration
- Display Configuration
- Load Configuration
- Load Boot Image
- Clear Monitor Stats
- Set to Factory Default
- Reboot

TOOLS

**Save Configuration**

Save current device configuration to the following location:

☒ Flash memory ☒ Save to the default configuration file (startup.cfg)  
☐ Save to the filename:

☐ Local disk drive Save to your local disk drive via http transfer.

☐ TFTP server IP address:   
File path/name:

☐ FTP server IP address:   
File path/name:   
User name:   
Password:

Save

2. Select one of the following:

Flash memory	Save the current configuration to <i>startup.cfg</i> in Flash memory or click <b>Save to the filename</b> and enter another name. The name can be up to eight characters, with no file extension (such as “myconfig”). Click Save.  Note that <i>startup.cfg</i> is loaded each time you reboot the device. Always save the standard configuration to <i>startup.cfg</i> . Saving to a backup location is also recommended.
Local disk drive	Save the current configuration to the disk of a local machine in your network. Select this option, click Save, and specify the file name and location.
TFTP server	Save the current configuration to a TFTP server in your network. Enter the server’s IP address and a path and file name on the server, such as “/peribit/config_save.cfg”. Click Save.
FTP server	Save the current configuration to an FTP server in your network. Enter the server’s IP address and a path and file name on the server, such as “/peribit/config_save.cfg”. If the FTP server does not accept anonymous user access, enter a user name and password with read/write privileges to the server. Click Save.

3. After saving the configuration, you can reboot the device and reload the configuration settings if necessary.



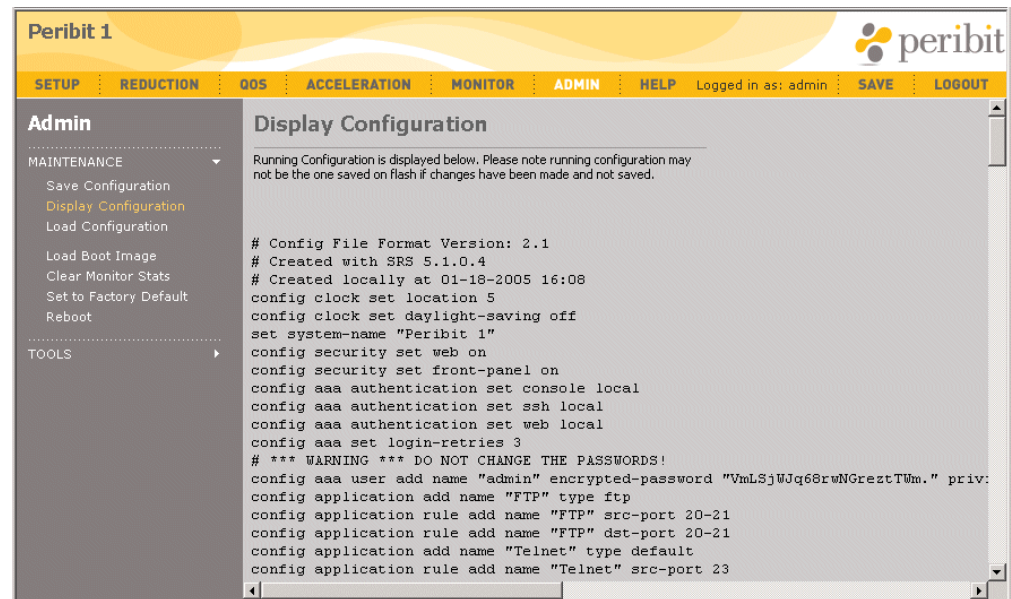
## Displaying the Running Configuration

The current configuration running on the device can be viewed through the Web console. The running configuration may be different from the configuration saved in Flash memory.

To view the running configuration:

1. Click ADMIN in the menu frame, and then click Display Configuration in the left-hand navigation frame.

**Figure 152: Displaying the Parameters of the Running Configuration**



2. Some configuration parameters can be set only through the CLI (refer to “Using the Command Line Interface (CLI)” on page 283).

## Loading a Device Configuration File

You can change a device's configuration by loading a configuration file that was previously saved to Flash memory, a local disk, or an FTP or TFTP server.



**NOTE:** A configuration file contains device-specific information, such as IP network addresses. Therefore, do not load the configuration file from one WX device to another.

To load a configuration file:

1. Click ADMIN in the menu frame, and then click Load Configuration in the left-hand navigation frame.

**Figure 153: Loading a Configuration File**



**NOTE:** Verify that the configuration file contains the correct configuration for the device. Loading an improper configuration file can have adverse effects on the device and on the other WX devices in the community.

2. Select the source for the configuration file (including location and file name), and then click Load.
3. To retain the configuration when the device is restarted, click SAVE in the menu frame to save the configuration to *startup.cfg* in Flash memory. This step is unnecessary if you load *startup.cfg* from Flash memory.
4. If the new configuration file changes the device's IP address, you MUST save the configuration to *startup.cfg* in Flash memory, and then reboot the device, as described in "Rebooting the Device" on page 270.

## Loading a Boot Image

To upgrade the WXOS operating system on a WX device, you can load a new boot image from a local disk or an FTP or TFTP server. You can then reboot the device to activate the new software. Loading a boot image does not affect the configuration settings stored in the *startup.cfg* file. All configuration information is preserved.

To load a boot image:

1. Click ADMIN in the menu frame, and then click Load Boot Image in the left-hand navigation frame.

**Figure 154: Loading a Boot Image**

2. Select the appropriate source and specify the software image (including location and file name), and click Load.



**NOTE:** To downgrade to a previous version of WXOS, select Allow image downgrade. Always save the current configuration file before upgrading to a new release so that you can reload the configuration if you must downgrade to the previous release.

3. Reboot the device to activate the new system software. Refer to “Rebooting the Device” on page 270 for more information.

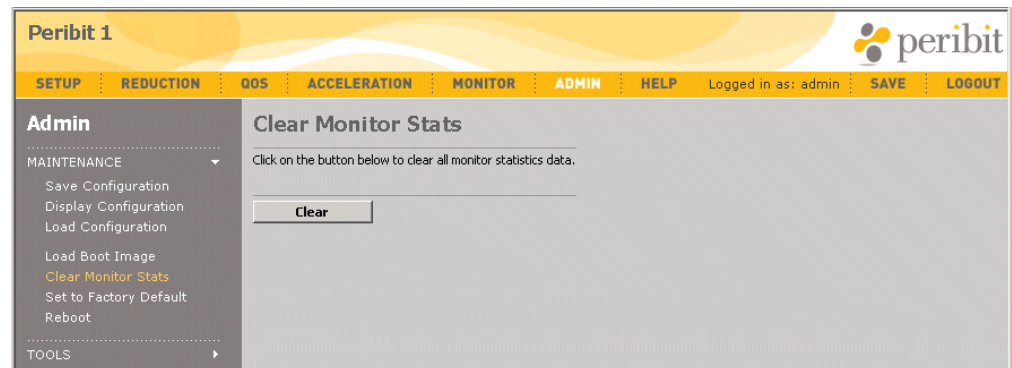
## Clearing Application Monitoring Statistics

At any time you can reset all the application monitoring statistics to zero. This may be useful during testing.

To clear the application monitoring statistics:

1. Click ADMIN in the menu frame, and then click Clear Monitor Stats in the left-hand navigation frame.

**Figure 155: Clearing Application Monitoring Statistics**



2. To clear the application monitoring statistics, click Clear.

## Setting the Device to the Factory Default Configuration

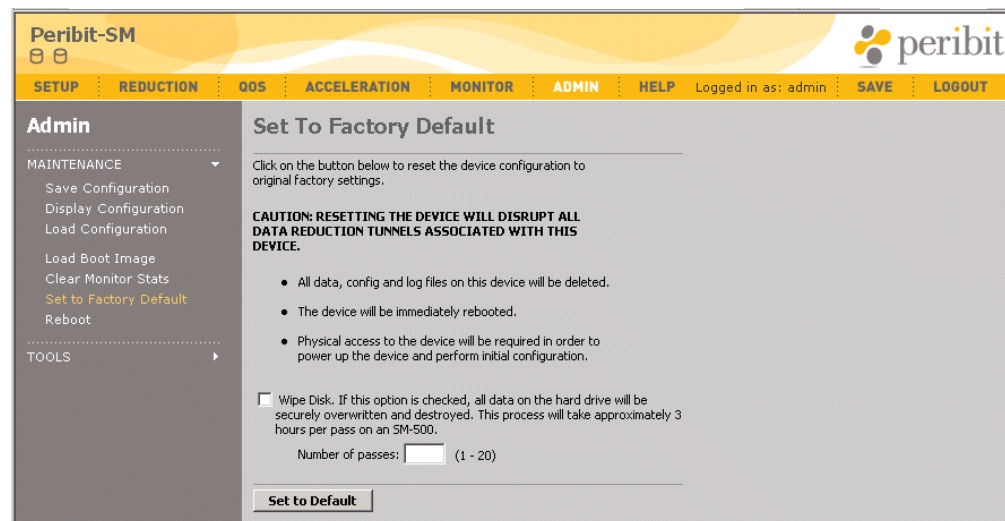
You can erase all device configuration information, including reduction statistics and network address information, by restoring the factory default configuration. This is useful when you must move the device to another location.



**NOTE:** Restoring the factory default configuration removes all data, configuration information and log files. It also disrupts the reduction tunnels associated with this device. Before you restore the factory default configuration, it is strongly recommended that you back up the configuration file to another location (refer to “Saving the Device Configuration” on page 263). In addition, you must have physical access to restart the device and do the initial configuration.

To set the device to the factory default configuration:

1. Click ADMIN in the menu frame, and then click Set to Factory Default in the left-hand navigation frame.

**Figure 156: Restoring the Factory Default Configuration Settings**

2. Before you set the device to its factory default configuration, verify that other devices in the community are not affected while this device is offline.
3. On a WXC device, you may want to wipe the hard disks for security purposes. Click the **Wipe Disk** check box, and enter the number of passes used to wipe the disks (up to 20). During each pass, a different value is written to each byte on the disks.

The first pass uses random numbers, the second pass writes a repeated pattern, the third pass uses zeros, the fourth pass writes another repeated pattern, while the fifth pass repeats the sequence with random numbers, shifted by one byte. Each pass takes about three hours. For maximum security, five passes are recommended. To stop the process, reboot the device.

4. Click Set to Default. If you elected to wipe the disks, the current pass number and the percent completion of the pass are displayed. After the disks are wiped, the factory defaults are loaded.
5. Wait until the LCD on the front panel displays the following:  
 “Factory Default. Power System Off”
6. Unplug the power cable from the back of the device, plug the cable back in, and then specify the IP address, subnet mask, and default gateway for the device (refer to “Installation” on page 29).

Rebooting the Device

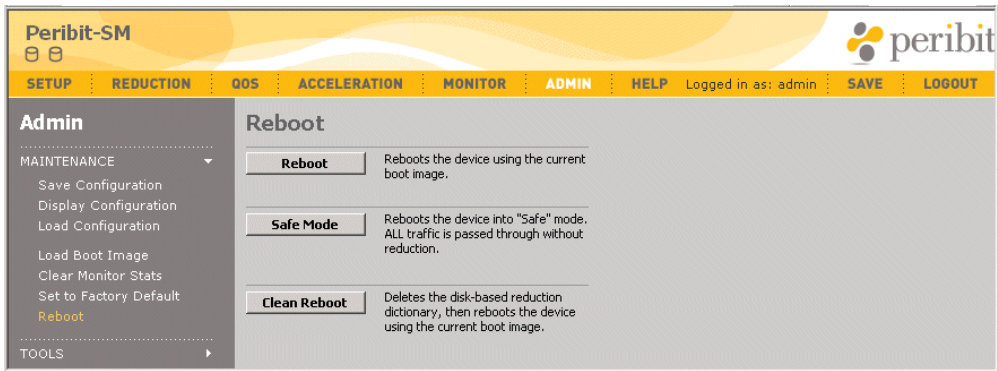
If you load a new boot image of the WXOS software on a device, you must reboot the device to activate the new software. During a reboot, the current boot image (*srs.os*) and the device configuration file (*startup.cfg*) are loaded from Flash memory into main memory. For a WX 15, you can use the CLI to select a specific boot image to be loaded (refer to “reboot” on page 295).

In addition, you can reboot a WX device in Safe Mode so that the power stays on, but traffic is passed through without reduction.

To reboot the device:

- 1. Click ADMIN in the menu frame, and then click Reboot in the navigation frame.

Figure 157: Rebooting the Device



- 2. Select one of the following:

Reboot	Performs a standard reboot of the device.
Safe Mode	Reboots the device so that the power is on, and the device can be configured, but traffic is passed through without reduction. Note the following: <ul style="list-style-type: none"><li>■ If IPSec is enabled, the default policy may cause traffic to be dropped (refer to “Defining the Default IPSec Policy” on page 224).</li><li>■ The warning “ERROR SW Passthru” is displayed in the front panel.</li><li>■ To exit Safe Mode, click <b>Reboot</b> to do a standard reboot.</li></ul>
Clean Reboot	Reboots the device and clears the reduction dictionary used for Network Sequence Monitoring. Available only on WXC devices.

## Using Maintenance Tools

The following topics describe how to use the maintenance tools through the Web console:

- “Pinging a Network Device” in the next section
- “Running a Traceroute to a Network Device” on page 272
- “Running a Packet Capture” on page 273
- “Generating NetFlow Records” on page 274
- “Entering CLI Commands from the Web Console” on page 275
- “Viewing and Saving System Logs” on page 277
- “Viewing and Saving the Access Control Log” on page 278
- “Exporting Performance Data” on page 279
- “Creating a Diagnostic File” on page 280
- “Viewing the WX 100 Server/Client Summary” on page 281

### Pinging a Network Device

You can use the ping utility to verify connections to other WX devices, or other network devices.

To use the ping utility:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Ping.

**Figure 158: Using the Ping Utility**

The screenshot shows the Peribit 1 web console interface. At the top, there is a navigation bar with tabs: SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The ADMIN tab is selected. Below the navigation bar, there is a left-hand navigation menu with sections: Admin (containing MAINTENANCE and TOOLS) and TOOLS (containing Ping). The Ping utility is selected in the TOOLS section. The main content area displays the Ping utility form with the following fields: Destination IP address (text input), Data Size (32 bytes), and Number of times (3). There are Submit and Reset buttons at the bottom of the form.

2. In the destination field, enter the IP address of a WX device or other network device.
3. Optionally, enter the size of each ping packet (8 to 4068 bytes), and the number of packets sent (1 to 50).



- Click **Submit** to ping the device. The results are shown in the Web console, including the round-trip time of each packet (in milliseconds).

For example:

```
PING 192.168.0.127: 32 data bytes
40 bytes from 192.168.0.127: icmp_seq=0. time=2. ms
40 bytes from 192.168.0.127: icmp_seq=1. time=2. ms
40 bytes from 192.168.0.127: icmp_seq=2. time=2. ms
—192.168.0.127 PING Statistics—
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms) min/avg/max = 2/2/2
```



**NOTE:** If you ping an address that is advertised by an off-path WX device that uses RIP for packet interception, the ping packets are routed through the WX device, which may affect the round trip times.

## Running a Traceroute to a Network Device

You can use the traceroute utility to determine the number of router hops and the route taken from the WX device to another network device. This tool can help you determine the point in your network that may be causing a connection failure.

To use the traceroute utility:

- Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Traceroute.

**Figure 159: Using Traceroute to Determine the Network Path to a Device**

- Enter the IP address of the destination device, and the maximum number of router hops (1 to 30) to search for that device.
- Click Submit. The results are displayed in the Web console, including the IP address of each device in the path, and the round-trip time (in milliseconds) of each of the three packets sent to identify each hop. For example:

```
traceroute to 192.168.0.127 (192.168.0.127), 10 hops max, 40 byte packets
```



```

1 192.168.53.130 2 ms 0 ms 0 ms
2 192.168.53.70 2 ms 2 ms 4 ms
3 192.168.53.1 0 ms 2 ms 2 ms
4 192.168.52.15 2 ms 2 ms 2 ms
5 192.168.0.127 2 ms 2 ms 2 ms

```



**NOTE:** If you trace an address that is advertised by an off-path WX device that uses RIP for packet interception, the trace packets are routed through the WX device, which may affect the number of hops and round trip times.

## Running a Packet Capture

The packet capture utility lets you capture raw network data from the device's Local and/or Remote interfaces. The packet capture information can then be exported to a file and analyzed by a protocol analyzer program or other hardware. The packet capture's file format is either "libpcap" or "snoop". You must enter the packet-capture password to run or save a packet capture.



**NOTE:** If tunnel switching is enabled, running a packet capture will capture intermediate assembled packets before they are recompressed. These packets have zeros for the source and destination, and may have checksum errors. These packets are internal to the device and can be ignored.

To use the packet capture utility:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Packet Capture.

**Figure 160: Using the Packet Capture Utility**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The ADMIN link is highlighted. Below the navigation bar, the left sidebar shows the 'Admin' menu with sub-items like MAINTENANCE and TOOLS. The 'TOOLS' menu is expanded, showing options like Ping, Traceroute, Packet Capture (highlighted), NetFlow, and Command Line Interface. The main content area displays the 'Packet Capture' configuration page. It includes a 'Status' section showing 'Ready'. The configuration section has several fields: 'Interface' set to 'Local', 'Size' (empty), 'Max. Number Packets' set to 'All', 'Snap Length' set to '1514 Bytes', 'Storage Format' set to 'libpcap', and 'Delete After' set to '1 hours'. There is a 'Packet Capture Password' field and buttons for 'Start', 'Stop', 'Save...', and 'Delete'. A note at the bottom states: 'To start a packet capture or to save a packet capture to a local disk, you must enter the packet capture password.'

2. Specify the following information:

Interface	Select the interface(s) where you want to capture packets (Local, Remote, or Both).
Size	Enter the number of bytes to be captured (minimum is 4096). Execution stops when the specified number of bytes are captured.
Max. Number Packets	To limit the capture to a maximum number of packets, select the second option and enter the number of packets.
Snap Length	Enter the maximum number of bytes captured for each packet (1 to 65535). The default is 1514. Select All to capture the entire packet.
Storage Format	Select the format of the captured data (libpcap or snoop). The default file name is <i>pkgdump.dmp</i> .
Delete After	Enter the number of hours that the packet capture file is retained (1 to 168).
Packet Capture Password	Enter the packet capture password. To change the password, refer to "Changing the Packet Capture Password" on page 88.

3. To start the packet capture, click Start. The status is displayed in the upper-right corner of the page. Click Stop at any time to stop the capture.
4. To save the packet capture, click Save, and specify a file name and location.
5. To manually delete the packet capture file, click Delete. You cannot run another packet capture until the previous one is deleted.

## Generating NetFlow Records

Traffic data is collected continuously 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. The collected statistics can be sent to a Cisco NetFlow server and displayed in the Web console.

To generate NetFlow records:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click NetFlow.

**Figure 161: Generating NetFlow Records**

2. Click Enable NetFlow, and enter the IP address and port number of a NetFlow server.
3. Click Submit to activate the changes, or click Reset to discard them.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

NetFlow data is sent in Version 5 format, as described in “NetFlow Version 5 Export” on page 405.

### **Entering CLI Commands from the Web Console**

Some options are available only through the command line interface (CLI). You can enter CLI commands from the Web console as described here, or from a Secure Shell (SSH) program or a terminal connected to the serial port, as described in “Accessing the CLI” on page 283.

Note the following when entering commands from the Web console:

- CLI configuration commands are applied to the candidate configuration when you click Submit. Use the “commit” command to apply changes to the running configuration.
- Only “show,” “configure,” and “commit” commands are supported, as well as “ls” and “pwd”. The “cd” command is not supported, nor are any commands that require user interaction, such as “rollback”, “save-config”, and “import-route-table”.
- The entire command must be entered on one line. For example, you cannot enter “configure” and “application” on separate lines, as you can in a Secure Shell (SSH) or terminal emulation program.
- Multiple commands can be entered together (one per line), as in a script. Up to 10 KB of commands can be entered at once.

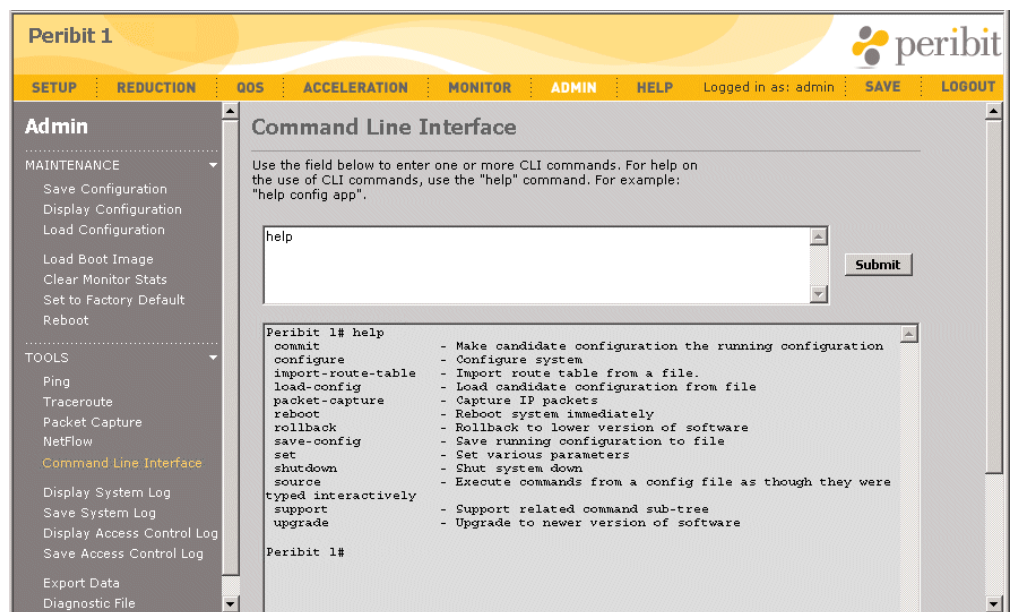
- To view the online help for a command, type “help” before the command. You cannot use “?” to view the help. Also, the CLI keyboard shortcuts are not supported.

For more information about the CLI commands, refer to “Using the Command Line Interface (CLI)” on page 283.

To enter CLI commands from the Web console:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Command Line Interface.

**Figure 162: Entering CLI Commands**



2. Enter one or more CLI commands (one per line) in the upper list box, and click Submit. The results are displayed in the lower list box. To clear the results, delete all the commands and click Submit.
3. To apply all configuration changes to the running configuration, enter the “commit” command and click Submit.
4. To retain your changes when the device is restarted, click SAVE in the menu frame.

## Viewing and Saving System Logs

The system log files can be displayed in the Web console. You can also download these log files to a local machine for use by third-party applications. If your network has dedicated Syslog servers, you can configure the WX device to send log messages to up to five Syslog servers, as described in “Enabling Syslog Reporting” on page 65.

To view or download system log files:

1. Click ADMIN in the menu frame, and click TOOLS in the left-hand navigation frame.
2. To display the system log file, click Display System Log in the left-hand navigation frame. The current system log is displayed in the Web console. The most recent entries are displayed last.
3. To download a system log file for a specific time period, click Save System Log in the left-hand navigation frame.

**Figure 163: Saving the System Log file for the Running Configuration**

The screenshot shows the Peribit 1 Web console interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The user is logged in as 'admin'. The left-hand navigation menu is expanded to 'Admin' > 'TOOLS', where 'Save System Log' is highlighted. The main content area is titled 'Save System Log' and contains a table of log files.

File name	First entry time
pnlog	2005-01-18 00:00:24
pnlog1	2005-01-17 00:00:25
pnlog2	2005-01-16 00:00:27
pnlog3	2005-01-15 00:00:29
pnlog4	2005-01-14 00:00:30
pnlog5	2005-01-13 00:00:32
pnlog6	2005-01-12 00:00:34
pnlog7	2005-01-11 00:00:37
pnlog8	2005-01-10 00:00:38
pnlog9	2005-01-09 00:00:42
pnlog10	2005-01-08 00:00:44

The pnlog file contains the most recent data. Each time pnlog reaches 1 MB in size, it is saved as pnlog1, and the existing log files are renumbered up to pnlog10 (older log files are discarded). The First entry time column shows the oldest entry in each log file.

4. Click the name of the log file you want to save, click Save, and specify a file name and location.

## Viewing and Saving the Access Control Log

The access control log shows the user name, date, and time of each user who accessed the device in the past five days, as well as the configuration changes made by each user. The access method is shown as SSH (CLI access), HTTPS (Web access), or CONSOLE (direct access). The workstation IP address is included for SSH and HTTPS.

For example, the following entries indicate that a user logged in from the Web console, changed the prime time setting, and committed the change by clicking Submit. The “Created locally” entries indicate the time stamp of the previous and current configuration. The “CHANGED” entries indicate the previous and current values.

```
HTTPS: 192.168.0.76 admin Login 2005-05-13 08:10:19 HTTP/1.1 POST / 0
HTTPS: 192.168.0.76 admin Commit config 2005-05-13 08:11:48 0
CHANGED:
< # Created locally at 05-13-2005 07:03
—
> # Created locally at 05-13-2005 08:11
ADDED:
> config prime-time set mode on
CHANGED:
< config prime-time set hours 0-24
—
> config prime-time set hours 7-18
```



**NOTE:** The access log has six files. Viewing or saving the access log concatenates the data from all the files.

---

To view or download an access control log file:

1. Click ADMIN in the menu frame, and click TOOLS in the left-hand navigation frame.
2. To display the access control log, click Display Access Control Log in the left-hand navigation frame. The access control log is displayed in the Web console. The most recent entries are displayed last.
3. To download the access control log, click Save Access Control Log in the left-hand navigation frame, click Save, and specify a file name and location.


## Exporting Performance Data

You can export the performance data for all time periods 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 performance data is similar to the data displayed in the Monitor page of the Web console (refer to “Monitoring and Reporting” on page 227).

To export performance data to CSV format:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Export Data.

**Figure 164: Exporting Performance Data to CSV Format**



2. In the Export Data page, select All (ZIP format) to export the data for all time periods as a “.zip” file. If you cannot open a “.zip” file (some browser versions cannot), select All (CSV format).

Refer to Appendix , “Understanding Exported Data Results” for a description of the CSV data file.

3. Click Submit, and then click Save and specify a file name and location.

## Creating a Diagnostic File

If you experience problems with a WX device, you can generate a diagnostic file to send to Technical Support. The diagnostic file contains current configuration, filter settings, system information, and the most recent log files. After you generate and save the diagnostic file, email it to [support@juniper.net](mailto:support@juniper.net).

To create and send a diagnostic file to Technical Support:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Diagnostic File.

**Figure 165: Creating a Diagnostic File**

**Peribit 1**

SETUP REDUCTION QOS ACCELERATION MONITOR ADMIN HELP Logged in as: admin SAVE LOGOUT

**Admin**

MAINTENANCE

- Save Configuration
- Display Configuration
- Load Configuration
- Load Boot Image
- Clear Monitor Stats
- Set to Factory Default
- Reboot

TOOLS

- Ping
- Traceroute
- Packet Capture
- NetFlow
- Command Line Interface
- Display System Log
- Save System Log
- Display Access Control Log
- Save Access Control Log
- Export Data
- Diagnostic File**

**Diagnostic File**

This page allows you to create a diagnostic file containing the current configuration, filter settings, and all logs for this device. This file can be emailed to Peribit Networks and can provide useful information to aid in the diagnosis of problems.

When you click the 'Submit' button, the file will be generated and downloaded to your computer. When the dialog box appears, choose 'Save this file to disk'. After the file is downloaded, attach it to an email and send it to [support@peribit.com](mailto:support@peribit.com).

For faster response, fill out the information below.

Name:

Company:

Phone:

Email:

Problem description:

2. Complete the form so that your contact information is included with the diagnostic file.
3. Click Submit to generate the diagnostic file, and then click Save and specify a file name and location. Note that a diagnostic file for an WX 100 also includes information for the client devices (if any).

Email the diagnostic file as an attachment to [support@juniper.net](mailto:support@juniper.net). A support representative will contact you.



## Viewing the WX 100 Server/Client Summary

An WX 100 can act as a server for up to six client devices. The Server/Client Summary page lets you view the port number, status, model number, and number of tunnels for each client connected to the WX 100.

To view the Server/Client Summary on an WX 100:

1. Click ADMIN in the menu frame, click TOOLS in the left-hand navigation frame, and then click Server/Client Summary.

**Figure 166: Viewing the WX 100 Server/Client Summary**

The screenshot shows the Peribit 1 web interface. The top navigation bar includes links for SETUP, REDUCTION, QOS, ACCELERATION, MONITOR, ADMIN, and HELP. The ADMIN link is selected. The left-hand navigation frame shows the TOOLS menu expanded, with Server/Client Summary selected. The main content area displays the Server/Client Summary page, which includes a table of connected clients.

Port	Status	Model	No. of Tunnels	
			OUT	IN
1	Active	WXC-500	3	1
2	Active	WXC-500	1	1
3	Active	WXC-500	0	1
4	Not Connected			
5	Not Connected			
6	Not Connected			
stack master	Active	SR-100	1	2

2. Review the following information:

Port	Port number on the WX 100 where a client device is connected. The port number becomes the client ID, and is shown on the front panel of the client device.
Status	Indicates the port status: <ul style="list-style-type: none"> <li>■ <b>Active.</b> Client connected and processing traffic.</li> <li>■ <b>Passive.</b> Client connected, but idle.</li> <li>■ <b>Not Connected.</b> No client installed.</li> </ul>
Model	Model number of the client device.
No. of Tunnels	Number of tunnels handled by each client and the WX 100. Note that remote devices see only the WX 100, not the client device that is actually processing the traffic.



## Chapter 11

# Using the Command Line Interface (CLI)

The following topics describe how to use the command line interface (CLI) to configure WX devices:

- Accessing the CLI on page 283
- Logging In Using the CLI on page 284
- CLI Basics on page 285
- CLI Command Summary on page 287
- System-Level Commands on page 290
- Configuration Commands on page 300
- Show Commands on page 368



**NOTE:** You should use the Web console for most configuration tasks. However, the CLI provides some additional options that may be useful in special circumstances.

---

## Accessing the CLI

---

The following sections describe two ways to access the CLI:

- Using a Secure Shell Program from a Remote Workstation on page 284
- Using a Terminal Connected to the Serial Port on page 284

You can also access the CLI from the Web console, as described in “Entering CLI Commands from the Web Console” on page 275.

### Using a Secure Shell Program from a Remote Workstation

Secure Shell (SSH) is an application program that provides authentication and encryption capabilities for secure Internet communications. You can download SSH client software from the following site:

<http://www.openssh.com>

Because there are many different types of SSH applications available, it is recommended that you read the instructions for your specific SSH application.



**NOTE:** WX devices support SSH version 2.5 with DES/3DES encryption.

---

### Using a Terminal Connected to the Serial Port

You can connect a terminal to the serial port on the WX device, and then use a terminal emulation program (such as HyperTerminal) to log in to the CLI and enter configuration commands. Some terminal emulation programs also include a Secure Shell.

Use a female/female DB-9 crossover cable (null-modem cable) to connect the serial port on the back of the WX device to the serial port on the terminal. The serial port is of type RS-232 (AT-compatible) with a male, DB-9 connector.

On the terminal, verify the following serial port settings:

- Baud rate: 9600 bps
- Data bits: 8
- Parity: none
- Stop bits: 1
- Flow control: none

### Logging In Using the CLI

---

Enter your user name and password to log in to the CLI. When a WX device is accessed for the first time, use **admin** and **peribit** for the user name and password.

The following prompt is displayed:

*device*#

Where *device* indicates the name of the WX device. To add or change the local user accounts, refer to “configure security” on page 360.

## CLI Basics

Note the following about the CLI:

- CLI commands are case sensitive.
- To view the online help, type “help” before the command or type “?” after it. Type just “help” or “?” at the command prompt to view the available commands or options.
- All configuration changes are made to a staged “candidate” configuration, not the “running” configuration. Changes take effect only when you type “commit”. To retain your changes after the next reboot, type “save-config”.



**NOTE:** If you decide not to commit your changes, you must undo them manually or reboot the WX device to erase them. Otherwise, the next “commit” command or the next update in the Web Console will apply the changes to the running configuration.

- To view all of the settings for the running configuration, type the following:

```
show -run all
```

- To view a specific configuration setting, type the following:

```
show -run <configuration setting>
```

For example, to view the device IP address, subnet mask, and default gateway:

```
show -run ip
```

- To view settings for the candidate configuration, omit “-run” from the show commands.

Table 8 summarizes the keyboard shortcuts.

**Table 8: Keyboard Shortcuts**

Action	Shortcut	Description
Complete commands	Tab or Ctrl + I	Completes a partially typed keyword if enough characters are entered to uniquely identify it.
Recall commands	Ctrl + P or ↑	Retrieves the previous command from the history buffer.
	Ctrl + N or ↓	Retrieves the next command from the history buffer.
Delete characters	Ctrl + D	Deletes the character at the cursor.
	Ctrl + H	Deletes the character before the cursor (same as Delete key).
	Ctrl + K	Deletes all characters from the cursor to the end of the line.
	Ctrl + W	Deletes the word before the cursor.
	Ctrl + U	Deletes all characters on the line.

Action	Shortcut	Description
Move cursor	Ctrl + A	Moves the cursor to the start of the line.
	Ctrl + B	Moves the cursor back one character.
	Ctrl + E	Moves the cursor to the end of the line.
	Ctrl + F	Moves the cursor forward one character.
Transpose characters	Ctrl + T	Transposes the character at the cursor with the preceding character.
Exit configuration mode	Ctrl + Z	Returns to the top level of the CLI.

## Command Modes

When you log in to the CLI, the prompt is the device name followed by “#”, which indicates that you are at the top level of the command hierarchy (also called EXEC mode):

```
device#
```

System level configuration commands, such as “commit” and “save-config” can be entered only at this level.

Configuration commands for specific features can be entered on one line (“config ...”) or they can be entered in stages in configuration mode. For example, to change an interface setting, you can access configuration mode, which changes the prompt to indicate the mode:

```
device# config
device(config)#
```

You can now enter the rest of the command (“interface ...”) or access the interface sub-mode, which again changes the prompt to indicate the mode:

```
device(config)# interface
device(config-interface)#
```

By entering a “?” at each level, you can review the available options and complete the command in stages. To back up one level in the configuration, type **exit**. To return directly to the top level, press **Ctrl + Z**.

## CLI Command Summary

Table 9 provides a summary of the available CLI commands. Note that the “ping” and “traceroute” commands can be entered at any level

**Table 9: CLI Command Summary**

Command	Description
<b>System-Level Commands</b>	
“commit” on page 290	Apply configuration changes to the running configuration.
“configure” on page 290	Enter configuration mode.
“copy” on page 290	Copy files to or from a WX device.
“embed” on page 291	Enable or disable hardware passthrough on a WX 100.
“import-route-table” on page 291	Import a Cisco routing table from an FTP or TFTP server.
“list” on page 292	List the files in a specific directory on the device.
“load-config” on page 292	Load a device configuration from Flash memory, reset to factory default settings, or securely wipe the hard disks on a WXC device.
“packet-capture” on page 293	Capture raw network traffic on the Local or Remote interface.
“ping” on page 294	Verify connectivity with other network devices.
“reboot” on page 295	Reload the WXOS software on the device.
“remove” on page 296	Delete files from the device.
“remove” on page 296	Delete files from the device.
“rollback” on page 296	Revert to an earlier version of WXOS.
“save-config” on page 297	Save the running configuration to Flash memory to preserve committed changes the next time the device is rebooted.
“set” on page 298	Change the device name, location, or administrator contact information.
“shutdown” on page 298	Turn off the device in preparation for disconnecting the power.
“source” on page 299	Execute CLI commands stored in a file.
“support” on page 299	Generate a diagnostic file for Technical Support.
“traceroute” on page 300	Trace the network path to a remote device.
“upgrade” on page 299	Load a new (later) version of WXOS from an FTP or TFTP server.
<b>Configuration Commands</b>	
“configure aaa” on page 300	Define user accounts locally, and specify how users are authenticated.
“configure acceleration” on page 302	Configure the various methods of acceleration WAN traffic.
“configure application” on page 309	Define application definitions, assign them to traffic classes, and enable or disable reduction and acceleration by application.
“configure arp” on page 312	Add static ARP entries or clear dynamic entries.
“configure backup” on page 312	Configure a WX device as a backup for one or more primary devices.
“configure clock” on page 315	Set the device time, time zone, and enable or disable daylight savings time.
“configure console” on page 315	Set the baud rate of the DB9 console port.
“configure dns” on page 316	Specify the DNS servers and domain name used to resolve IP addresses on the Traffic report.
“configure filter” on page 316	Exclude applications or address pairs from data reduction, and specify whether fragmented packets are reduced.

Command	Description
"configure interface" on page 318	Specify interface speeds and duplex modes, run a test to detect a mode mismatch, enable the reduction of VLAN traffic, and reset the interface traffic statistics to zero. You can also enable high-availability support so that when a failure is detected on one interface, the other interface is turned off.
"configure ip" on page 320	Change the device IP address, subnet mask, or default gateway.
"configure ipsec" on page 320	Configure IP security (IPSec) to authenticate and encrypt traffic between WX devices in the same community.
"configure license" on page 324	Enter a new license key.
"configure mon-apps" on page 325	Specify the applications to be monitored on reports.
"configure multi-node" on page 326	Configure two WX devices to work together.
"configure multi-path" on page 326	Configure primary and secondary paths between WX devices.
"configure ospf" on page 330	Use OSPF to discover routes on the LAN side of the WX device.
"configure packet-interception" on page 331	Configure RIP, WCCP, or external routing to route traffic to an off-path device.
"configure prime-time" on page 333	Specify the days and times during the week when network performance is most important.
"configure profile-mode" on page 334	Enable Profile Mode and configure virtual devices for non-intrusive testing.
"configure qos inbound" on page 335	Configure inbound QoS settings.
"configure qos outbound" on page 337	Configure outbound QoS settings.
"configure radius" on page 342	Specify RADIUS servers and server groups used to authenticate users.
"configure reduction" on page 343	Configure data reduction settings.
"configure reduction-subnet" on page 350	Specify the subnets advertised to other WX devices for reduction.
"configure reg-server" on page 352	Configure the WX device that acts as the registration server.
"configure remote-routes" on page 355	Specify how often remote routes are fetched from other devices, and whether each remote route is validated.
"configure rip" on page 356	Use RIP to discover routes on the LAN side of the WX device.
"configure route" on page 356	Add static routes to the routing table, enable router load balancing, and specify the ICMP age-out interval.
"configure route-poll" on page 359	Obtain dynamic routes by periodically polling a Cisco router
"configure security" on page 360	Restrict access by IP address, change the packet capture password, lock the front-panel keypad, or disable the Web console and/or the SSH interface.
"configure snmp" on page 361	Enable or disable SNMP, generate traps, and specify community strings.
"configure sntp" on page 362	Specify primary and secondary SNTP servers to maintain the device time.
"configure stack-group" on page 363	Configure a WX 100 to support client devices.
"configure syslog" on page 364	Send Syslog messages to one or more Syslog server
"configure top-talker" on page 365	Export top traffic statistics to a file and/or send to a NetFlow server.
"configure wan-performance-monitor" on page 366	Configure WAN performance monitoring to measure latency and loss between the current device and one or more remote WX devices.
<b>Show Commands</b>	
"show aaa" on page 368	Display authentication methods for the Console, Web, and CLI.
"show acceleration" on page 368	Display application acceleration configuration.
"show access-log" on page 369	Display management access log.
"show all" on page 369	Display all system configuration information.



Command	Description
“show application” on page 370	Display application definition.
“show arp” on page 370	Display ARP entries.
“show backup-sr” on page 370	Display backup mode configuration.
“show clock” on page 370	Display time related parameters.
“show connection” on page 370	Display list of current reducer connections.
“show console” on page 371	Display console (serial) port parameters.
“show contact” on page 371	Display contact information for the device.
“show dns” on page 371	Display DNS server addresses and default domain name.
“show filter” on page 371	View reduced applications and excluded address pairs.
“show flow-details” on page 372	View details of a specific traffic flow.
“show import-route-table” on page 373	Display import route table information.
“show interface” on page 373	Display network interface parameters.
“show ip” on page 374	Display the IP parameters.
“show ipsec” on page 374	Display IPSec configuration and security associations.
“show license” on page 374	Display license information.
“show location” on page 375	Display location description for this system.
“show log” on page 375	Display system log.
“show mon-apps” on page 375	Display list of monitored applications.
“show multi-node” on page 375	Display whether multi-node is enabled.
“show multi-node-status” on page 375	Display the status of multi-node configuration (master node only).
“show multi-path” on page 375	Display multi-path configuration.
“show ospf” on page 376	Display OSPF settings.
“show packet-capture” on page 376	Display packet capture settings.
“show packet-interception” on page 376	Display packet interception parameters for off-path devices.
“show prime-time” on page 377	Display prime time settings.
“show profile-mode” on page 377	Display Profile Mode settings.
“show qos inbound” on page 377	Display QoS settings.
“show radius” on page 378	Display RADIUS configuration.
“show reduction” on page 378	Display reduction status.
“show reduction-subnet” on page 379	Display reduction subnet status.
“show reg-detail” on page 379	Display detailed information about the registration database.
“show reg-server” on page 380	Display registration server parameters.
“show reg-server” on page 380	Display summary information about the registration database.
“show remote-routes” on page 380	Display remote route information.
“show rip” on page 381	Display RIP parameters.
“show route” on page 381	Display routing table.
“show route-poll” on page 381	Display routing poll table.
“show security” on page 382	Display security related parameters.
“show snmp” on page 382	Display SNMP related parameters.

Command	Description
"show snmp" on page 382	Display SNMP related parameters.
"show stack-group" on page 382	Display model and status of clients connected to an WX 100.
"show syslog" on page 383	Display Syslog parameters.
"show system" on page 383	Display general system information.
"show system-name" on page 383	Display the system name.
"show top-talker" on page 384	Display top-talker parameters.
"show uptime" on page 384	Display system uptime.
"show version" on page 384	Display version information.
"show wan-performance-mon" on page 384	Display WAN performance monitoring configuration.

## System-Level Commands

This section describes the system-level CLI commands. Most of these commands must be entered at the top level of the command hierarchy. Note that the **ping**, **traceroute**, and file commands can be entered at any level.

### **commit**

The **commit** command applies the "candidate" configuration to the "running" configuration. The candidate configuration is a staged configuration that includes all the configuration changes made since the last **commit** command.

To commit the candidate configuration as the running configuration, type:

```
commit
```

### **configure**

The **configure** command is used to access the configuration commands. The command can be entered by itself or followed by specific configuration parameters:

```
config console set baud-rate <number>
```

### **copy**

Use the Copy command to copy files on the WX device or between the device and an FTP or TFTP server.

1. To copy files between directories on the device:

```
copy <source path and file name> <destination path and file name>
```

If you omit the path name, the current directory is assumed.

2. To copy files between the device and a remote location, either the source or destination can be an FTP or TFTP server:

```
copy <full path and file name> ftp://<IP address>[:username:password]/<path and file name>
```

or:

```
copy tftp://<IP address>/<path and file name> <full path and file name>
```

You must specify the full path name on the device. If the user name or password includes a “#”, enclose the entire string in quotation marks. For example:

```
copy /ata0/cfg/startup.cfg “ftp://192.168.0.7:user1:pass#/startup.cfg”
```

## ***embed***

To enable or disable hardware passthrough on a WX 100 (enabled by default):

```
embed bypass-capability <on | off>]
```

Disabling hardware passthrough will block all traffic through the device during a power failure. In high-availability environments, this allows power failures to be detected and the traffic routed to an alternate device. This command is available only on later non-fiber versions of the WX 100 that have the network connectors on the front panel.

## ***import-route-table***

You can import routes from a Cisco router if you first export the routes to a file, and save the file to an FTP or TFTP server. The routes displayed when you enter a “show ip route” command on the Cisco router are added to the local routes on the WX device.

The router must be in the same subnet as the WX device, and it is preferable to use the router connected to the Remote interface. The following types of imported routes are recognized:

**B** - BGP routes, **C** - Connected routes, **D** - EIGRP routes, **E** - EGP derived, **I** - IGRP derived, **i** - IS-IS derived, **O** - OSPF derived, **R** - RIP derived, **S** - Static routes



**NOTE:** You should not import a routing table if dynamic routing is enabled (RIP, OSPF, or route polling). Also, on an off-path device that uses RIP for packet interception, be careful not to import RIP routes that were advertised by the off-path device (traffic to those destinations will be dropped).

1. On the Cisco router, export the routing table and save it to an FTP or TFTP server.
2. To import the routing table from the FTP or TFTP server:

```
import-route-table route-file ftp://<IP address>[:<user>:<pass>]/<path and file name>
```

or:

```
import-route-table route-file tftp://<IP address>/<path and file name>
```

The routing table is stored in the Flash memory and applied to the “candidate” configuration. You are prompted for the FTP user name and password if you omit them from the command line.

3. To delete the last imported route table file:

```
import-route-table delete
```

4. To commit the candidate configuration as the running configuration:

```
commit
```

## **list**

To list the files in a directory:

```
ls [<directory path>]
```

If you omit the directory path, the files in the current directory are listed.

## **load-config**

Configurations can be saved and loaded at any time. The loaded configuration becomes the running configuration. You can also reload the factory default configuration, such as when you must move the device to another location, and securely wipe all data from the hard disks on a WXC device.

To load a configuration using the CLI, the configuration must be stored in the Flash memory. Use the Web console to load a configuration from a local disk or an FTP or TFTP server, as described in “Loading a Device Configuration File” on page 266.



**NOTE:** The configuration file contains information specific to the device, such as IP network settings. Therefore, you cannot load a configuration file from one device to another.

1. As a precaution, save the running configuration to an FTP or TFTP server, as described in “save-config” on page 297
2. To load a device configuration file:

```
load-config <filename | factory-default [-wipe-disk <n>]> [-echo]
[-preserve-ip]
```

Where:

- **filename.** Name of the configuration file (up to 8 characters) without the “.cfg” extension.
- **factory-default.** Reloads the factory settings and restores the temporary license. When the reload is done, unplug the power cable from the back of the WX device, plug the cable back in, and then specify the IP address, subnet mask, and default gateway for the device.



**NOTE:** Restoring the factory default configuration removes all data, configuration information and log files. It also disrupts the reduction tunnels associated with this device. Before you restore the factory default configuration, you should back up the configuration file to another location (refer to “Saving the Device Configuration” on page 263).

- **wipe-disk <n>**. On a WXC device, when you reload the factory defaults you can specify the number of passes used to perform a secure wipe of the hard disks. During each pass, a different value is written to each byte on the disks.

The first pass uses random numbers, the second pass writes a repeated pattern, the third pass uses zeros, the fourth pass writes another repeated pattern, while the fifth pass repeats the sequence with random numbers, shifted by one byte. Each pass takes about three hours (to stop the process, reboot the device). For maximum security, five passes are recommended. To view the progress of a secure wipe, enter the “show reduction” command.

- **-echo**. Displays each command as it is executed.
- **-preserve-ip**. Retains the device IP addresses when you reload the factory defaults.

3. Type “y” to confirm loading the configuration file.
4. To retain a loaded configuration when the device is restarted, save the configuration to *startup.cfg* in Flash memory:  
  
save-config
5. If a loaded configuration file changes the IP address, you MUST save the configuration to *startup.cfg*, and then reboot the device.

## packet-capture

The packet capture feature lets you capture raw network data from the Local and/or Remote interfaces. The packet capture information can then be exported to a file and analyzed by a protocol analyzer program or other hardware. The format of the captured file is either “libpcap” or “snoop”. Packet captures are logged in the Access Log file.

You are prompted for the packet-capture password when you start or copy a packet capture. The default password is “peribit”. To change the password, refer to “configure security” on page 360.



**NOTE:** If tunnel switching is enabled, running a packet capture will capture intermediate assembled packets before they are recompressed. These packets have zeros for the source and destination, and may have checksum errors. These packets are internal to the device and can be ignored.

1. To start the Packet Capture:

```
packet-capture start interface <local | remote | both> size <number> [packets
<number>] [format <libpcap | snoop>] [snaplen <max size>] [savetime <time>]
```

Where:

- **interface <local | remote | both>**. Indicates the interfaces where data is collected.

- **size <number>** . Number of bytes to capture (4096 is the minimum).
  - **packets <number>** . Maximum number of packets to capture.
  - **format <libpcap | snoop>** . File format of the collected data. The default is `libpcap`.
  - **snaplen <max size>** . Maximum number of bytes captured for each packet (0 to 65535). The default is 1514. A zero captures the entire packet.
  - **savetime <time>** . Number of seconds that a completed packet capture is available in memory. The default is 3600.
2. To stop a packet capture:
 

```
packet-capture stop
```
  3. To copy a packet capture to an FTP or TFTP server:
 

```
packet-capture copy ftp://<IP address>[:<user>:<pass>]/<path and file name>
[startpkt <number>] [numpkts <number>]
```

or:

```
packet-capture copy tftp://<IP address>/<path and file name> [startpkt
<number>] [numpkts <number>]
```

Where:

    - **startpkt <number>** . Starting packet number. The default is "0".
    - **numpkts <number>** . Number of packets to copy in addition to the start packet. The default is zero, which copies all packets.
  4. To delete the packet capture data:
 

```
packet-capture delete
```

## ***ping***

You can use the `ping` command to verify connections to other WX devices, or other devices in your network. To ping a WX device or other network device:

```
ping <IP address>
```

```
PING 192.168.5.150 (192.168.5.150): 56 data bytes
64 bytes from 192.168.5.150: icmp_seq=0. time=16. ms
64 bytes from 192.168.5.150: icmp_seq=1. time=4. ms
64 bytes from 192.168.5.150: icmp_seq=2. time=2. ms
64 bytes from 192.168.5.150: icmp_seq=3. time=4. ms
64 bytes from 192.168.5.150: icmp_seq=4. time=4. ms
```

```
—192.168.5.150 PING Statistics—
5 packets transmitted, 5 packets received, 0% packet loss
round-trip (ms) min/avg/max = 2/6/16
```

## reboot

If you load a new boot image of the system software, you must reboot the device. Rebooting the device loads the configuration information in the “startup.cfg” file, along with the current boot image. If you can reboot the device in Safe Mode, the power is on, and the device can be configured, but all traffic is passed through without reduction.

1. To immediately reboot the system:

```
reboot [-all | -client-id <0-6>] | [[-safe-mode] [-no-resync]]
```

Where:

- **-all.** Reboots an WX 100 server and all of its client devices. Does not support the safe-mode option.
- **-client-id <0-6> .** Reboots all of an WX 100’s clients (0) or just a specific client (1-6).
- **-safe-mode.** Reboots the device so that the power is on, and the device can be configured, but all traffic is passed through without reduction. Note that the warning “ER SW Passthru” is displayed in the front panel of devices that have an LCD.
- **-no-resync.** Reboots the device and clears the reduction dictionary used for Network Sequence Monitoring. Available only on WXC devices.

Type “y” to confirm that you want to reboot.

2. If you issue the reboot command from a terminal connected to the serial port, the following options are displayed after the reboot. You can enter a number (1 to 4) to specify the boot image to be loaded.
  - **1. Primary image.** The most recent image loaded on the device (*srs.os*). If you do not select another boot image within a few seconds, the primary image is loaded.
  - **2. Secondary image.** The previous image loaded on the device (*srs1.os*). If you have not upgraded the software, the primary and secondary boot images are the same.
  - **3. Recovery image.** The image loaded in the read-only area of Flash memory (WX 15 only). Load the recovery image only if you suspect that the read/write area of Flash memory has been corrupted (it is loaded automatically if the primary and secondary images are not found). After the recovery image is loaded, you must:
    - a. Enter the IP address, subnet mask, and default gateway for the device. Skip all other Quick Setup prompts.
    - b. Use the “upgrade” command to load a full boot image on the device (refer to “upgrade” on page 299). Do NOT use the recovery image for normal operation.
    - c. Reconfigure the device (the device configuration is reset to the factory defaults)

- **4. Specify boot image name.** If you have copied the primary or secondary image to another name on the device (for example, “copy srs.os test.os”), you can specify the name of the image to be loaded. Intended primarily for testing purposes.

## **remove**

To delete a file on the device:

```
rm <path and file name>
```

If you omit the path name, the current directory is assumed.

## **rollback**

The `rollback` command is used to install a previous version of the WXOS operating system. To rollback to a previous version of WXOS, you must have the “.bin” or “.zip” file installed on an FTP or TFTP server in your network. You can also enter “rollback” without parameters to undo the outbound QoS settings in the candidate configuration.

Due to hardware changes, some WX devices require a minimum version of WXOS. Before rolling back a WX 100 server to a previous release, check the following table. If you roll back a WX 100 server to a release that a client does not support, the client and server will have different versions of WXOS (the client is not rolled back).

To view the WX hardware version or serial number, use the “show system” command or click Help > About in the Web console (the hardware version is next to the model number).

Model	Minimum WXOS 5.x Version
WX 15	5.0.2.20h
WX 20	Hardware v1.0 (below S/N 0020007100)—Any Hardware v2.0—5.0.11.0
WX 50	Below S/N 0050003660—Any S/N 0050003660 to 0050004620—5.0.6.0 Above S/N 0050004620—5.0.12.0
WX 55	Any
WX 60	5.0.10.0
WX-80	Any
WX 100 (data ports on front panel)	5.1.7
WX 100 (data ports on back panel)	Standalone—5.0.0.56 Stack—Minimum version required by clients. For example, to attach a WX 60 client, first upgrade the WX 100 to WXOS 5.0.10.0 or later.
WXC 250	5.0.8.0g
WXC 500	Below S/N 0500002000—5.0.2.21g S/N 0500002000 or higher—5.0.10.0g



Before you rollback to a previous version of WXOS, note the following:

- WXOS 5.3 is supported only on WX 100 servers in stack mode— it is not supported on any WX device in standalone mode. If you install WXOS 5.3, and then want to change the WX 100 server to standalone mode (no clients), you must roll back to the previous release.
- If you rollback to a previous version of WXOS, you will lose the features introduced in the later version (the associated settings in the configuration file are ignored).



**NOTE:** If you roll back WXOS 5.0 to a previous version, all configuration data (except the IP information) will be reset to factory defaults. Always save the configuration file before upgrading to a new release so that you can reload the configuration after a roll back.

1. To rollback to a previous version of WXOS:

```
rollback ftp://<IP address>[:<user>:<pass>]/<path and name of the WXOS file>
```

or:

```
rollback tftp://<IP address>/<path and name of the WXOS file>
```

2. Type “y” to confirm the rollback.
3. Reboot the device to activate the new software.

## save-config

After you commit configuration changes, you must save the configuration if you want to preserve the settings the next time the device is rebooted. When you save the configuration through the CLI, it is stored in the Flash memory. Use the Web console to save the configuration to a local disk or an external FTP or TFTP server, as described in “Saving the Device Configuration” on page 263.



**NOTE:** A configuration file contains information specific to the device, such as IP network settings. Therefore, you cannot load the configuration file from one device to another.

1. To view the current configuration:

```
show all
```

2. To save the configuration with the default name:

```
save-config
```

The configuration file is saved as *startup.cfg* and is used when you reboot the device.

3. To save the configuration with another name:

```
save-config <file name>
```

The name can be up to 8 characters. Do not include a file name extension (such as ".txt").

4. Type "y" to confirm saving the configuration file.

## set

The **set** command lets you specify the device name, location, and an administrator's contact information. Text that includes spaces must be enclosed in double quotation marks.

1. To view the current system settings:

```
show -run system
```

2. To specify a device name (up to 30 characters):

```
set system-name <device name>
```

Do not use colons (:), asterisks (\*) question marks (?) or angle brackets (< >) in device names. Device name changes are propagated to the other WX devices in the community the next time the device checks in with the registration server for updates.

3. To set an administrator contact information:

```
set contact <contact name, phone, etc.>
```

4. To set a location for the device:

```
set location <location>
```

## shutdown

Run the **shutdown** command before removing the power cord from a WX device.

1. To shut down the device:

```
shutdown [-all | -client-id <0-6>] | [[-reset-all] [-reset-monitor] [-reset-reg] [-reset-log] [-reset-access-log]]
```

Where:

- **-all**. Shuts down an WX 100 server and all of its client devices. Does not support reset options (everything is reset).
- **-client-id <0-6>**. Shuts down all of an WX 100's clients (0) or just a specific client (1-6).
- **-reset-all**. Resets everything (the default).
- **-reset-monitor**. Resets the monitoring statistics.
- **-reset-reg**. Resets the information from the registration server.
- **-reset-log**. Deletes the system log files.
- **-reset-access-log**. Deletes the access log files.

2. Type “y” to confirm the shutdown.

## source

The **source** command executes a file of configuration commands as if they were typed interactively. To execute a file of configuration commands:

```
source [-echo] <file path and name>
```

The file name can be up to 8 characters. The full path name is required, but the file name extension is optional. The **-echo** option displays each command as it is executed.

## support

You can create a diagnostic file containing the current configuration, system information, filter settings, and log files. You can then email this file to Technical Support to assist in the diagnosis of problems. The CLI command sends the diagnostic file to an FTP or TFTP server. Use the Web console to save the diagnostic file to a local disk, as described in “Creating a Diagnostic File” on page 280.

1. To create a diagnostic file and copy it to an FTP or TFTP server:

```
support export <label> ftp://<IP address>[:<username>:<password>]/<path and  
file name>
```

or:

```
support export <label> tftp://<IP address>/<path and file name>
```

2. Press Enter.
3. Type a description for the file and press Enter.
4. Type “.” on a line by itself and press Enter.
5. When the command prompt returns, the file was successfully created and sent to the TFTP or FTP server. Make a copy of the file and send it to *support@juniper.net*.

## upgrade

To upgrade the system software to a later version, you can load a new boot image of the WXOS operating system from a TFTP or FTP server. Upgrading the system software does not affect the configuration information stored in the *startup.cfg* file. All configuration information is preserved.



**NOTE:** Your monitoring statistics may be corrupted if you use this command to install a previous version of the software. Always use the “rollback” command to restore a previous version of WXOS (refer to “rollback” on page 296).

1. To upgrade system software from an FTP or TFTP server:

```
upgrade ftp://<IP address>[:username:password]/<path and file name>
```

or:

```
upgrade tftp://<IP address>/<path and file name>
```

2. Type “y” to confirm upgrading the system software.
3. Reboot the device to activate the new software.

## ***traceroute***

You can use the trace route utility to determine the number of router hops and the route taken from the current WX device to another device in your network. This tool can help you determine the point in your network that is causing a connection failure.

To run a trace route to a WX device or other network device:

```
traceroute <IP address>
```

The trace route results are displayed in the CLI.

## **Configuration Commands**

---

### ***configure aaa***

The **aaa** command is used to define up to 25 users, and to specify how users are authenticated to access the device through the Web, SSH (CLI), and the console. You can also enable or disable authorization checking. To define the RADIUS servers and server groups, refer to “configure radius” on page 342.

1. To view the current AAA settings:

```
show -run aaa
```

2. Type the following command to enter the configure AAA mode:

```
config aaa
```

3. To add a user account that can be authenticated locally:

```
user add name <name> [idle-timeout <seconds>] [privilege-level <read-only | read-write>]
```

Where:

- **name <name>** . User name (up to 32 characters). If the name includes spaces, enclose the name in quotation marks.
- **idle-timeout <seconds>** . Number of seconds before an idle user is logged out (the default is 1800). A zero indicates the user is never logged out. The timeout setting is ignored if authorization checking is disabled (see Step 5).

- **privilege-level.** Indicate whether the user has read-only or read-write access (the default is read-write). The read-only setting is ignored if authorization checking is disabled (see Step 5).

and then press Enter. Type the new password (from 4 to 64 characters) and press Enter, and then repeat to verify.

To change a user's password:

```
user set name <name> password
```

and then press Enter. Type the new password and press Enter, and then repeat to verify.

To change a user's idle timeout and/or access level:

```
user set name <name> idle-timeout <seconds> privilege-level <read-only | read-write>
```

To delete a user account:

```
user remove <name>
```

4. To specify up to four authentication methods for each management interface:

```
authentication set [console | ssh | web] [local | none | <server group>]
```

Where:

- **console | ssh | web.** Indicates the management interface (**ssh** is for Secure Shell access to the CLI).
- **local | none | <server group> .** Indicates whether users are authenticated locally by the device (the default), by a group of one or more RADIUS servers, or not at all. The **none** option is valid only for the console interface, and it can be used alone or after the last RADIUS group, but it cannot be used directly after **local**.

You can specify up to four methods (separated by spaces), which are then 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 in the local database.

In the following command, if a user is not in the local database, the RADIUS servers in *group1* are tried in sequence. If none of them responds, the servers in *group2* are tried, and so on. If all of the RADIUS servers are down or do not respond, access is denied.

```
authentication aaa web local group1 group2 group3
```

5. By default, authorization checking is disabled, 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 set authorization checking to "same-as-authentication".

```
authorization set mode {off | same-as-authentication}
```

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.

6. To specify the number of unsuccessful login attempts allowed on the SSH interface (1-10 or unlimited) before the user is disconnected (the default is three):

```
set login-retries {<number> | unlimited}
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure acceleration**

The **acceleration** command is used to configure the various acceleration methods. Each method is most effective in high-latency environments, as described below. For more information about acceleration, refer to “Accelerating WAN Traffic” on page 189.

- **Fast Connection Setup (FCS).** The sending device locally acknowledges session requests for destinations known to be active. Intended for applications that have many short sessions, such as HTTP 1.0 and NetBios.
- **Active Flow Pipelining (AFP).** The sending and receiving WX devices terminate the TCP session and acknowledge all data transmissions locally. This results in three independent sessions—between the source and the sending device, between the two WX devices, and between the receiving device and the destination. Intended for applications that do a large volume of data transfers over satellite connections or other high-latency environments. On a given path between two devices, AFP cannot be used simultaneously with Fast Connection Setup.
- **Forward Error Correction (FEC).** The sending WX device sends recovery packets with the data so that the receiving device can reconstruct lost packets without requesting a retransmission. You can specify the number of recovery packets per block of data packets. Intended for use in high-loss, high-latency environments, such as satellite connections.
- **Application Flow Acceleration (CIFS, Exchange, and HTTP).** The sending WX device locally acknowledges the multiple requests needed to read or write large files, and provides HTTP caching and pre-fetch for static Web objects (.css, .gif, .jpeg, and .js). CIFS, Exchange, and HTTP acceleration require Active Flow Pipelining to be enabled.

To accelerate traffic between two WX devices, the following conditions must be met:

- The selected applications must be reduced, and a reduction tunnel must exist in both directions between the WX devices.
- Outbound QoS must be enabled and the WAN circuit speed must be specified for each remote WX device for which you want to accelerate traffic (refer to “configure qos outbound” on page 337).

- To use AFP or Application Flow Acceleration, you must enable the related topology features (refer to “configure reduction” on page 343). Enable clustering for AFP if the outbound and return traffic does not always traverse the same two WX devices.



**NOTE:** PFA is most effective in networks with high-speed connections and high latency. However, PFA may have no effect if traffic must traverse high-latency or low-speed connections that are one or more hops beyond the receiving WX device.

1. To view the current acceleration configuration and status, use the following commands:

```
show -run acceleration application <cifs | exchange | http> <configuration | status>
```

```
show -run acceleration cluster <configuration | status>>
```

```
show -run acceleration packet-flow <configuration | status>
```

Where:

- **application <cifs | exchange | http> <configuration | status>** . Lists the configuration or status for CIFS, Exchange, or HTTP application acceleration.
  - The **CIFS** status shows the current number of active flows, passive flows, and number of files being tracked, along with several totals since the device was last reset, such as the total number of CIFS flows, the total reads and writes, and the number of reads and writes accelerated. Most active flows are accelerated; passive flows and flows for unsupported clients or servers are not. For example:
 

```
Active flows: 1
Passive flows: 6
Flows from unsupported clients: 2
Flows to unsupported servers: 2
Total flows: 32
Files currently tracked: 0
Accelerated writes: 0
Total writes: 0
Accelerated reads: 0
Total reads: 2
```
  - The **Exchange** status shows the current number of active flows, and several totals since the device was last reset: the Packet Data Units (PDUs) reduced (cc) and decompressed (dc), the number of read, write, and “other” operations (“r/w/o”), and the number of reads and writes accelerated (“other” operations cannot be accelerated). For example:
 

```
Flows: Active:    5
PDUs : cc/dc:    367/456
      r/w/o:    723/543/342
Accel: Total:    612 (392 reads, 220 writes)
```
  - The **HTTP** status shows the current cache usage for pre-fetched objects (items), cached static objects (datablocks), cookies, HTTP servers (hosts), and URLs (host-paths). For example:

```
***** Database Usage *****
      Total  Used
Items:      4096   0
Data Blocks: 8192   0
Cookies:     384   0
Hosts:       512   0
Host Paths: 16384   0
```

- **cluster** < **configuration** | **status** > . Lists the other WX devices in the same cluster (if any) or the last heartbeat sent and received by each device in the cluster. Clusters of devices can be defined for AFP if the outbound and return traffic does not always traverse the same two WX devices (asymmetric routing support).
- **packet-flow** < **configuration** | **status** > . Lists the global configuration settings for FCS, AFP, and FEC, or the configuration status for each remote endpoint.

2. Type the following command to enter the configure acceleration mode:

```
config acceleration
```



3. The following table describes the acceleration settings.

Setting	Commands
Packet Flow Acceleration methods	<p>To enable or disable the PFA methods to be used for one or more endpoints (all methods are disabled by default):</p> <pre>active-flow-pipelining set mode &lt;on   off&gt; fast-connection-setup set mode &lt;on   off&gt; forward-error-correction set mode &lt;on   off&gt;</pre> <p>To clear the counters shown on the Forward Error Correction report:</p> <pre>forward-error-correction clear-counters</pre>
Endpoints	<p>You can enable the same PFA methods for all endpoints or add each endpoint and specify its PFA methods (all methods are disabled by default).</p> <h3>Enabling Specific Endpoints</h3> <p>To add an endpoint, a reduction tunnel must exist for the device and outbound QoS must be configured correctly:</p> <pre>endpoint add ip-address &lt;IP address&gt; [mode &lt;on   off&gt;] [active-flow-pipelining &lt;on   off&gt;]   {[fast-connection-setup &lt;on   off&gt;] [forward-error-correction &lt;on   off&gt;] [data-pkts &lt;4-25&gt;] [recovery-pkts &lt;0-5&gt;]]}</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>■ <b>mode &lt;on   off&gt;</b> . Enables or disables the endpoint for PFA (disabled by default).</li> <li>■ <b>&lt;method&gt; &lt;on   off&gt;</b> . Enables or disables a PFA method. Active Flow Pipelining cannot be used with Fast Connection Setup.</li> <li>■ <b>data-pkts &lt;4-25&gt; recovery-pkts &lt;0-5&gt;</b> . For Forward Error Correction, one recovery packet is sent for every nine data packets. Increasing the ratio of recovery packets to data packets reduces retransmissions, but requires more overhead. Zero recovery packets disables error correction.</li> </ul> <p>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.</p> <p>To change the PFA settings for an endpoint:</p> <pre>endpoint set ip-address &lt;IP address&gt; [mode &lt;on   off&gt;] [active-flow-pipelining &lt;on   off&gt;]   {[fast-connection-setup &lt;on   off&gt;] [forward-error-correction &lt;on   off&gt;] [data-pkts &lt;4-25&gt;] [recovery-pkts &lt;0-5&gt;]]}</pre> <p>To remove an endpoint from the list shown by the “show acceleration” command:</p> <pre>endpoint remove &lt;IP address&gt;</pre> <p>In the Web console, the endpoint is disabled, but is not deleted.</p>

### Enabling All Endpoints

To use the same PFA methods for all endpoints, enable all endpoints (disabled by default):

```
set enable-all-endpoints on
```

To set or change the methods that apply to all endpoints (the “default” IP address indicates all endpoints):

```
endpoint set ip-address default [active-flow-pipelining <on | off>] | {[fast-connection-setup <on | off>] [forward-error-correction <on | off>] [data-pkts <4-25>] [recovery-pkts <0-5>]]}
```

Traffic is accelerated to all remote WX devices for which a reduction tunnel exists and outbound QoS is configured correctly. The specified PFA methods are applied to all qualifying endpoints, and to all qualifying endpoints added to the same community in the future.


Setting	Commands
PFA Applications	<p>To specify the applications that use Fast Connection Setup:  <b>fast-connection-setup application add &lt;name&gt;</b></p> <p>To add an application that is included or excluded from Active Flow Pipelining:  <b>active-flow-pipelining application add &lt;name&gt;</b></p> <p>To indicate whether the specified applications are included or excluded from Active Flow Pipelining (included by default):  <b>active-flow-pipelining application mode {include   exclude}</b></p> <p>Note that “include” mode excludes traffic for undefined applications.</p> <p>To remove an application from the list shown by the “show acceleration” command:  <b>active-flow-pipelining application remove &lt;name&gt;</b>  <b>fast-connection-setup application remove &lt;name&gt;</b></p> <p>An application removed from the “include” list is shown as disabled in the Web console, but it is not deleted.</p>
CIFS Acceleration	<p>To enable or disable CIFS acceleration (disabled by default):  <b>cifs set mode &lt;on   off&gt;</b></p> <p>To enable or disable CIFS acceleration between clients and servers that have SMB signing enabled, but not required (disabled by default):  <b>cifs set disable-signing &lt;on   off&gt;</b></p> <p>To accelerate a CIFS application (must be an application of type CIFS):  <b>cifs application add &lt;name&gt;</b></p> <p>Note that acceleration is enabled for all remote endpoints for which AFP is enabled.</p> <p>To disable acceleration for a CIFS application:  <b>cifs application remove &lt;name&gt;</b></p> <p>For more information about CIFS acceleration, refer to “Microsoft CIFS and Microsoft Exchange Acceleration” on page 201.</p>
Exchange Acceleration	<p>To enable or disable Exchange acceleration (disabled by default):  <b>exchange set mode &lt;on   off&gt;</b></p> <p>To accelerate an Exchange application (must be an application of type Exchange):  <b>exchange application add &lt;name&gt;</b></p> <p>Note that acceleration is enabled for all remote endpoints for which AFP is enabled.</p> <p>To disable acceleration for a Exchange application:  <b>exchange application remove &lt;name&gt;</b></p> <p>For more information about Exchange acceleration, refer to “Microsoft CIFS and Microsoft Exchange Acceleration” on page 201.</p>

Setting	Commands
HTTP Acceleration	<p>To enable or disable HTTP acceleration (disabled by default):</p> <pre>http set mode &lt;on   off&gt;</pre> <p>To accelerate an HTTP application (the application must have an application type of HTTP):</p> <pre>http application add &lt;name&gt;</pre> <p>Note that acceleration is enabled for all remote endpoints for which AFP is enabled.</p> <p>To disable acceleration for an HTTP application:</p> <pre>http application remove &lt;name&gt;</pre> <p>To enable or disable pre-fetch for HTTP acceleration (enabled by default):</p> <pre>http set &lt;name&gt; pre-fetch &lt;on   off&gt;</pre> <p>WXC devices support “header-and-body” caching, which stores each page’s static objects (“.css”, “.gif”, “.jpeg”, and “.js”) in the WXC cache (enabled by default). WX devices support “header-only” caching. To specify the type of caching:</p> <pre>http set data-types &lt;header-only   header-and-body&gt;</pre> <p>To enable or disable caching for both headers and static objects (enabled by default):</p> <pre>http set cache &lt;on   off&gt;</pre> <p>For more information about HTTP acceleration, refer to “HTTP Acceleration” on page 202.</p>

AFP clusters (asymmetric routing support)

For AFP to accelerate a traffic flow, the traffic 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 WX 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 WX device.

In the following example, if D3 receives a SYN packet from D1, the SYN and its source are advertised to D4. If D4 receives the SYN ACK, it can establish the flow with D1.



Note the following:

- Load balancing on the router or switch must be flow- or destination based (not packet-based)
- If you have a cluster on both sides of the WAN, reduction tunnels must be enabled in both directions between all the WX devices in the two clusters.
- A device in a cluster can accelerate traffic only to remote devices that have WXOS 5.1 or later.
- If Multi-Path is enabled on one peer, it must be enabled for all devices in the cluster. Also, traffic is accelerated only if the same path is used in both directions (primary or secondary).
- To ensure that traffic flows are accelerated, asymmetric routing support takes precedence over preferred assemblers and tunnel load balancing settings defined on the WX device.

To specify up to three cluster peers for the current device, enter the IP address of each of the other WX devices in the cluster (multiple addresses must be separated by spaces):

```
cluster set <IP-addresses> | none
```

Setting	Commands
	<p>Note that these are device IP addresses, not virtual addresses. Specify “none” to disable clustering without removing the peer addresses.</p> <p>To add one or more peers to the cluster (multiple addresses must be separated by spaces):</p> <pre>cluster add &lt;IP-address1 IP-address2 ...&gt;</pre> <p>To remove one or more devices from the cluster:</p> <pre>cluster remove &lt;IP-address1 IP-address2 ...&gt;</pre>
AFP buffers	<p>For optimum performance of Active Flow Pipelining, you can adjust the size of the buffer used to receive traffic. For example, if most of the traffic is sent from a hub to the spokes, you may want to adjust the buffer size on the spoke devices (default is “medium”).</p> <pre>active-flow-pipelining set buffer-size {small   medium   large   huge}</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>■ <b>small.</b> Recommended for links with speeds of 512 Kbps (or higher) and round-trip-times under 150 ms, and for links under 512 Kbps and RTTs up to 300 ms.</li> <li>■ <b>medium.</b> Recommended for links with speeds of 512 Kbps (or higher) and round-trip-times from 150 to 600 ms, and for links under 512 Kbps and RTTs above 300 ms.</li> <li>■ <b>large.</b> Recommended for links with speeds of 512 Kbps (or higher) and round-trip-times from 600 to 1200 ms. Not recommended for slower links.</li> <li>■ <b>huge.</b> Recommended for links with speeds of 512 Kbps (or higher) and round-trip-times above 1200 ms. Not recommended for slower links.</li> </ul> <p><b>NOTE:</b> Larger buffer sizes use more memory, but smaller buffer sizes may restrict throughput. The recommended buffer sizes allow up to about 100 Mbps of unreduced throughput.</p>
Heartbeats	<p>On a high-loss link, reduction may be terminated if heartbeat packets are lost. By default, when AFP or FEC is enabled for a remote endpoint, the local device stops reducing data for the remote device if it fails to respond to 15 consecutive heartbeats (passthrough mode). If 30 consecutive heartbeats get no response, the reduction tunnel to the remote device is disabled. The local device attempts to reestablish the tunnel every three minutes.</p> <p>To increase the number of consecutive missed heartbeats that stop data reduction and disconnect the tunnel:</p> <pre>set heartbeat-misses passthru &lt;number   default&gt; disconnect &lt;number   default&gt;</pre> <p>This setting applies only to remote endpoints for which AFP or FEC are enabled. To change the heartbeat settings for all other endpoints, refer to “configure reduction” on page 343.</p>
MSS override	<p>In some cases, the maximum segment size (MSS) negotiated by TCP may be too high for some environments (such as a VPN). To specify a maximum MSS value for AFP sessions (“default” is 1460):</p> <pre>active-flow-pipelining set mss-override &lt;64-1460   default&gt;</pre> <p>This value is used only if the negotiated value is higher.</p>

- To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure application**

The Application command is used to manage application definitions. Application definitions allow WX devices to identify the traffic of up to 256 applications (the WX 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”.

For each application you define, you can:

- Enable or disable data reduction. To conserve system capacity, you should disable reduction for applications whose traffic is encrypted or already compressed.
- Enable TCP acceleration (if data reduction is enabled).
- Assign the application to a traffic class. Traffic classes are used by outbound QoS to allocate WAN bandwidth, and by Multi-Path to direct traffic to a specific network path.
- View application reduction and acceleration statistics.

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 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.

For a list of the default application definitions, refer to Table 1 on page 89.

1. To view the current definitions for one or all applications:

```
show -run application [name]
```

2. Type the following command to enter the configure application mode:

```
config application
```

3. To add an application definition:

```
add name <name> [type <default | cifs | citrix | exchange | ftp | http>] [precedence <number>]
```

Where:

- **name <name>** . Application definition name. If the name includes spaces, enclose the name in quotation marks.
- **type <type>** . Application type. Specify one of the following.
  - **Default**. No special processing (default).
  - **cifs**. To use CIFS application acceleration (refer to “configure acceleration” on page 302), apply to the CIFS application and each application that uses CIFS.

- **citrix**. To add a Citrix client or application name for pattern matching, apply to the ICA application.
  - **exchange**. To use Exchange application acceleration, apply to the Exchange application and each application that uses Exchange. Also allows Exchange ports to be learned dynamically.
  - **ftp**. Apply to the FTP application to allow FTP ports to be learned dynamically.
  - **http**. To use HTTP application acceleration, apply to the HTTP application and each application that uses HTTP. Also allows a URL to be specified for pattern matching.
  - **precedence < number >** . Packets are compared against the definitions in ascending order by precedence number. The comparison stops on the first match, so if two definitions are similar, the more specific definition must have a lower precedence number. By default, a new definition receives the next highest precedence number. If you specify a lower value, the existing definitions are renumbered (you cannot exceed the current highest precedence number).
4. To add a rule to an application definition (omitting an option allows a match to occur on any value):

```
rule add name <name> [src-port <number>] [dst-port <number>] [proto <string>]
[src-addr <IP address>[/<mask>]] [dst-addr <IP address> [/<mask>]]
[dscp <number>] [url <string>] [citrix-app <name>] [citrix-client <name>]
[ip-precedence <number>]
```

Where:

- **name < name >** . Application definition name. If the name includes spaces, enclose the name in quotation marks.
- **src-port < number >** . Source port number, a range of port numbers separated by a hyphen (-), or a series of comma-separated port numbers and ranges. For a list of common application port numbers, refer to Appendix , “Common Application Port Numbers”.
- **dst-port < number >** . Indicates the destination port (same format as the source port). Typically, source and destination ports are specified in separate rules so that a match occurs on packets that specify either port. A rule that includes both ports will match only packets that specify both ports.
- **proto < string >** . Indicates the protocol is “tcp”, “udp”, or a protocol number (0 to 134). To match on a URL or Citrix name, the protocol must be TCP. By default, a match can occur on any TCP or UDP packet. Any protocol you define by number becomes a default (like TCP and UDP) that applies to any rule that does not specify a protocol.
- **src-addr < IP address > [/mask]**. Source IP address or subnet.

- **dst-addr <IP address> [/mask]**. Destination IP address or subnet. Typically, source and destination addresses are specified in separate rules so that a match occurs on packets that specify either address. A rule that includes both addresses will match only packets that specify both addresses.
- **dscp <number>**. Differentiated Services Code Point (DSCP) value (0 to 63).
- **url <string>**. A URL of up to 127 characters (application type must be HTTP). The general format is:

<host>/<uri>

Where:

**<host>** is up to eight strings separated by periods. An asterisk (\*) by itself indicates any string. For example:

www.peribit\*.com

**<uri>** is up to eight strings separated by slashes. An asterisk (\*) by itself indicates any string. For example:

www.peribit\*.com/\*/index.htm

Note that an asterisk is treated as a single character when it is part of another string, such as “www.peribit\*.com”.

- **citrix-app <name>**. Name of a Citrix application (ICA application definition only).
- **citrix-client <name>**. Name of a Citrix client (ICA application definition only).
- **ip-precedence <number>**. ToS IP precedence value (0 to 7).

5. To change an application's precedence number or type:

```
modify name <name> [precedence <number>] [type <default | cifs | citrix |
exchange | ftp | http>]
```

To change an application rule, specify the application name and the rule ID of the rule you want to change (1 to 10):

```
rule modify name <name> rule-id <1-10> [src-port <number>]
[dst-port <number>] [proto <string>] [src-addr <IP address>[/<mask>]]
[dst-addr <IP address>[/<mask>]] [dscp <number>] [url <string>]
[citrix-app <name>] [citrix-client <name>] [ip-precedence <number>]
```

To delete a value from an application rule, specify a “-” for the value. The following example deletes the protocol so that a match can occur on any protocol:

```
rule modify name <application name> rule-id <1-10> proto -
```

To delete an entire rule:

```
rule remove name <application name> rule-id <1-10>
```

6. To delete an application definition:

```
remove <application name>
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure arp**

The ARP command is used to communicate with devices that do not respond to Address Resolution Protocol (ARP) requests. Using the ARP command, you can configure static ARP entries that map the IP addresses of those devices to their MAC addresses.

1. To view a list of static and dynamic ARP entries:

```
show -run arp
```

2. To add a new static ARP entry, type

```
config arp add <IP address> <ethernet address> <local | remote>
```

Where <IP address> is the IP address, <ethernet address> is the MAC address (the format is xx:xx:xx:xx:xx:xx), and <local | remote> indicates the device's Local or Remote interface.

To delete a static ARP entry:

```
config arp remove <IP address>
```

3. To clear all dynamic ARP entries:

```
config arp flush
```

4. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure backup**

The Backup command allows a WX device to serve as a backup for up to 10 primary (active) WX devices. The primary devices can reside in different communities, provided that the backup device belongs to each community. Each hour, the backup device downloads the configuration of each primary device using SSL on TCP port 3577.

The backup and primary devices exchange UDP heartbeats every five seconds on port 3578. If 12 heartbeats are missed, the backup device is activated if the primary's configuration was received at least once. An activated backup continues to send heartbeats to the failed primary, and returns to standby mode when the primary recovers.




The backup feature supports both in-line and off-path deployments, but the backup and primary devices must be the same type (an off-path primary cannot have an in-line backup).

Note the following when configuring a WX device as a backup:

- Do not use a WX 15 as a backup device.
- In general, a backup device must be in the same data path as the primary device so that the backup can advertise the same reduction subnets as the primary. It is recommended that the backup and primary devices be on the same subnet. If a primary device uses static routes, the backup device **MUST** be on the same subnet.

Figure 167 shows a deployment where the backup can advertise the same reduction subnets as the primary, even though it is not in the same data path.

**Figure 167: Sample Backup Deployment**



- To back up an off-path device, the reduction subnets must be static routes (not learned through RIP or OSPF), and the off-path backup device must be on the same subnet.
- The primary and backup devices must all have the same versions of WXOS. Also, it is highly recommended that the registration server has WXOS 5.0.4 or later. If a primary device uses Multi-Path, the primary and backup devices must have WXOS 5.0.4 or later.
- To back up a registration server, the backup device must be the secondary registration server and must be on the same subnet.
- Do not manually configure IPsec on the backup device.



**NOTE:** When a backup is active, do not change the community or save the configuration. Changes to the backup configuration are NOT exported to the primary devices.

To configure a backup device:

1. To install a new WX device as a backup, do not obtain a permanent license. The temporary 30-day license supports all features at the maximum device speed. Only the time the device is active is counted against the 30-day limit (WXOS 5.1 or later required).

To convert an active device to a backup, reload the factory default configuration to restore the temporary license and erase potential configuration conflicts with the primary devices (such as IPsec passwords):

**load-config factory-default**

When the reload is done, unplug the power cable from the back of the device, and plug the cable back in.

2. On the registration server, assign the backup device to each community that contains primary devices to be backed up. The backup device must NOT belong to the Default community (inactive devices are purged from the Default community in 24 hours). If a primary device belongs only to the Default community, assign it to another community.
3. On each remote WX device, if reduction, acceleration, QoS, IPsec, or Multi-Path is enabled for a primary device, verify that the same feature is enabled for the backup device. Note that if reduction is enabled for "ALL discovered Peribit devices", then an outbound reduction tunnel to the backup device is formed automatically when the backup becomes active.
4. Disable load balancing on all devices that reduce data for a primary device.
5. To enable or disable backup mode on the backup device (disabled by default):

**config backup-sr set mode <on | off>**



**NOTE:** If you later add a new WX device to the community, you can disable backup mode on the backup device, verify that each feature is configured correctly from the new device to the backup device, and then re-enable backup mode. For reduction, this step is unnecessary if the new device enables reduction for "ALL" devices.

6. To specify a primary device supported by this backup device:

**config backup-sr remote-sr add <IP address>**

To remove a primary device:

**config backup-sr remote-sr remove <IP address>**

7. To view the current backup configuration:

**show -run backup-sr**

8. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure clock**

If an NTP server is used to set device times in your network, refer to “configure sntp” on page 362. If your network does not use an NTP server, you should manually configure the time settings for each WX device. The date and time is saved with each entry in the system log files, which can help you troubleshoot problems if they arise.

1. To view the current clock settings:

```
show -run clock
```

2. To set the data and time:

```
config clock set time <YYYYMMDDhhmm>
```

For example, to set the time to 12:30 p.m. March 16, 2003:

```
config clock set time 200303161230
```

3. To set the time zone for the device:

```
config clock set location <id>
```

Where < id > is the ID number of the time zone (1 to 74). To view the list of time zones:

```
config clock set location ?
```

4. To set daylight savings time on (if applicable):

```
config clock set daylight-saving on
```

5. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure console**

You can configure the baud rate for the DB9 console port on the back of the WX device. The default baud rate is 9600.

1. To view the current baud rate:

```
show -run console
```

2. To set the baud rate:

```
config console set baud-rate <number>
```

3. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure dns**

You can specify the DNS servers used to resolve IP addresses on the Traffic report, and the local DNS domain name of the WX device. When an IP address in the local domain is resolved by one of the DNS servers, the domain name is prepended to the host name shown on the Traffic report.

1. To view the current DNS settings:

```
show -run dns
```

2. To specify up to three DNS servers:

```
config dns server set <IP-address1 IP-address2 IP-address3> | none
```

Specify “none” to remove all DNS server addresses.

To add up to three DNS servers:

```
config dns server add <IP-address1 IP-address2 IP-address3>
```

To delete one or more DNS servers:

```
config dns server remove <IP-address1 IP-address2 IP-address3>
```

3. To specify or change the local domain name (up to 256 characters):

```
config dns set domain-name <name> | none
```

The domain name must include at least one period, but not as the first or last character. If the local domain is not specified, 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. Specify “none” to remove the local domain name.

4. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure filter**

By default, all applications running over TCP or UDP (except Groupwise, HTTPS, SNMP, SSH, and Traceroute) are enabled for data reduction. The **filter** command lets you specify the applications, protocols, or source and destination address pairs to be reduced. You can also disable the reduction of packet fragments. Note that a source/destination filter, which applies to all traffic, is applied before the application filter, and is more efficient.

For example, to conserve system capacity, you should exclude applications whose traffic is encrypted or already compressed because the reduction will be minimal. Note that applications must be defined before they can be filtered. To create application definitions, refer to “configure application” on page 309. Undefined applications are reduced by default.

Note the following:

- If you disable data reduction between a source and destination, traffic acceleration between those points is also disabled. Also, in oversubscribed mode, 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.

1. To view the current filter settings:

```
show -run filter
```

2. Type the following command to enter the configure filter mode:

```
config filter
```

3. To include or exclude one or more applications from data reduction:

```
add application <name1 name2 ...>
```

Names that include spaces must be enclosed in quotation marks. Multiple applications must be separated by commas (no spaces). Use the `show application` command to view the names of your currently defined applications.

Indicate whether the specified applications are included or excluded from data reduction:

```
set mode-applications <off | include | exclude>
```

Set the mode to “off” to reduce all applications (the default).

To remove one or more applications from the filter:

```
remove application <name1 name2 ...>
```

4. To include or exclude all traffic between two addresses or subnets:

```
add bi-address-pair <IP address>[/<mask>]-<IP address>[/<mask>]
```

An asterisk (\*) can be used alone (no subnet mask) to indicate any IP address, such as  
“\*-192.168.1.2”.

To include or exclude traffic in just one direction:

```
add address-pair <from IP address>[/<mask>]-<to IP address>[/<mask>]
```

Indicate whether the address pairs are included or excluded from data reduction:

```
set mode-address-pair <off | include | exclude>
```

Set the mode to “off” to reduce traffic between all eligible addresses (the default).

To remove one or all address pairs from the filter:

```
remove address-pair {all | <IP address>[/<mask>]-<IP address>[/<mask>]}
```

5. To enable or disable the reduction of packet fragments:

```
set ip-fragments {on | off}
```

All packet fragments are reduced by default. Fragments may not be associated with the correct application, but disabling reduction may cause fragments to arrive before the reduced packets that should precede them.

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure interface**

The Interface command lets you set the interface speeds and duplex modes, run a test to detect a mode mismatch on the Local or Remote interface, enable the reduction of VLAN traffic that adheres to the IEEE 802.1Q specification, and reset the interface traffic statistics to zero.

In addition, you can enable high-availability support so that when a failure is detected on one interface, the other interface is turned off. This allows the switch or router to detect the failure, and ensures that the routing mechanisms work as expected.

1. To view the Local and Remote interface MAC addresses, configuration, and traffic statistics:

```
show -run interface -verbose
```

2. Type the following command to enter the configure interface mode:

```
config interface
```

3. To reset the Local or Remote interface statistics to zero:

```
reset-stats <local | remote>
```

4. To set the speed and duplex mode setting for the Local or Remote interface (Gigabit speeds are available for the WX 60, WX 100, and WXC 500):

```
set speed-duplex local <auto | 10-half | 10-full | 100-half | 100-full> | <1000-full>
set speed-duplex remote <auto | 10-half | 10-full | 100-half | 100-full> | <1000-full>
```

The fiber-optic WX 100 interfaces support only 1000 Mbps with full-duplex.

5. To enable link status propagation from the Local interface to the Remote interface:

```
set propagate-failure local-to-remote on
```

If the switch fails, the Remote interface is turned off so that the router can detect the loss of connectivity with the switch.

To enable link status propagation from the Remote interface to the Local interface:

```
set propagate-failure remote-to-local on
```

If the router fails, the Local interface is turned off so that the switch can detect the loss of connectivity with the router.

Specify the number of seconds that the interface is shut down (the default is 15) or use “forever” to shut down the interface indefinitely:

```
set down-time local-to-remote <seconds | forever>
set down-time remote-to-local <seconds | forever>
```

6. To test the duplex settings between the local or remote interface and an IP address:

```
test <local | remote> <IP address>
```

This test sends test packets to the specified IP address.

7. To enable or disable a periodic test of the duplex settings on both interfaces (enabled by default):

```
set enable-periodic-test <on | off>
```

This test does not send any packets. If mismatched duplex settings are detected, an error message is displayed above the menu frame in the Web Console, and when you log in to the CLI. A mismatch can be detected only when data is sent and received at the same time.

8. To enable or disable the reduction of 802.1q VLAN traffic (disabled by default):

```
set vlan mode <on | off>
```

To specify the default VLAN ID (1 through 4095) used for untagged frames in the VLAN environment where the WX device is installed:

```
set vlan native-id <1-4095>
```

Specify the VLAN ID (1 through 4095) for the port where the Local interface of the WX device is connected. On ports that have multiple VLANs, specify the VLAN that has the largest number of hosts.

```
set vlan id <1-4095>
```

To preserve the VLAN ID in the header of reduced packets for routers that use the ID for QoS, MPLS, or other functions (disabled by default):

```
set vlan preserve <on | off>
```

In some VLAN environments, local routes may be discovered on the WAN side of the device. To add WAN-side routes to the list of reduction subnets, enable the WAN reduction subnet option (refer to “configure reduction-subnet” on page 350).

9. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## configure ip

During the installation process, you entered an IP address, subnet mask, and a default gateway so that the WX device can communicate with other devices in your network. You can use the following CLI commands to change any of these settings.

1. To view the current IP address, subnet mask, and gateway:

```
show -run ip
```

2. To set the IP address for the device:

```
config ip set ip-address <IP address>
```

3. To set the subnet mask for the device:

```
config ip set subnet-mask <subnet mask>
```

4. To set the default gateway for the device:

```
config ip set default-gateway <gateway ip address>
```

5. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.




---

**NOTE:** If you change the IP address or subnet mask, you must reboot the device. Also, if this device is a registration server, you must first transfer the registration server to another WX device before changing the IP address (refer to “configure reg-server” on page 352).

---

## configure ipsec

IP security (IPSec) can be used to authenticate and encrypt traffic between any pair of WX devices (endpoints) in the same community. IPSec must be enabled on both devices, and both devices must be configured with the same pass phrase (preshared key) and security algorithms. Encryption can also be enabled based on the traffic path (refer to “configure multi-path” on page 326). Each WX device can encrypt traffic for up to 100 other WX devices (the WX 15 and WX 20 are limited to two and five devices, respectively).

When IPSec is enabled, all compressed and passthrough traffic destined for the peer device is encrypted. For traffic sent to unadvertised subnets (no remote WX device), you can define a default IPSec policy that specifies the remote subnets for which traffic is sent unencrypted or dropped and logged.

To manage IPSec configurations, you define templates that specify the security algorithms and key lifetimes for outgoing traffic, and then apply a template to each of the remote WX devices that support IPSec. The predefined template named “Wizard” has the following properties:



- **Encryption.** Advanced Encryption Standard with a 128-bit key (AES-128)
- **Authentication.** Secure Hash Algorithm (HMAC/SHA-1)
- **Key lifetimes.** Keys are limited to 24 hours or 100 MB of traffic.

When you configure IPsec for the first time, you should use the Setup Wizard (refer to “Using the IPsec Setup Wizard” on page 215). The Setup Wizard updates the Wizard template.

To change the IPsec settings:

1. To view the current IPsec settings:

```
show -run ipsec [sa [<ip-address>]]
```

Where:

- **sa [<ip-address>]**. Displays the inbound and outbound security associations (SAs) for each endpoint or just the specified endpoint. Each SA specifies the algorithms and generated keys used to protect traffic in one direction. The SA information includes:
  - **SA Index.** Number that identifies each SA, also called the Security Parameter Index (SPI). To establish a secure connection, the outbound SA index on the sender must match an inbound SA index on the receiver.
  - **State.** Indicates whether an SA is “mature” (active) or “dying” (the key lifetime has expired). A new SA is negotiated when the key lifetime reaches 80 % of the time limit or 50 % of the data limit. After the first key expires, each endpoint has four SAs: two active (inbound and outbound) and two that are “dying.”
  - **Sequence #.** Indicates the sequence number of the last packet received. A packet is dropped if its sequence number is a duplicate or is not within 32 of the last received sequence number. Used for anti-replay protection.

2. Type the following command to enter the configure IPsec mode:

```
config ipsec
```

3. To enable or disable IPsec (disabled by default):

```
set mode <on | off>
```

4. To add a new IPsec template (only the name is required):

```
template add name <name> [key-time-lifetime <hours>] [key-data-lifetime <MB>]
[encryption any | AES-128 | AES-192 | AES-256 | 3DES] [authentication any |
HMAC/SHA-1 | HMAC/MD5]
```

Where:

- **name <name>**. Template name (up to 20 characters). If the name includes spaces, enclose the name in quotation marks.

- **key-time-lifetime** < hours > . Number of hours (up to 2160) before the generated security keys are renegotiated (default is 24). A zero indicates that the keys have no time limit.
- **key-data-lifetime** < MB > . Number of megabytes of traffic (up to 4000) before the generated security keys are renegotiated (default is 100). A zero indicates that the keys have no data limit. If both lifetimes are set, keys are renegotiated when 75 % of either limit is reached.
- **encryption** any | **AES-128** | **AES-192** | **AES-256** | **3DES**. Algorithm used to encrypt outbound traffic. Specify “any” to use the algorithm selected for the other endpoint.  
If both endpoints specify “any,” AES with a 128-bit key is used (the default). Note that triple Digital Encryption Standard (3DES) is slower and less secure than AES.
- **authentication** any | **HMAC/SHA-1** | **HMAC/MD5**. Algorithm used to authenticate outbound traffic. Specify “any” to use the algorithm selected for the other endpoint.  
If both endpoints specify “any,” HMAC/SHA-1 is used (the default). HMAC/SHA-1 provides more security, but HMAC/MD5 is two to three times faster.

To change a template, specify the template name and the settings you want to change:

```
template set name <name> [new-name <name>] [key-time-lifetime <hours>]
[key-data-lifetime <MB>] [encryption any | AES-128 | AES-192 | AES-256 | 3DES]
[authentication any | HMAC/SHA-1 | HMAC/MD5]
```

To delete an IPSec template:

```
template remove <name>
```

If the deleted template was applied to an endpoint, the endpoint reverts to the Wizard template. The Wizard template can be changed, but not deleted.

5. To assign an IPSec template to a remote endpoint:

```
endpoint add ip-address <address> [template <name>] [mgmt-traffic-mode <on | off>] [pass-phrase]
```

Where:

- **ip-address** < address > . IP address of a WX device that supports IPSec.
- **template** < name > . Name of an IPSec template (default is “Wizard”).
- **mgmt-traffic-mode** < on | off > . Indicates whether management traffic for the remote endpoint is encrypted (disabled by default). Should be disabled during testing. Management traffic includes SNMP, Syslog, and registration server traffic.
- **pass-phrase**. Prompts you for a password when you press **Enter**. The password is used to generate a pre-shared key of the appropriate length. Type the pass phrase (4 to 64 characters), press **Enter**, and then repeat to verify. The same pass phrase must be specified on the remote device.

Alternatively, you can specify the same pass phrase for all endpoints:

```
set common-pass-phrase
```

and press **Enter**. Type the password (at least four characters), press **Enter**, and then repeat to verify. You must then enable the common pass phrase (disabled by default):

```
set common-pass-phrase-mode <on | off>
```

To change an endpoint, specify the endpoint address and the settings you want to change:

```
endpoint set ip-address <address> [template <name>] [mgmt-traffic-mode <on | off>] [pass-phrase]
```

To disable IPSec for an endpoint:

```
endpoint remove <IP address>
```

Traffic to a deleted endpoint will be unencrypted.

6. The default IPSec policy is applied to traffic sent to unadvertised subnets (no remote WX device), and to traffic between WX devices where IPSec is enabled, but the key negotiation has failed. By default, all such traffic is unencrypted.

To add a destination address or subnet to the default policy for which traffic must be dropped and logged:

```
encrypt-required-subnets add <address>[/mask]
```

After you verify that IPSec is working correctly, all subnets advertised IPSec-enabled peers should be added to the encryption-required list to avoid sending unencrypted traffic to those subnets if a remote WX device fails.

To specify an address or subnet where encryption is optional:

```
encrypt-optional-subnets add <address>[/mask]
```

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, the traffic is sent unencrypted.

To remove a required or optional subnet from the default policy:

```
encrypt-required-subnets remove <address>[/mask]
encrypt-optional-subnets remove <address>[/mask]
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure license**

Each non-backup WX device requires a permanent license key for normal operation. The license key determines the licensed modules and throughput for the device, and properly registers the product. Initially, each device has a temporary 30-day license with access to all features. When the temporary license expires, all traffic will pass through without reduction. Temporary licenses are used for backup WX devices because only the active device time is counted against the 30-day limit (WXOS 5.1 or later required).

To obtain a permanent license key, you need:

- Device serial number displayed in the License Key page (also displayed in the About box and on the back of the device)
- One or more Right To Use (RTU) keys that were emailed to you in a PDF file
- User ID and password to access the License Key server at:

<http://license.peribit.com>

If you do not have Internet access, please call Technical Support at +1-888-314-JTAC (U.S, Canada, and Mexico) or +1-408-745-9500.

The speed RTU key specifies the licensed speed and level of support for the device. A separate RTU is needed for each optional feature (such as IPSec encryption). If you do not enter an RTU key, the device is licensed for the base speed with no customer support. If you lose the license key, you can use the License Key server to retrieve your current license key.

To view or change the license key:

1. To view the current license and device serial number:

```
show -run license
```

2. If you have a temporary license, obtain and apply a permanent license key:

- a. Go to the Online License Server at <http://license.peribit.com>.
- b. Log in to the server, enter your contact information, the device serial number, and click Submit.
- c. Enter the RTU keys for the desired device speed, level of support, and optional features, and click Yes. If you omit the RTU keys and click No, the device is licensed for the base speed and modules, with no technical support.
- d. On the WX device, enter the license key displayed by the server:

```
config license set license-key <new license key>
```

## **configure mon-apps**

You can select the applications to be monitored, as well as enable or disable the monitoring of WAN traffic. If an application is monitored, you can view performance statistics for the application (up to 40 applications can be monitored). WAN traffic monitoring is required to view the WAN reports generated when WAN performance monitoring is configured (refer to “configure wan-performance-monitor” on page 366). For more information about monitoring statistics, refer to “Monitoring and Reporting” on page 227.

Only defined applications can be monitored. Application definitions are provided for applications with well-known port numbers. All other applications are grouped together as “Undefined” or “Others”. Undefined applications are monitored automatically. To define additional applications, refer to “configure application” on page 309.



**NOTE:** If you disable monitoring for an application, its historical monitoring statistics are moved to the “Others” application category on reports. If monitoring is re-enabled, the historical statistics remain in the “Others” category.

1. To view the applications being monitored:

```
show -run mon-apps
```

2. Type the following command to enter the configure monitored applications mode:

```
config mon-apps
```

3. To clear the current list of monitored applications:

```
clear
```

To specify one or more applications to be monitored:

```
add <application1 application2 ...>
```

Multiple applications must be separated by spaces. If an application name contains spaces, enclose the entire name in quotation marks.

To remove one or more applications from the monitoring list:

```
remove <application1 application2 ...>
```

4. To enable or disable WAN traffic monitoring (enabled by default):

```
wan-traffic <on | off>
```

5. To specify whether traffic reports show application port names for reserved port numbers (up to 1024) or all port numbers (the default is “all”):

```
set port-map <all | reserved>
```

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

**configure multi-node**

In multi-node configuration, WX devices work in conjunction to provide greater reduction of data on higher-speed or heavily congested networks. For more information on installing and administering multi-node configurations, refer to “Multi-Node Configurations” on page 433.

**configure 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. Selected traffic can be sent over a preferred path under normal conditions, and dynamically switched to the alternate path when the preferred path fails or when congestion or latency exceed a specified threshold. Note that each Multi-Path endpoint counts as two reduction tunnels.

For example, if you normally send database traffic over Frame Relay, and email traffic over the Internet, you can automatically divert the database traffic to the Internet if Frame Relay fails, and divert email traffic to Frame Relay if the Internet becomes congested. Traffic is switched back to the preferred path when conditions return to normal.

To use Multi-Path, you must:

- Configure a secondary source IP address to be used for outgoing packets intended for the secondary path. You can also specify a primary and secondary gateway address or ToS/DSCP value. Note that ToS/DSCP values override the ToS/DSCP settings defined for outbound QoS.
- Define templates that specify the preferred path (primary or secondary) for each outbound QoS traffic class and the conditions when the traffic for each class can be switched.
- Apply a template to the remote WX devices that support Multi-Path, and specify the congestion and latency thresholds for each path to the remote device.
- If necessary, configure the WAN routers to route the marked packets to the appropriate path.

Note that data reduction must be enabled on each WX device that supports Multi-Path (refer to “configure reduction” on page 343). Encryption can be enabled for one or both paths (refer to “configure ipsec” on page 320).

1. To view the current multi-path settings:

```
show -run multi-path
```

To view the last 32 events when traffic was switched between primary and secondary paths:

```
show multi-path events [<address>]
```

2. Type the following command to enter the configure multi-path mode:

```
config multi-path
```

3. To enable or disable multi-path processing (disabled by default):

```
set mode <on | off>
```

4. On the subnet where the WX device is installed, reserve a unique secondary IP address to be used as the source address on packets sent on the secondary path (on packets sent on the primary path, the device address is the source address). To specify the secondary address:

```
set sec-ip-address <address>
```



**NOTE:** If you must change the secondary address, enter “set mode off” and “commit” commands, and then enable Multi-Path again and specify the new address.

Optionally, packets sent on the primary and secondary paths can be marked with different ToS/DSCP values or gateway addresses. You can specify values for both marking methods, but only one method can be used for each endpoint.

If the WAN routers for the primary and secondary paths are on the same subnet as the WX device, enter their IP addresses. In this case, no additional router configuration is needed. To specify primary and secondary gateway addresses:

```
set gateway-ip <primary-address> <secondary-address>
```

If the WAN routers for the two paths are on separate subnets, the default gateway must be configured to route traffic to the appropriate WAN link (refer to “Configuring Routers to Support Multi-Path” on page 123).

To specify primary and secondary ToS IP precedence values (0 to 7) or DSCP values (0 to 63), set the mode (default is IP precedence), and then set an IP precedence or DSCP value:

```
set ip-precedence-dscp-mode <ip-precedence | dscp>
set ip-precedence <primary> <secondary>
set dscp <primary> <secondary>
```



**NOTE:** These values override the IP precedence or DSCP settings defined for:

- Outbound QoS (refer to “configure qos outbound” on page 337)
- WX control packets (refer to “configure reduction” on page 343)

Also, multi-path DSCP values override ToS type-of-service settings used for Cisco router balancing (refer to “configure route” on page 356)

5. To add a new multi-path template:

```
template add name <name>
```

Where:

- **name <name>** . Template name (up to 20 characters). If the name includes spaces, enclose the name in quotation marks.

By default, a new template specifies that each QoS traffic class uses the primary path and is never switched to the alternate path. To change the preferred path and bypass condition for a traffic class:

```
template class set name <name> class-name <name> [preferred-path <primary | secondary>] [bypass-condition <never | failure-only | performance-failure>]
```

Where:

- **name <name>** . Template name. If the name includes spaces, enclose the name in quotation marks.
- **class-name <name>** . Traffic class name. To view the current traffic classes, enter the “show qos outbound” command. To add a new traffic class, refer to “configure qos outbound” on page 337.
- **preferred-path <primary | secondary>** . Indicates the default path used by the traffic class (default is primary).
- **bypass-condition <never | failure-only | performance-failure>** . Indicates when this traffic class is switched: never, only when the other path fails, or when the path fails or the specified congestion or latency thresholds are exceeded (default is never).

To change a template name:

```
template set name <name> new-name <name>
```

To delete a multi-path template:

```
template remove <name>
```

A template assigned to an endpoint cannot be deleted until the endpoint is deleted.

6. To apply a multi-path template to a remote endpoint:

```
endpoint add ip-address <address> template <name> [marking-method <ip-only | gateway-ip | tos-dscp>]
```

Where:

- **ip-address <address>** . IP address of a remote WX device that supports multi-path processing.
- **template <name>** . Name of a multi-path template.



- **marking-method** < **ip-only** | **gateway-ip** | **tos-dscp** > . Indicates whether packets on the primary and secondary paths are distinguished only by the source IP address (the default) or also by the gateway IP address or ToS/DSCP value specified in Step 4.

To change an endpoint's template and/or marking method:

```
endpoint set ip-address <address> [template <name>] [marking-method <ip-only | gateway-ip | tos-dscp>]
```

To disable multi-path processing for an endpoint:

```
endpoint remove <IP address>
```

7. To change the default loss and latency thresholds for the primary or secondary paths to a remote endpoint:

```
endpoint path set ip-address <address> latency-tolerance <20-5000>
probes-per-minute <1-60> probes-above-latency <1-60> probes-lost <1-60>
minutes-to-divert-la <1-32> minutes-to-divert-lo <1-32>
minutes-to-return-la <1-32> minutes-to-return-lo <1-32>
```

Where:

- **ip-address** < **address** > . Primary or secondary IP address of a remote WX device that supports multi-path processing.
  - **latency-tolerance** < **20-5000** > . Latency in milliseconds that must be exceeded before traffic is switched to the alternate path (default is 5000).
  - **probes-per-minute** < **1-60** > . Number of times per minute that the path is tested (default is 12).
  - **probes-lost** < **1-60** > . Number of probes that must be lost per minute before the minute is marked as above the loss threshold (default is 2).
  - **minutes-to-divert-la** < **1-32** > . Number of consecutive minutes that the median latency must exceed the latency threshold before traffic is switched to the alternate path (default is 4).
  - **minutes-to-divert-lo** < **1-32** > . Number of consecutive minutes that must exceed the loss threshold before traffic is switched to the alternate path (default is 4).
  - **minutes-to-return-la** < **1-32** > . Number of consecutive minutes of acceptable latency required before traffic is switched back to the primary path (default is 4).
  - **minutes-to-return-lo** < **1-32** > . Number of consecutive minutes of acceptable loss required before traffic is switched back to the primary path (default is 4).
8. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

**configure ospf**

If your network uses OSPF, you can enable OSPF support on the WX device. The OSPF routes on the local side of the device are discovered and added to the Local Routes table.

1. To view the current OSPF settings:

```
show -run ospf [all | neighbor [detail]]
```

The “all” option shows all configuration and neighbor information. The “neighbor detail” option shows details of the neighboring OSPF-enabled routers, such as the designated router (DR) and backup designated router (BDR). For example:

```
===== OSPF Neighbors =====
```

ID	Pri	State	Dead Time	Address	Interface
13.13.13.1	1	2-Way	00:00:37	10.200.1.1	fei
14.14.14.2	1	Full	00:00:39	10.200.1.3	fei
15.15.15.2	1	2-Way	00:00:39	10.200.1.16	fei
11.11.11.2	1	2-Way	00:00:40	10.200.1.2	fei
16.16.16.2	1	Full	00:00:39	10.200.1.25	fei

```
===== OSPF Neighbors' Details =====
```

```
Neighbor 13.13.13.1, interface address 10.200.1.1
```

```
In the area 0 via interface fei
Neighbor Priority is 1, State is 2-Way, 2 state changes
DR is 10.200.1.25
BDR is 10.200.1.3
Options is DC N/P (0x15)
Dead timer due in 37 seconds
Authentication: none
```

2. Type the following command to enter the configure OSPF mode:

```
config ospf
```

3. To enable or disable OSPF:

```
set ospf <on | off>
```

4. To enter an OSPF area ID:

```
set area <IP address in dotted-decimal notation or a number>
```

5. To specify the type of OSPF authentication (the default is none):

```
set auth-type <crypt | password | none>
```

If you set OSPF authentication to “crypt,” specify the MD5 key ID (1 to 255) and encryption key (up to 16 characters):

```
set crypt <key-id> <key>
```

If you set OSPF authentication to “password,” specify the password (up to 8 characters):

```
set password <password>
```

6. To change the number of seconds (1 to 65535) between the sending of OSPF hello packets (the default is 10):

```
set hello-interval <number>
```

To change the number of seconds (1 to 65535) before adjacent routers assume the WX device is down when no hello packets are received (the default is 40):

```
set dead-interval <number>
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.



**NOTE:** If you change the dead interval, you must stop and restart the OSPF service for the change to take effect.

### **configure packet-interception**

If the WX device is deployed off-path, where only the Local port is connected to the network, you can use one of the following methods to route traffic to the Local port for reduction.

- **Route injection.** The Routing Information Protocol (RIPv2) is used to advertise the off-path WX device as the lowest cost “router” for the remote routes advertised by the other devices in the community. Requires that surrounding routers give the highest priority to RIP routes. When RIP is used, note the following:
  - To advertise the subnet where a remote WX device is installed, several new subnets are generated to exclude the IP address of the remote device. This prevents the router from returning the traffic tunneled to the remote device.
  - The off-path WX device has no passthrough data. Both reduced and unreduced traffic is sent through the reduction tunnel.
- **WCCP.** The Web Cache Communication Protocol is used to redirect specific types of traffic from the router to the off-path device. The router must support WCCP version 2. Refer to the sample router commands in “WCCP Router Configuration Commands” on page 112.
- **External.** The WAN edge router is configured to route traffic to the off-path device. The off-path device must be connected directly to the router. Refer to the sample router commands in “External Policy-Based Router Commands” on page 113.

In each case, the redirected traffic is reduced (if eligible) and returned to the WAN edge router over the Local interface. Note that off-path WX devices do not support multi-node configurations. Also, outbound bandwidth management is limited to the WAN traffic that is routed through the off-path device.

1. To view the current packet interception settings:  
`show -run packet-interception`
2. Type the following command to configure off-path interception:  
`config packet-interception`
3. To enable or disable off-path interception (disabled by default):  
`set mode {rip | wccp | external | off}`



**CAUTION:** Enabling packet interception disables the Remote interface. If the WX device is installed in the data path, data transmission through the device will stop.

4. If you use RIP, you can specify the frequency of RIP updates, the delay between each route in an update, and the cost (metric) assigned to each route.
  - a. To change the number of seconds between RIP updates (the default is 30):  
`rip set update-timer <1-7200>`  
This value must match the update timer setting on the router.
  - b. To reduce the load on slower routers, you can specify a delay between each packet in a RIP update (default is 0). To specify the number of milliseconds between each packet (0 through 50):  
`rip set delay <0-50>`
  - c. Each route has a default metric (cost) of two. To change the metric (1 through 15):  
`rip set metric <1-15>`
5. If you use WCCP, you must specify the router IP address, authentication, WCCP priority, and a service ID for each protocol whose traffic you want redirected to the off-path device.
  - a. To specify the router address:  
`wccp set router-ip-address <iP address>`
  - b. If the Cisco router requires a WCCP password:  
`wccp set auth-type password`  
To specify the password:  
`wccp set password`  
At the prompts, enter and verify the password.
  - c. A WCCP priority value (0 through 255) is required to indicate the order in which packets are compared against the services (protocols) you specify, relative to the other services redirected by the router. Higher values have a higher priority. The default is 230. To specify the WCCP priority:

```
wccp set priority <0-255>
```

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 priority to avoid “stealing” traffic from the Web cache.

- d. To specify a protocol whose traffic you want redirected to the off-path device:

```
wccp protocol add {tcp | udp | <protocol-number>} <service-id>
```

Where:

- **<protocol-number>** . IP protocol number (0 to 255). The standard protocol numbers are listed at:  
*<http://www.iana.org/assignments/protocol-numbers>*.
- **<service-id>** . WCCP service ID number for the protocol (51 through 99). The number must be unique among all the WCCP services defined on the router.

- e. To stop the redirection of a protocol’s traffic:

```
wccp protocol remove {tcp | udp | <protocol-number>}
```

6. If you use external mode, passthrough traffic is returned to the router or switch. If this causes routing loops, you can disable passthrough mode so that passthrough traffic is included in the reduction tunnels. However, if an appropriate reduction tunnel does not exist, the traffic is dropped. To enable or disable passthrough mode (enabled by default):

```
external set pass-through {on | off}
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure prime-time**

The prime time command 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. For example, to view reduction and acceleration statistics during business hours, you could set the prime time to 9:00 AM to 5:00 PM on Monday through Friday.

1. To view the current prime-time settings:

```
show -run prime-time
```

2. Type the following command to enter the configure prime-time mode:

```
config prime-time
```

3. To enable or disable prime time:

```
set mode {on | off}
```

Prime time is disabled by default, which means the effective “prime time” is 24-hours a day, seven days a week.

4. To specify the days of the week in prime time:

```
set days {mon,tue,wed,thu,fri,sat,sun}
```

The days must be separated by commas (no spaces).

5. To specify the prime-time hours for the selected days of the week:

```
set hours <hour-hour>
```

Where the time range is in 24-hour format (such as “9-17” for 9 AM to 5 PM).

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure profile-mode**

When a WX device is installed in Profile Mode, you can view the performance for specific remote subnets by defining “virtual” WX devices and associating one or more subnets with each virtual device. If you also specify a circuit speed, it is used to estimate the maximum possible acceleration of TCP traffic that might be obtained using Packet Flow Acceleration.

On the reduction and acceleration reports, you can select a virtual device from the Destination menu to view the performance for the associated remote subnets (refer to “Monitoring and Reporting” on page 227).

1. To view the current Profile Mode settings:

```
show -run profile-mode
```

2. Type the following command to configure Profile Mode:

```
config profile-mode
```

3. To enable or disable Profile Mode (disabled by default):

```
set mode <on | off>
```



**CAUTION:** Enabling Profile Mode disables the Remote interface. If the WX device is installed in the data path, all data transmission through the device will stop.

---

4. To add a “virtual” WX device and its circuit speed (in Kbps):

```
remote-sr add <IP address> [<speed>]
```

To remove a virtual device:

```
remote-sr remove <IP address>
```

5. To associate a remote subnet with a virtual device:

```
remote-sr subnet add <IP address> <subnet/mask>>
```

To delete a subnet from a virtual device:

```
remote-sr subnet remove <IP address> <subnet/mask>>
```

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.


### **configure qos inbound**

Inbound bandwidth management lets you specify maximum bandwidths for four classes of incoming WAN traffic destined for the Local Area Network (LAN). Setting a maximum bandwidth for each class (and optionally the queue length) ensures that low-priority traffic, such as Web traffic, does not interfere with mission-critical applications. Bandwidths are specified as percentages of the inbound speed (aggregate local WAN speed), and traffic that exceeds the maximum bandwidths is dropped.

The following table describes the four traffic classes for inbound bandwidth management:

Class	Description
Reduced	Reduced traffic from other WX devices.
Intranet	Unreduced TCP traffic from a specified list of IP subnets (such as the subnets that have no WX device). Use the Traffic report to help create the list of subnets (refer to “Traffic Statistics” on page 256).
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.

In the following example, to enable inbound QoS on D1, you set the local inbound speed to 1500 Kbps (1.5 Mbps), and then set maximum bandwidth percentages for one or more of the four traffic classes. In this example, you might set the maximum bandwidth for the Default class to 10 % to limit low-priority traffic from the public Internet.



1. To view the current inbound QoS settings:

```
show -run qos inbound
```

2. Type the following command to enter the configure inbound QoS mode:

```
config qos inbound
```

3. To specify the local inbound speed in Kbps (8 to 1000000) for the WAN edge router associated with the WX device:

```
aggregate-wan-speed <8-1000000>
```

4. To configure the bandwidth limits and queue lengths (optional) for each class:

```
class-default max-bw <percentage> [queue-len <1-512>]
class-intranet max-bw <percentage> [queue-len <1-512>]
class-reduced max-bw <percentage> [queue-len <1-512>]
class-tcp max-bw <percentage> [queue-len <1-512>]
```

Where:

**max-bw <percentage>** . Maximum percentage of the inbound speed allowed for traffic in the specified class. A zero indicates that all traffic in the class will be dropped. A value of 100 (the default) effectively disables inbound bandwidth management for the class.

**queue-len <1-512>**. Maximum number of packets allowed in the queue for this class (the default is 40).



**NOTE:** Please contact Technical Support for assistance before changing the queue lengths.

---

5. For the Intranet class, the maximum bandwidth setting applies only to traffic from the subnets you specify. To define a subnet in the Intranet class:

```
define-intranet add <IP Address/Subnet Mask>
```

To remove a subnet from the Intranet class:

```
define-intranet remove <IP Address/Subnet Mask>
```

6. To enable or disable inbound bandwidth management (disabled by default):

```
bw-mgmt <on | off>
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.



## **configure qos outbound**

Outbound bandwidth is managed by assigning applications to traffic classes, defining templates that specify a priority, guaranteed bandwidth, and maximum bandwidth for each traffic class, and then applying a template to the remote WX devices for which you want to manage outbound traffic. You can also specify LAN/WAN address or subnet pairs to be excluded from bandwidth management. For an overview of outbound bandwidth management, refer to “Understanding Outbound Bandwidth Management” on page 154.

When you configure the outbound QoS settings for the first time, you should use the Setup Wizard (refer to “Using the Outbound QoS Setup Wizard” on page 164). The Setup Wizard creates two templates with the same settings:

- Wizard-PrimeTime
- Wizard-NonPrimeTime

If you use the Web console to customize the settings for specific endpoints, new templates are created whose names include the IP address of the endpoint:

- PTO- <IP\_address> for customized prime-time templates
- NTO- <IP\_address> for customized nonprime-time templates

To change the outbound QoS settings:

1. To view the current outbound QoS settings:

```
show -run qos outbound
```

2. Type the following command to enter the configure inbound QoS mode:

```
config qos outbound
```

To undo your outbound QoS changes by copying the running configuration to the candidate configuration, exit from configuration mode and type:

```
rollback
```

3. To enable or disable outbound bandwidth management, select one of the prioritization methods (disabled by default):

```
set mode <bw-weighted-fair-queueing | bw-strict-priority | off>
```

Where:

- **bw-weighted-fair-queueing.** Queues are created for each traffic class, and processed according to their priority and guaranteed bandwidth.
- **bw-strict-priority.** Queues are created for each priority, and processed according to their priority. For traffic classes that have the same priority, processing is weighted by the guaranteed bandwidth.

4. The following table describes the outbound QoS settings.

Settings	Commands
Outbound speed and Oversubscribed Mode	<p>To specify the local outbound speed in Kbps (8 to 1000000) for the WAN edge router associated with the WX device:</p> <pre>set aggregate-wan-speed &lt;8-1000000&gt;</pre> <p>To specify the local WAN as oversubscribed (disabled by default):</p> <pre>set oversubscribed-mode &lt;on   off&gt;</pre>
Traffic classes	<p>Initially, all applications belong to the "Default" class. To add the name of a new traffic class (up to 20 characters):</p> <pre>class add &lt;name&gt;</pre> <p>To move an application to a new class (an application can belong to only one class):</p> <pre>class application move &lt;class&gt; &lt;application&gt;</pre> <p>To change the name of a class:</p> <pre>class set name &lt;oldname&gt; &lt;newname&gt;</pre> <p>To delete a class (any applications in the class are moved to the Default class):</p> <pre>class remove &lt;name&gt;</pre>
Templates	<p>To add a new template (up to 20 characters):</p> <pre>template add &lt;name&gt;</pre> <p>Different templates can be defined for prime-time and nonprime-time hours. To specify a template's guaranteed bandwidth percentage (0 to 80) for a traffic class (default is zero):</p> <pre>template set bw-guaranteed &lt;template&gt; &lt;class&gt; &lt;percent&gt;</pre> <p>The total guaranteed bandwidth percentage for all classes cannot exceed 80%. To specify a template's maximum bandwidth percentage (0 to 100) for a traffic class (the default is 100%):</p> <pre>template set bw-max &lt;template&gt; &lt;class&gt; &lt;percentage&gt;</pre> <p>Traffic is dropped when the maximum bandwidth is exceeded. A zero indicates that all traffic in the class is dropped.</p> <p>To specify a template's priority (0 to 7) for a traffic class, where 7 is the highest priority (the default is zero):</p> <pre>template set priority &lt;template&gt; &lt;class&gt; &lt;priority&gt;</pre> <p>Priority settings are used by the Strict Priority and Weighted Fair Queueing prioritization methods.</p> <p>To delete a template:</p> <pre>template remove &lt;name&gt;</pre> <p>If the deleted template was applied to an endpoint, all priority and guaranteed bandwidth values are set to zero for that endpoint. Maximum bandwidth values are set to 100%.</p> <p>To specify a template's maximum queue length (1 to 512) for a traffic class (the default is 80 packets) or the maximum number of milliseconds that a packet can be in the queue before it is dropped (the default is "no-limit"):</p> <pre>template set queue-len &lt;template&gt; &lt;class&gt; &lt;packets&gt;</pre> <pre>template set age-out &lt;template&gt; &lt;class&gt; &lt;2-5000   no-limit&gt;</pre> <p><b>NOTE:</b> Please contact Technical Support for assistance before changing queue lengths or age-out times.</p>

Settings	Commands
Endpoints	<p>To manage the outbound bandwidth to a remote WX device (endpoint), specify the device's IP address and its associated WAN circuit speed in Kbps (8 to 1000000):</p> <pre>tunnel add &lt;IP address&gt; &lt;speed&gt;</pre> <p><b>CAUTION:</b> Unless congestion control is enabled, be sure to test the WAN circuit speed. The actual WAN speed is typically less than the rated speed (refer to “WAN Circuit Speeds and Router Overhead” on page 156). Exceeding the actual WAN speed effectively shifts bandwidth management to the router, and may cause the router to drop traffic.</p> <p>To change an endpoint's circuit speed (8 to 1000000, in Kbps):</p> <pre>tunnel set &lt;address   other-traffic&gt; &lt;8-1000000&gt;</pre> <p>The predefined “other-traffic” endpoint is used to manage the bandwidth for all traffic that is not sent to one of the specified WX devices. The circuit speed for “other-traffic” defaults to the aggregate local WAN speed.</p> <p>To assign a template to an endpoint for prime-time or nonprime-time hours:</p> <pre>tunnel set prime-time &lt;address   other-traffic&gt; &lt;template&gt; tunnel set non-prime-time &lt;address   other-traffic&gt; &lt;template&gt;</pre> <p>To remove a template from an endpoint, replace the template name with a “-” in the above commands (sets all priority and guaranteed bandwidth values to zero and all maximum bandwidths to 100%).</p> <p>To delete an endpoint from outbound bandwidth management:</p> <pre>tunnel remove &lt;IP address&gt;</pre> <p>Traffic to the deleted endpoint will be managed by the “Other-traffic” endpoint. You cannot delete an endpoint for which acceleration is enabled.</p>
Congestion control	<p>If the WAN bandwidth to a remote WX device is variable, such as for Frame Relay or shared satellite links, you can enable congestion control for traffic sent to that device. This dynamically adjusts the bandwidth allocation for each endpoint based on the latency measured for the ACKs returned for each reduced meta packet.</p> <p>To enable or disable congestion control (disabled by default):</p> <pre>set congestion-control-mode &lt;on   off&gt;</pre> <p>To enable congestion control for all QoS-enabled endpoints or a list of included endpoints (default is all endpoints):</p> <pre>set congestion-control-endpoint-policy &lt;all   include&gt;</pre> <p>To enable or disable congestion control for a specific endpoint:</p> <pre>tunnel set congestion-control-mode &lt;IP address&gt; &lt;on   off&gt;</pre> <p>To specify the minimum bandwidth (in Kbps) for an endpoint (optional):</p> <pre>tunnel set min-bandwidth &lt;IP address&gt; &lt;bandwidth&gt;</pre> <p>The minimum WAN speed depends on the network. 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.</p> <p>When packet loss is detected, the TCP fast backoff method is used to reduce data transmission, and then gradually increase it. In some environments, primarily in satellite networks where packet-level load balancing is used across multiple WAN links, out-of-order packet reception may be mistaken for packet loss. In this case you can enable the SCPS backoff method to reduce data transmission more slowly (“default” enables the TCP method):</p> <pre>set congestion-control-action-on-loss &lt;tcp-fast-backoff   scps-slow-backoff   default&gt;</pre>
Tunnel passthrough	<p>When congestion control is enabled, it applies only to the tunneled traffic sent between WX devices. For remote endpoints that have congestion control enabled, you can enable or disable the inclusion of passthrough traffic in the reduction tunnel (disabled by default):</p> <pre>set tunnel-sr-passthrough &lt;on   off&gt;</pre> <p>Congestion control is not applied to passthrough traffic sent to non-WX endpoints.</p>

Settings	Commands
Virtual endpoints	<p>By default, traffic sent to non-WX destinations is managed by the QoS settings for the “Other-traffic” endpoint. To manage such traffic more closely, you can define virtual endpoints for specific remote subnets. The maximum number of virtual endpoints (up to 320) depends on the device type (2 for the WX 15, 5 for the WX 20 and WXC 250).</p> <p>You can also view the WAN Throughput and WAN Application Summary reports for each virtual endpoint (refer to “WAN Statistics” on page 228).</p> <p>To add virtual endpoints:</p> <pre>non-sr add name &lt;name&gt; bandwidth &lt;8-1000000&gt; [prime-time &lt;template&gt;] [non-prime-time &lt;template&gt;] [mode &lt;on   off&gt;]</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>■ <b>name &lt;name&gt;</b> . Virtual endpoint name (up to 20 characters). If the name includes spaces, enclose the name in quotation marks.</li> <li>■ <b>bandwidth &lt;8-1000000&gt;</b> . WAN circuit speed associated with this endpoint (in Kbps).</li> <li>■ <b>prime-time &lt;template&gt;</b> . Optional name of the template used for prime-time hours. Default is none (all priority and guaranteed bandwidth values are zero, and all maximum bandwidths are 100 %).</li> <li>■ <b>non-prime-time &lt;template&gt;</b> . Optional name of the template used for nonprime-time hours. Default is none.</li> <li>■ <b>mode &lt;on   off&gt;</b> . Enables or disables the endpoint (enabled by default). If you disable a virtual endpoint, any traffic to its associated subnets is managed by the “Other-traffic” endpoint.</li> </ul> <p><b>NOTE:</b> If you do not assign a template to a virtual endpoint, traffic to that endpoint has the lowest priority.</p> <p>To add destination addresses or subnets to a virtual endpoint (multiple addresses/subnets must be separated by spaces and enclosed in double quotation marks):</p> <pre>non-sr subnets add name &lt;endpoint&gt; subnets &lt;ip-address&gt;[/mask],&lt;ip-address&gt;[/mask]...</pre> <p>Note that adding an address/subnet to one endpoint automatically deletes it from any other virtual endpoint. If a subnet is also advertised by a WX device, the subnet here is ignored.</p> <p>To change a virtual endpoint, specify the name and the properties you want to change:</p> <pre>non-sr add name &lt;name&gt; [new-name &lt;name&gt;] [bandwidth &lt;8-1000000&gt;] [prime-time &lt;template&gt;] [non-prime-time &lt;template&gt;] [mode &lt;on   off&gt;]</pre> <p>To remove a template from an endpoint, replace the template name with a “-” in the above command (sets all priority and guaranteed bandwidth values to zero and all maximum bandwidths to 100 %).</p> <p>To remove addresses or subnets:</p> <pre>non-sr subnets remove &lt;ip-address&gt;[/mask]...</pre> <p>To delete a virtual endpoint:</p> <pre>non-sr remove &lt;name&gt;</pre> <p>Traffic to the subnets associated with the deleted endpoint will be managed by the settings for the “Other-traffic” endpoint.</p>

Settings	Commands
ToS/DSCP	<p>The ToS/DSCP values on incoming LAN traffic can be changed to support other QoS devices in the network. For each traffic class, you can specify a ToS IP precedence value or a DSCP value. The specified ToS/DSCP values apply to all traffic in the class, regardless of whether the traffic is reduced or outbound QoS is enabled.</p> <p>To specify whether ToS or DSCP values are changed (disabled by default):</p> <pre>tos-dscp set mode &lt;tos   dscp   off&gt;</pre> <p>For applications whose traffic is reduced, specify whether the ToS/DSCP value is restored to its original value after the traffic is assembled by the remote WX device (enabled by default):</p> <pre>tos-dscp set restore &lt;on   off&gt;</pre> <p>To enable or disable ToS/DSCP changes for a traffic class (disabled by default):</p> <pre>tos-dscp class set mode &lt;class&gt; &lt;on   off&gt;</pre> <p>To set a 6-bit DSCP value (0 to 63) for a traffic class:</p> <pre>tos-dscp class set dscp &lt;class&gt; &lt;0-63&gt;</pre> <p>To set an 8-bit DSCP value (0 to 255) for a traffic class (only the lower six bits are used):</p> <pre>tos-dscp class set dscp-byte &lt;class&gt; &lt;0-255&gt;</pre> <p>To set a ToS IP precedence value (0 to 7) for a traffic class:</p> <pre>tos-dscp class set ip-precedence &lt;class&gt; &lt;0-7&gt;</pre>
Excluded subnets	<p>To avoid managing traffic addressed to the router on the WAN side of the WX device, all LAN traffic sent to the WX device's subnet is excluded from outbound bandwidth management. This ensures that we manage only the traffic sent across the WAN.</p> <p>To view the current filter settings:</p> <pre>show -run qos excl-filter</pre> <p>To access the configuration mode for the exclusion filter:</p> <pre>configure qos excl-filter</pre> <p>To enable or disable the exclusion filter (enabled by default):</p> <pre>set mode &lt;on   off&gt;</pre> <p>To exclude additional LAN/WAN address or subnet pairs from outbound bandwidth management:</p> <pre>add &lt;LAN address&gt;[/mask]-&lt;WAN address&gt;[/mask]</pre> <p>Use an asterisk (*) by itself to indicate any LAN/WAN address or subnet.</p> <p><b>NOTE:</b> Traffic bursts between excluded addresses are unrestrained by priority or bandwidth considerations, and may cause other traffic to be dropped by the router.</p> <p>To delete one or all excluded LAN/WAN address or subnet pairs:</p> <pre>remove {&lt;LAN address&gt;[/mask]-&lt;WAN address&gt;[/mask]   all}</pre> <p>Use an asterisk (*) by itself to indicate any LAN or WAN address or subnet.</p>

5. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure radius**

The RADIUS command is used to define RADIUS servers and server groups. At least one server group is required. To specify how the server groups are used to authenticate users, refer to “configure aaa” on page 300.

1. To view the current RADIUS settings:

```
show -run radius
```

2. Type the following command to enter the configure RADIUS mode:

```
config radius
```

3. You can define up to 20 RADIUS servers. To add a RADIUS server:

```
server add name <name> ip-address <address> auth-port <number>
timeout <seconds> retransmit <number> dead-time <minutes>
```

Where:

- **name <name>** . RADIUS server name (up to 32 characters). If the name includes spaces, enclose the name in quotation marks.
- **ip-address <address>** . IP address of the server.
- **auth-port <number>** . Authentication UDP port number on the server (default is 1812).
- **timeout <seconds>** . Number of seconds (1 to 65535) that the WX device waits for the server to respond (default is three).
- **retransmit <number>** . Number of times (1 to 100) that a request is sent to the server (default is three).
- **dead-time <minutes>** . Number of minutes (0 to 1440) after all retransmissions fail that the WX device waits before trying to access the server again (default is zero).

and then press **Enter**. Type the secret key (up to 31 characters) used to access the server and press **Enter**, and then repeat to verify. The same key must be configured on the RADIUS server.

To change the key used to access the server:

```
server set name <name> key
```

and then press **Enter**. Type the new key and press **Enter**, and then repeat to verify. Make the same change on the RADIUS server.

To change other server properties, specify the server name and the settings you want to change:

```
server set name <name> new-name <name> ip-address <address> auth-port
<number> timeout <seconds> retransmit <number> dead-time <minutes>
```

To remove a server definition:

```
server remove <name>
```

4. You can define up to four server groups, with up to five servers per group (each server can belong to multiple groups). To add a server group name (up to 32 characters):

```
server-group add name <name>
```

The servers in a group are accessed in the order specified. For example, if the first server does not respond, the second server is accessed. To add a RADIUS server to a server group:

```
server-group server add <group-name> <server-name>
```

To change the name of a server group:

```
server-group set name <name> new-name <name>
```

To delete a server group (does not delete the associated servers):

```
server-group remove <name>
```

5. To change the source IP address used in RADIUS packets (defaults to the device's IP address):

```
set client-source <IP address>
```

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## ***configure reduction***

The Reduction command enables you to configure the reduction and assembly engines.

1. To view the current reduction settings:

```
show -run reduction [all | network-sequence-mirroring | pre-sync status]
```

The “all” option includes reductions statistics since the last time the device was reset. For example:

```
===== Reduction Statistics =====
Packets: Total=79837075 - Accept=43450309
Overflow=0 FilterPassthru=75 Default Assembler=0 No Assembler=45723

Reject Protocol=57
Accept Protocol=58040
SR Traffic=1107
Local=36902
Mid Watermark packets=1351
Mid Watermark reached=5
Hi Watermark reached=1
```

The following table describes the reduction settings and statistics:

Keyword	Description
Total	Number of unreduced packets into the device.
Accept	Number of packets into the reduction engine.
Overflow	Packets not reduced because the reduction queue is full (the device is too busy or the WAN link is too slow).
FilterPassthru	Packets not reduced due to application or address filter settings.
Default Assembler	Packets reduced and sent to the default assembler.
No Assembler	Packets not reduced because of no remote WX device.
<b>The following statistics are shown only if they are non-zero.</b>	
Reject Protocol	Packets for IP protocols that are not reduced.
Accept Protocol	Packets for additional IP protocols that are enabled for reduction (does not include TCP and UDP packets, which are reduced by default).
Exclude Address	Packets not reduced due to source/destination filter settings.
TTL Expired	Packets not reduced because the Time to Live value was zero.
Accept Fragmented	Fragmented packets reduced (fragment reduction is enabled by default).
Reject Fragmented	Fragmented packets not reduced (fragment reduction is disabled).
Malformed	Malformed packets not reduced.
SR Traffic	Management packets sent to other WX devices (not reduced).
Local	Packets destined for the local subnet (not reduced).
Mid Watermark packets	Packets that received less reduction processing because the reduction queue exceeded the optimum level (the device is busy or the WAN link is slow).
Mid Watermark reached	Number of times that the reduction queue exceeded the optimum level.
Hi Watermark reached	Number of times the reduction queue became full. Packets received while the queue is full are counted as overflow (not reduced).

2. Type the following command to enter the configure reduction mode:

```
config reduction
```

3. To specify the acceleration features to be used:

```
set topology-features <all | all-but-app | all-but-app-and-afp>
```

For WXOS 5.3, use “all” (the default).

4. To specify the community topology (the default is hub for WXOS 5.3):

```
set topology-type <hub | mesh | spoke>
```

For WXOS 5.3, use only hub or mesh.

To set the range of devices for a hub or mesh topology:

```
set topology-size <range number | max-mem>
```



For WXOS 5.3, use only “5” for hub or “0” for mesh. Table 10 shows the numbered ranges of devices supported by the WX 100 and its client devices (the WX 60 and WXC 500),

**Table 10: Device Ranges by Model, Topology, and Feature Set**

Device	Mesh Ranges – All Features	Hub Ranges – All Features
WXC 500	0 = Up to 10 1 = Up to 15 2 = Up to 20 3 = Up to 25 4 = Up to 30 5 = Up to 40	0 = Up to 10 1 = Up to 15 2 = Up to 20 3 = Up to 22 4 = Up to 32 5 = Up to 50
WX 60	0 = Up to 25 1 = Up to 35 2 = Up to 50 3 = Up to 65 4 = Up to 75 5 = Up to 85	0 = Up to 25 1 = Up to 45 2 = Up to 65 3 = Up to 85 4 = Up to 105 5 = Up to 125
WX 100 (with WX 60 or WXC 500 clients)	When a WX 100 has one or more client devices (stack configuration), the client model type is detected, and the topology ranges displayed in the Web console are for the maximum number of devices when all six clients are connected to the server. For example, if a WX 100 hub using all features has one WXC 500 client, the highest topology range is 300 devices, which is the number of devices supported by six WXC 500 clients (6 x 50).	

5. The following table describes the reduction settings.

Settings	Commands
Reduction and Assembly	<p>You can disable the assembly engine to stop other devices in the community from sending reduced data to this device: To enable or disable the assembly engine (enabled by default):</p> <pre>set assembler &lt;on   off&gt;</pre> <p>To enable or disable data reduction (enabled by default):</p> <pre>set reducer &lt;on   off&gt;</pre>
Endpoints	<p>By default, each WX device attempts to form reduction tunnels with all other devices in the community. To form tunnels with only specific devices:</p> <pre>add assembler-list &lt;IP addresses&gt;</pre> <p>Multiple IP addresses must be separated by spaces. To replace the current list of assemblers with a new list:</p> <pre>set assembler-list &lt;IP addresses&gt;</pre> <p>To specify which devices to form tunnels with (the default is all):</p> <pre>set assembler-mode &lt;all   list   hub-only&gt;</pre> <p>To remove devices from the current list of assemblers:</p> <pre>remove assembler-list &lt;IP addresses&gt;</pre>

Settings	Commands
Network Sequence Caching	<p>On WXC devices, you can enable or disable the use of Network Sequence Caching (NSC) for data reduction. NSC uses disk storage to identify longer patterns of repeated traffic, and retains those patterns for longer periods of time (even when a reduction tunnel is down). NSC is most effective where large files are often sent over the WAN, such as for database backups.</p> <p>To use NSC between two WXC devices, standard reduction must be enabled, and you must enable Active Flow Pipelining (AFP) for the appropriate devices and applications. AFP also requires outbound QoS (refer to “configure acceleration” on page 302).</p> <p>To enable NSC on a WXC device (disabled by default):</p> <pre>network-sequence-mirroring set mode &lt;on   off&gt;</pre> <p>To specify how heavily the disk is used for reduction:</p> <pre>network-sequence-mirroring set disk-access-policy &lt;0-3&gt;</pre> <p>The number (0 to 3) indicates how heavily the disk (and CPU) is used during data reduction. Lower values indicate a higher level of disk access (default is 1). For high-bandwidth links, you may want to increase the value to maximize throughput.</p> <p>Under heavy traffic loads, NSC processing for some types of data may reduce overall throughput. In this case, you can enable NSC overflow mode to allow MSR to take over some processing from NSC (disabled by default). Please contact Technical Support before using this option.</p> <pre>network-sequence-mirroring set overflow-mode &lt;on   off&gt;</pre> <p><b>NSC Endpoints</b></p> <p>To enable NSC for all WXC devices in the community or a specific list of devices (default is all):</p> <pre>network-sequence-mirroring endpoint set mode &lt;all   list&gt;</pre> <p>To add or remove devices from the list enabled for NSC (include a space between IP addresses):</p> <pre>network-sequence-mirroring endpoint add &lt;IP addresses&gt; network-sequence-mirroring endpoint remove &lt;IP addresses&gt;</pre> <p><b>NSC Applications</b></p> <p>To define a list of applications that are included or excluded from NSC (default is excluded):</p> <pre>network-sequence-mirroring application mode &lt;include   exclude&gt;</pre> <p>To add or remove an application from the list that is included or excluded from NSC:</p> <pre>network-sequence-mirroring application add &lt;name&gt; network-sequence-mirroring application remove &lt;name&gt;</pre>
Pre-Synchronization	<p>Large files, such as database files and software updates, can be preloaded on remote NSC-enabled devices. The repeated patterns in the files are added to the reduction dictionaries, so that when a user requests the files, the response time is much faster. The files must be on an FTP server. Be sure to enable NSC for the application that users will access to retrieve the preloaded files. To pre-synchronize a file for a remote NSC-enabled device:</p> <pre>network-sequence-mirroring pre-sync &lt;NSC-device-address&gt; ftp://&lt;host:port&gt;[:&lt;username&gt;:&lt;password&gt;]/&lt;path and file name&gt;</pre> <p>The <i>host</i> is the FTP server name or IP address. If the FTP server allows anonymous access, and the default port (port 21) is used, enter:</p> <pre>network-sequence-mirroring pre-sync &lt;NSC-device-address&gt; ftp://&lt;host&gt;/&lt;path and file name&gt;</pre>
Default assemblers	<p>To create a list of up to six default assemblers (for more information about default assemblers, refer to “Defining Default Assemblers” on page 141):</p> <pre>set def-assembler-list &lt;IP address&gt;</pre> <p>Multiple IP addresses must be separated by a space.</p> <p>To add subnets to be excluded from the default assemblers:</p> <pre>add excl-subnet-list &lt;IP address&gt;/&lt;subnet mask&gt;</pre> <p>To remove all or specific subnets from the exclude list:</p> <pre>remove excl-subnet-list &lt;all   IP address&gt;/&lt;subnet mask&gt;</pre>

Settings	Commands
Preferred assemblers	<p>When two or more WX devices in a community can reach a single subnet, and no other policies apply, traffic is routed to each device on an arbitrary basis. To use a specific device when there is more than one path, you can specify that device as a preferred assembler. To designate one or more (up to 80) preferred assemblers:</p> <pre>set pref-assembler-list &lt;space separated list of WX IP addresses&gt;</pre>
Load balancing	<p>The load balancing policy enables two or more WX devices to share the transmission of reduced data to a common destination with equal cost paths. To enable or disable load balancing:</p> <pre>set lb-policy &lt;off   per-packet   per-destination   per-flow&gt;</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>■ <b>off.</b> All traffic is routed to one of the available tunnels. No load balancing (default).</li> <li>■ <b>per-packet.</b> Traffic is distributed on a per-packet basis (round robin).</li> <li>■ <b>per-destination.</b> Traffic is distributed based on destination IP address.</li> <li>■ <b>per-flow.</b> Traffic is distributed based on source and destination IP addresses and ports.</li> </ul>
Tunnel mode	<p>To change how traffic is sent in a reduction tunnel to a remote device:</p> <pre>set tunnelmode &lt;udp   multi-flow   visibility   ipcomp&gt;</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>■ <b>udp.</b> Uses UDP (port 3577) to send meta packets as a single traffic flow.</li> <li>■ <b>multi-flow.</b> 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. The default maximum number of flows is 256 (used to allocate resources—not a hard limit). To change the default: <pre>set max-flows &lt;integer between 256 and 1024&gt;</pre> </li> <li>■ <b>visibility.</b> Uses UDP and preserves the source and destination ports of all packets so that performance monitoring tools can identify the various devices responsible for the traffic in the reduction tunnel. Verify that your tools are configured to monitor UDP traffic.</li> <li>■ <b>ipcomp.</b> Uses the IP payload compression protocol (protocol number 108) to send meta packets as a single traffic flow. Provides optimum reduction in most environments (default).</li> </ul>
Tunnel switching	<p>To enable tunnel switching on selected devices, such as to send reduced traffic between communities (disabled by default):</p> <pre>set tunnel-switching {on   off}</pre> <p><b>NOTE:</b> This feature must be implemented carefully to avoid unnecessary compression. For more information, refer to “Configuring Tunnel Switching” on page 148. Tunnel switching cannot be used on hubs where the topology setting specifies the maximum range of devices.</p>
Heartbeat packets	<p>By default, UDP keep-alive “heartbeat” packets are sent every 5 seconds to confirm the operability of the reduction tunnels between WX devices. To change the heartbeat frequency:</p> <pre>set heartbeat-frequency &lt;1-300&gt;</pre> <p><b>NOTE:</b> All WX devices in the same community must have the same heartbeat frequency.</p> <p>If a device fails to respond to four consecutive heartbeats, the other WX devices stop reducing data for the device (passthrough mode). If 10 consecutive heartbeats get no response, the other devices disable their reduction tunnels and attempt to reestablish the tunnel as follows: every three minutes for the first hour, every 15 minutes for the second hour, every hour for the next 22 hours, and once a day thereafter.</p> <p>To change the number of missed heartbeats that stops reduction and disconnects tunnels:</p> <pre>set heartbeat-misses passthru &lt;number   default&gt; disconnect &lt;number   default&gt;</pre> <p>The number of missed heartbeats allowed is higher for remote endpoints for which Active Flow Pipelining or Forward Error Correction is enabled (refer to “configure acceleration” on page 302).</p>

Settings	Commands																		
Heartbeat ToS/DSCP values	<p>The UDP tunnel keep-alive packets sent between WX devices have normal priority (zero) and may be dropped in heavily congested networks. To change the ToS/DSCP value (0 to 255) for these packets:</p> <pre>set tos-bit &lt;0-255&gt;</pre> <p>To set a ToS IP precedence value (0 to 7), enter a CLI value that sets the upper three bits (bits 6, 7, and 8) of the ToS/DSCP byte:</p> <table> <tr> <th>IP Precedence</th><th>CLI value</th></tr> <tr> <td>0</td><td>0</td></tr> <tr> <td>1</td><td>32</td></tr> <tr> <td>2</td><td>64</td></tr> <tr> <td>3</td><td>96</td></tr> <tr> <td>4</td><td>128</td></tr> <tr> <td>5</td><td>160</td></tr> <tr> <td>6</td><td>192</td></tr> <tr> <td>7</td><td>224</td></tr> </table> <p>To set a ToS type-of-service value (0 to 15), enter a CLI value that sets bits 2 through 5. For example, a CLI value of 2 equals a type-of-service value of 1.</p> <p>To set a DSCP value (0 to 63), enter a CLI value that sets the upper six bits of the ToS/DSCP byte. For example, a CLI value of 4 equals a DSCP value of 1.</p> <p><b>NOTE:</b> These values are overridden by the IP precedence or DSCP settings defined for Multi-Path (refer to “Configuring Policy-Based Multi-Path” on page 115).</p>	IP Precedence	CLI value	0	0	1	32	2	64	3	96	4	128	5	160	6	192	7	224
IP Precedence	CLI value																		
0	0																		
1	32																		
2	64																		
3	96																		
4	128																		
5	160																		
6	192																		
7	224																		
LAN/WAN check	<p>The LAN-WAN check is an important safety feature that helps prevent routing configuration errors. However, if the default gateway is on the LAN side of the WX device, or if you want to allow reduction tunnels on the LAN side of the device (such as for tunnel switching), you must disable the LAN-WAN check.</p> <pre>set lan-wan-check &lt;on   off&gt;</pre>																		
Dynamic Resource Allocation (DRA)	<p>DRA enhances reduction on low-speed WAN links (such as 128 Kbps). During good network conditions (such as low CPU load), the WX device attempts to further reduce the data without compromising latency or packet loss. To enable or disable DRA (enabled by default):</p> <pre>set modes DRA &lt;on   off&gt;</pre> <p><b>NOTE:</b> It is strongly recommended that you enable outbound QoS and specify the WAN circuit speed for each remote WX device. For more information, refer to “configure qos outbound” on page 337.</p>																		
Fast reduction tunnels	<p>Fast reduction tunnels increase throughput to remote WX devices by decreasing the resources devoted to MSR reduction (NSC is not supported). This feature may lower reduction percentages, but it allows Application Flow Acceleration, AFP, QoS, and Multi-Path to exceed the licensed speed of the device.</p> <p>The local WX device must be licensed at 20 Mbps or higher (the local or remote device cannot be an WX 15, WX 20, or WXC 250). To enable fast reduction tunnels for one or more remote endpoints (multiple IP addresses must be separated by spaces):</p> <pre>set fast-reduction-tunnel &lt;IP-addresses&gt;   none</pre> <p>Specify “none” to disable all fast reduction tunnels.</p> <p><b>NOTE:</b> On a WXC device, do not enable fast reduction tunnels to remote WXC devices. Throughput may be reduced when NSC is enabled.</p>																		

Settings	Commands
MSR symbol size	<p>The symbol size is the number of bytes that the Molecular Sequence Reduction (MSR) algorithm analyzes at one time to discover repeated traffic patterns. In general, larger symbol sizes require less processing, but achieve lower data reduction rates. The default symbol size depends on the licensed device speed (eight for 20 Mbps or higher, four for lower speeds).</p> <p>To change the symbol size:</p> <pre>set modes msr &lt;number&gt;</pre> <p>The valid symbol sizes are: 1-8, 10, 12, 16, 32, and 64 (use -1 to restore the default). On an WX 15, do not exceed a symbol size of 8.</p>
Meta-packets	<p>Multiple packets of reduced data are encapsulated in “meta” packets of up to 1500 bytes. If a device on the WAN side of the WX device adds to the packet (such as a VPN device), you can reduce the maximum meta packet size to avoid packet fragmentation by the router. Before you adjust the maximum meta packet size, verify the approximate number of bytes that are added by the network device.</p> <p>To set the maximum meta packet size:</p> <pre>set max-meta-pkt-size &lt;number between 576 and 1500&gt;</pre> <p>By default, the amount of time each meta-packet of reduced data is held is based on the round-trip time (RTT) to the destination device. To change the meta-packet wait time:</p> <pre>set meta-packet-wait &lt;mode&gt;</pre> <p>Where &lt; mode &gt; is “default”, “absolute-time”, or “rtt”.</p> <p>If you enter “absolute-time” as the mode, enter an amount of time (in 2 ms increments) for the meta-packet to wait before transmitting across the network. By default, this setting is 8 ms. For example,</p> <pre>set meta-packet-wait absolute-time 4</pre> <p>If you enter “rtt” as the mode, enter a percentage number that will be calculated by the RTT and used to hold the meta-packet before being transmitted across the network. For example,</p> <pre>set meta-packet-wait rtt percent-rtt 15</pre> <p>In addition to the RTT percentage, you can set an upper limit for which a packet will be held (in 2 ms. increments). The default is 8 ms. For example,</p> <pre>set meta-packet-wait rtt percent-rtt 15 limit 20</pre> <p>In extreme latency-sensitive networks, you can disable the grouping of reduced packets into meta packets so that reduced data is sent on a per-packet basis.</p> <pre>set multi-packet off</pre>
Policy routes	<p>Policy routes let you vary the default gateway used for reduced traffic based on the application. To add a policy route:</p> <pre>policy-route add &lt;application name&gt; &lt;gateway_IP_address&gt;</pre> <p>To remove a policy route:</p> <pre>policy-route remove &lt;application name&gt;</pre>

Settings	Commands
MAC addresses	<p>By default, the source hardware (MAC) address of an assembled packet is the MAC address of the WX device that assembled the packet. To change the source MAC address of assembled packets:</p> <pre>set assembly-source-mac-mode [default   copy-source   user-defined &lt;mac&gt;]</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>■ <b>default</b>. Uses the MAC address of the WX assembler.</li> <li>■ <b>copy-source</b>. Uses the source MAC address received in the reduced packet.</li> <li>■ <b>user-defined &lt;mac&gt;</b>. Specify a MAC address (the format is <code>xx:xx:xx:xx:xx:xx</code>).</li> </ul>
Reduction tradeoff for speed	<p>Under high traffic loads, reduction is scaled back to increase throughput. To specify the relative tradeoff of reduction for speed ("default" indicates the "standard" tradeoff):</p> <pre>set reduction-tradeoff-for-speed &lt;minimum   standard   maximum   default&gt;</pre> <p>Use "minimum" to ensure optimum reduction under all traffic loads. On high-speed WAN links (over 20 Mbps), "maximum" is recommended for optimum throughput.</p>

- To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure reduction-subnet**

Reduction subnets are the subnets on the LAN side of the WX device that you can selectively advertise to the other WX 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 WX device is installed. To identify more LAN-side subnets, you can:

- Add static routes manually (refer to "configure route" on page 356)
- Add dynamic routes using one of the following methods:
  - Enable the Open Shortest Path First (OSPF) and/or the Routing Information Protocol (RIPv1, RIPv2), as described in "configure ospf" on page 330 and "configure rip" on page 356
  - Periodically poll the routing table of a Cisco router (refer to "configure route-poll" on page 359)
  - Import a file of routes from an FTP server (refer to "import-route-table" on page 291)
- Enable the WAN reduction subnet option to include routes discovered on the Remote interface. In some environments, local routes may be discovered on the WAN side of the WX device.

The set of subnets advertised by each device is called a "netmap." By default, only the subnets you specify are advertised. You can enable the advertisement of all subnets or just selected subnets. To advertise specific subnets, you can create an Enabled list and a Disabled list of local IP subnets, and then set the mode for the lists to All, Include, or Exclude.

For example, if you have five subnets in the Enabled list and one subnet in the Disabled list, and the mode is set to “Include,” only the subnets in the Enabled list are advertised. If the mode is set to “Exclude,” only the subnet in the Exclude list is advertised. If the mode is set to “All,” all subnets are advertised and the lists are ignored.

1. To view the current reduction subnets and subnet settings:

```
show -run reduction-subnet
```

```
Mode: include
```

```
Wan-reduction-subnet Mode: off
```

Destination	Netmask	Cost	Enabled	Interface
192.168.0.0	255.255.255.0	1	no	Local

The Enabled column indicates whether the subnet is advertised.

2. Type the following command to enter the configure reduction subnet mode:

```
config reduction-subnet
```

3. To add entries to the Enabled list of reduction subnets:

```
add enable <IP address/subnet mask>
```

To add entries to the Disabled list of reduction subnets:

```
add disable <IP address/subnet mask>
```

If a subnet is on both the Enabled and Disabled lists, the subnet is disabled.

To remove entries from the Enabled or Disabled lists:

```
remove enable <IP address/subnet mask>
```

```
remove disable <IP address/subnet mask>
```

To set the reduction mode (the default is “include”, which advertises subnets on the Enabled list):

```
set mode <all | include | exclude>
```

4. By default, only routes discovered on the LAN side of the WX device (the Local interface) can be advertised as reduction subnets. For example, in VLAN environments, some LAN-side routes can be discovered only on the WAN side.

When the WX 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 page and cannot be advertised for reduction.

To include routes discovered on the Remote interface as potential reduction subnets:

```
set wan-reduction-subnet on
```

This option is enabled by default if the WX device is installed off-path (refer to “configure packet-interception” on page 331).



**NOTE:** Allow up to one minute for the remote routes to be added to the list of reduction subnets. After the routes are added, be careful to advertise only the LAN-side routes. WAN-side subnets are excluded if their next hop is the default gateway.

5. Each WX device dynamically adjusts its advertised subnets to exclude (carve out) any hosts or gateways that become unreachable. To enable or disable this feature (enabled by default):

```
set carveout <on | off>
```

6. To verify the mode and reduction subnets before committing these changes:

```
show reduction-subnet
```

7. To apply the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.



**NOTE:** If you disable an advertised subnet, you must reboot the device for the change to take effect.

## configure reg-server

When you install a WX device, you must designate one device as the registration server. The registration server stores network information for all devices that report to it, and identifies a community for each device. Every WX device periodically contacts the registration server to obtain information about other devices in the community. Initially, all WX devices are in the Default community.

Data reduction can occur only between devices in the same community. You can define separate communities to control how data is reduced, and you can add the same device to multiple communities for backup and redundancy. A device that belongs to multiple communities can reduce and assemble data for the devices in all of its communities.

If you are logged in to the registration server, you can change the password of the registration server, or designate a different WX device as the registration server. You can also add and change communities, and assign a secondary registration server to act as a backup when the primary registration server is not available.

1. To view the current registration server settings (the communities are shown only if the device is a registration server):

```
show -run reg-server
```

```
Registration server: 192.168.55.22
Secondary registration server: not set
This system is currently the registration server
Connection timeout (seconds): 2
Connection retry count: 1
```



## 2 Communities

Community "default-192.168.55.22" has 0 entries:

Community "Main" has 4 entries:

192.168.52.22 192.168.53.22 192.168.54.22 192.168.55.22

2. On a registration server, enter the following command to list the registered WX devices:

```
show -run reg-summary
```

To view the details for all registered devices, a specific device, or just the reducers or assemblers:

```
show -run reg-detail [<IP address> | -assemblers | -reducers]
```

Number of registered nodes: 4

Number of reducers: 4

Number of assemblers: 4

Node list:

IP-Address	Type	Duty	Proto	SW-Ver	Errors	Last-Register-Time	Name
192.168.52.22	SA/SR		0	4	0	JAN 07 13:07:30 2005	52/22-SR20
192.168.53.22	SA/SR		0	7	0	JAN 07 09:56:43 2005	53/22-SR80
192.168.54.22	SA/SR		0	6	0	JAN 07 14:41:21 2005	54/22-SR15
192.168.55.22	SA/SR	R	0	7	0	JAN 07 09:51:42 2005	55/22-SR100

Key for 'Duty': H=Hub R=RegServer S=SecondaryRegServer

Key for 'Type': SA=Sequence Assembler SR=Sequence Reducer

The **Proto** and **SW-Ver** columns identify the registration protocol for each device (internal use only). The **Errors** indicate the number of times that the server failed to propagate registration updates to a device.

To reset all the error counts to zero:

```
config reg-server clear-error-count
```



**NOTE:** Each device obtains all the latest registration information, including any missed updates, when it checks in with the registration server (every eight hours).

3. Type the following command to enter the configure registration server mode:

```
config reg-server
```

4. The following table describes the registration server settings.

Settings	Commands
Primary and secondary servers	<p>To specify the IP address of the registration server:</p> <pre>set ip-address &lt;registration server IP address&gt;</pre> <p>If this device is not the registration server, enter the IP address of the current (or future) registration server. If you have not yet configured the registration server, enter the future IP address of the registration server and specify the default password, "peribit".</p> <p>To specify the registration server password, enter the following command, and then enter and confirm the password at the prompts (also applies to the secondary registration server):</p> <pre>set password</pre> <p><b>NOTE:</b> Changing the password disrupts communication with all WX devices that use the registration server. To restore communication with the registration server, you must update the registration server password on each WX device.</p> <p>To specify a secondary registration server to act as a backup when the primary registration server is not available:</p> <pre>set sec-ip-address &lt;secondary registration server IP address&gt;</pre>
Timers	<p>To specify how often a device attempts to check in with the registration server (the default is every 8 hours):</p> <pre>set registration-frequency &lt;3-hours   8-hours   24-hours   7-days   once-only&gt;</pre> <p>To specify the number of times the WX device attempts to access the primary registration server before switching to the secondary (default is 1):</p> <pre>set connect-retries &lt;1-5&gt;</pre> <p>On a registration server, the retry count is also the number of times that the server attempts to send registration updates to a device.</p> <p>To specify the number of seconds between retries (the default is 2):</p> <pre>set connect-timeout &lt;2-60&gt;</pre> <p>To specify the number of days before a device is purged if it has not checked in (default is 1):</p> <pre>set ageout-time &lt;days&gt;</pre>
Communities	<p>To add the name of a new community (up to 31 characters):</p> <pre>community add &lt;name&gt;</pre> <p>To change a community name:</p> <pre>community set name &lt;old name&gt; &lt;new name&gt;</pre> <p>To add a WX device to a community:</p> <pre>community remote-sr add &lt;community&gt; &lt;IP address&gt;</pre> <p>To delete a WX device from a community:</p> <pre>community remote-sr remove &lt;community&gt; &lt;IP address&gt;</pre> <p>To delete a community:</p> <pre>community remove &lt;name&gt;</pre> <p>If you delete a community or remove devices from a community, the devices are moved to the Default community if they do not belong to any other user-defined communities.</p>
Database	<p>If you add or change communities on a secondary registration server, you can export the community database to the primary server:</p> <pre>community export-to-primary</pre> <p>To prevent the primary registration server from overriding the community database on the secondary server, enter the following command on the secondary server:</p> <pre>community set disable-import</pre>

5. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure remote-routes**

Remote routes are the reduction subnets advertised by other WX devices in the community. You can specify how often remote routes are fetched from other devices, and whether each remote route is validated.

1. To view the current remote routes and settings:

```
show -run remote-routes
```



**NOTE:** Each WX device dynamically adjusts its advertised subnets to exclude unreachable addresses. In this case, multiple remote routes must be advertised for the same subnet to exclude unresponsive addresses.

2. Type the following command to enter the configure remote route mode:

```
config remote-routes
```

3. To validate the remote routes advertised by other WX devices, you can enable route validation. 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 WX 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.

To enable or disable remote route validation (disabled by default):

```
set validation <on | off>
```



**NOTE:** Enable route validation only if the validity of the remote routes is in question. You should not use this option if load balancing is enabled.

4. To specify how often remote routes are fetched from the other WX devices in the community (the default is every 3600 seconds):

```
set frequency <once | 3600 | 7200 | 10800 | 86400>
```

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

5. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure rip**

If your network uses RIP, you can configure the WX device to use RIP to dynamically discover routes on both the Local and Remote interfaces. In this case, the WX device only receives routes, it does not send them. Off-path WX devices can be configured to both send and receive RIP routes (refer to “configure packet-interception” on page 331).

1. To view the current RIP configuration:

```
show -run rip
```

2. Type the following command to enter the configure RIP mode:

```
config rip
```

3. To specify the number of seconds before a route is aged out (the default is 300):

```
set ageout <1 - 8400>
```

4. To specify whether a RIP password is used in your network (default is none):

```
set auth-type <password | none>
```

To specify the RIP password (up to 15 characters):

```
set password <password>
```

5. To specify whether RIP version 1 or 2 is used (the default is 2)):

```
set version <1 | 2>
```

6. To enable or disable RIP (disabled by default):

```
set rip <on | off>
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure route**

When you first install a WX device, its routing table contains the local subnet where the device is installed, a route to the default gateway (the default route), and the loopback address. Use the Route command to add static routes to the routing table, enable router load balancing, and specify the ICMP age-out interval. A total of 8192 IP routes (static and dynamic) are supported (the WX 15 is limited to 1000).

1. To view the current routes and route settings (all routes are shown by default):

```
show -run route [protocol <ospf | rip | static>] [subnet <subnet/mask>]
```

2. Type the following command to enter the configure route mode:

```
config route
```

3. To add a new static route:

```
add <IP address> mask <subnet mask> gateway <gateway IP address> cost
<cost>
```

Use dotted-decimal notation for the IP address, mask, and gateway. The <cost> is an optional value from 0 to 65535. The default is 1000.

To delete a static route:

```
delete <IP address> mask <mask>
```

4. To set the precedence between static and dynamic routes (dynamic routes take precedence by default):


```
set precedence <static | dynamic>
```

5. To specify the number of minutes (1 to 143165) that routes redirected by ICMP are retained before being aged out (the default is 10):

```
set icmp-redirect-ageout <number>
```

6. **Route-based load balancing.** You can configure the WX device to distribute traffic for equal-cost paths across up to four different gateways. In Figure 168, 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 168: Load Balancing**



To specify the route-based load balancing policy:

```
set lb-policy <off | per-packet | per-destination | per-flow>
```

Where:

- **off.** Traffic is routed to any one of the available routers (default).
- **per-packet.** Traffic is distributed 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.

- **per-destination.** Traffic is distributed based on destination IP address.

- **per-flow.** Traffic is distributed based on source and destination IP addresses and ports.
7. **ToS marking for router-based load balancing.** You can configure the local router(s) to distribute reduced traffic based on the ToS type-of-service values set by the WX device. The type-of-service values (0 to 15) are set in bits 2 through 5 of the ToS/DSCP field. This method can be used together with route-based load balancing.



**NOTE:** You cannot use ToS marking for router-based balancing if DSCP values are set by Multi-Path or outbound QoS. However, the ToS IP precedence values set by these features do not interfere with the type-of-service values defined here.

To configure ToS marking for router-based balancing:

- a. Enable ToS marking for load balancing (disabled by default):

```
rtr-based-lb set mode <off | type-of-service>
```

- b. Specify two or more (up to 16) type-of-service values (0 to 15), separated by spaces:

```
rtr-based-lb set tos <0 - 15>
```

- c. Specify the load balancing policy (default is per destination):

```
rtr-based-lb set lb-policy <per-packet | per-destination | per-flow>
```

Where:

- **per-packet.** ToS values are assigned to each reduced meta-packet in round robin fashion.
- **per-destination.** ToS values are assigned to each reduced meta-packet by applying a hash function to the destination IP address.
- **per-flow.** ToS values are assigned to each reduced meta-packet by applying a hash function to the source and destination IP addresses and ports.

Note that per-flow and per-destination router balancing will lower the percentage of data reduction because each meta-packet contains traffic for only one flow or destination.

- d. Configure the router(s) to distribute the meta-packets based on the type-of-service values. On the inbound interface from the WX device, define a route map for router-balancing:

```
interface FastEthernet1/0
ip address 10.129.30.5 255.255.255.0
ip policy route-map router-balance
```

Define access lists that specify each ToS value set by the WX device:

```
access-list 101 permit ip 10.129.30.1 0.0.0.0 0.0.0.0 255.255.255.255 tos 10
access-list 102 permit ip 10.129.30.1 0.0.0.0 0.0.0.0 255.255.255.255 tos 11
```

Match the ToS values with the appropriate next-hop gateways:

```
route-map router-balance permit 10
match ip address 101
set ip next-hop 10.129.20.1
```

```
route-map router-balance permit 20
match ip address 102
set ip next-hop 10.129.50.1
```

8. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure route-poll**

The WX device can obtain dynamic routes by periodically polling a Cisco router. The Cisco router must be configured to allow Remote Shell Protocol (*rsh*) access by the WX device. The *rsh* protocol allows a user or device to execute commands on a remote system without having to log in. For more information on enabling *rsh* on your Cisco router, refer to the Cisco IOS documentation.



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

---

1. To view the current route poll settings:

```
show -run route-poll
```

2. Type the following command to enter the configure route poll mode:

```
config route-poll
```

3. To set the IP address of the Cisco router:

```
set remote-host <IP address>
```

To specify the port number (1 to 1024) defined on the router (the default is 514):

```
set remote-port <1-1024>
```

To set the local user name to match the remote user name defined on the router:

```
set local-user <user name>
```

To set the remote user name to match the local user name defined on the router:

```
set remote-user <user name>
```

To enable or disable route polling (disabled by default):

```
set mode <rsh | none>
```

4. To specify the IP address of a secondary Cisco router to be used when the primary is not available:

```
set sec-remote-host <IP address>
```

To specify the port number (1 to 1024) on the secondary router (the default is 514):

```
set sec remote-port <1-1024>
```

5. To enable or disable the extraction of BGP routes (disabled by default):

```
set allow-bgp-routes <on | off>
```

6. To change the polling frequency (default is every five minutes):

```
set frequency <number of minutes>
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure security**

The Security command can be used to restrict access to the WX device by IP address, change the packet capture password, lock the front-panel keypad, and disable the Web console and/or the SSH interface.

To define users locally and specify how users are authenticated (locally and/or through RADIUS), refer to “configure aaa” on page 300. To define the RADIUS servers, refer to “configure radius” on page 342.

1. To view the current security settings:

```
show -run security
```

2. Type the following command to enter the configure security mode:

```
config security
```

3. To change the packet-capture password:

```
set packet-capture
```

and then press Enter. Type the current password (the default is “peribit”) and press Enter. Next, type the new password and press Enter.

4. To enable or disable the lock on the front panel keypad on devices that have an LCD (disabled by default):

```
set front-panel <on | off>
```

5. To enable or disable the Web console (enabled by default):

```
set web <on | off | cmsonly>
```

Use **cmsonly** to allow Web access only through the Central Management System (CMS).



To enable or disable the SSH interface (enabled by default):

```
set ssh <on | off >
```

6. To restrict operator access, you can create lists of IP addresses or subnets that are allowed or denied access to the WX device. If you enter one allowed IP address, users can log in only from the specified address.

To add an IP address or subnet that is allowed access to this device:

```
add allow-ip-address <IP-address>[/<subnet mask>]
```

Multiple IP addresses must be separated by spaces.

To add an IP address or subnet that is denied access to this device:

```
add deny-ip-address <IP-address>[/<subnet mask>]
```

To remove one or all IP addresses that have access to the device:

```
remove allow-ip-address {all | <IP-address>[/<subnet mask>]}
```

To remove one or all IP addresses that are denied access to the device:

```
remove deny-ip-address {all | <IP-address>[/<subnet mask>]}
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## ***configure snmp***

WX devices support SNMP, the Management Information Base (MIB) II public objects, and private MIB objects. Your Network Management System (NMS) can use the private MIB to monitor the performance of the WX devices in your network. In addition, enabling SNMP traps on a WX device allows the device to send traps and alarms to the NMS as they occur.

1. To view the current SNMP settings:

```
show -run snmp
```

2. Type the following command to enter the configure snmp mode:

```
config snmp
```

3. To enable or disable support for SNMP (enabled by default):

```
set snmp <on | off>
```

4. To enter read and write community strings:

```
set read-community <string>
```

```
set write-community <string>
```

If the string value has spaces, enclose it in double quotation marks.

5. To enable or disable SNMP traps (disabled by default):

```
set trap <on | off>
```

To enter a trap community string:

```
set trap-community <string>
```

To enable or disable traps for authentication failures (disabled by default):

```
set auth-failure-trap <on | off>
```

To enter a trap destination:

```
set trap-destination <IP address>
```

Multiple trap destinations must be separated by spaces.




---

**NOTE:** For a description of each trap, refer to “SNMP Traps and Syslog Messages” on page 397.

---

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## **configure sntp**

WX devices support the Simple Network Time Protocol (SNTP). An SNTP server provides a common time base for devices within your network. If your network does not use an SNTP server, you can manually configure the time settings for each WX device (refer to “configure clock” on page 315).




---

**NOTE:** Before enabling SNTP, use the “show clock” command to verify that the time zone settings are correct. If necessary, use the “configure clock” command to change the time zone settings.

---

To enable SNTP on this device:

1. To view the current SNTP settings:

```
show -run sntp
```

2. Type the following command to enter the configure SNTP mode:

```
config sntp
```

3. To set the SNTP server address:

```
set ip-address <IP address>
```

4. To specify the number of minutes between updates from the time server (the default is 1440):

```
set interval <number>
```

5. To add a secondary SNTP server to be used if the primary is not available:

```
set sec-ip-address <IP address>
```

6. To enable or disable SNTP (disabled by default):

```
set sntp <on | off>
```

7. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure stack-group**

The WX 100 can act as a server to distribute the processing load to a “stack” of up to six client WX devices. The client devices are connected directly to the WX 100, and all clients must be the same device type: WX 55s, WX 60s, WX 80s, or WXC 500s (refer to “Connecting Client Devices to the Server” on page 43). Client mode must be enabled manually on each WX device connected to a WX 100. For WXC 500 clients, the WX 100 must be configured as a WXC server.

You can increase throughput by configuring up to six reduction tunnels between two WX 100 servers with the same number of clients (one tunnel for each client). You can also disable/enable the WX 100 interface for one or all clients.

1. To view the stack configuration on a WX 100:

```
show -run stack-group
```

2. To enable or disable client mode on a client device (disabled by default):

```
config stack-group set client-mode <on | off>
```

If the client devices are WXC 500s, on the WX 100 you must enable support for WXC devices (disabled by default):

```
config stack-group set sequence-mirror-server <on | off>
```

3. To enable up to 6 reduction tunnels between two WX 100 servers with the same number of clients, specify the same number of tunnels (one for each client) on both WX 100 servers:

```
config stack-group endpoint set ip-address <IP address> max-tunnels <1-6>
```



**NOTE:** Be sure to disable congestion control on the WX 100 (congestion control reduces throughput if multiple tunnels are enabled). Also, configuring Policy-Based Multi-Path for a remote endpoint overrides multiple reduction tunnels. After Multi-Path is configured, the command to establish multiple tunnels is ignored.

To disable multiple reduction tunnels, set the maximum tunnels to one:

```
config stack-group endpoint set ip-address <IP address> max-tunnels 1
```

- By default, in WXOS 5.3 tunnels are hosted only on the clients (not on the WX 100). To run the WX 100 in standalone mode (no clients), you must roll back WXOS 5.3 to the previous release. The following CLI commands to allow tunnels to be hosted on the WX 100 server or both the server and the clients are not supported in WXOS 5.3:

```
config stack-group set host-session server-only
config stack-group set host-session all
```

- To specify the maximum number of endpoints that are monitored for the WAN Throughput and WAN Application Summary reports (default is 320):

```
config stack-group set monitor-endpoints <320 | 1000>
reboot
```

You must reboot the device for the change to take effect. Note that the virtual endpoints monitored on these reports do not count against the maximum.

- To disable or enable the WX 100 interface to one or all clients (0 indicates all clients):

```
config stack-group interface disable <0-6>
config stack-group interface enable <0-6>
```

## **configure syslog**

WX devices can send Syslog messages to one or more Syslog servers. A Syslog server allows you to centrally log and analyze configuration events and system error messages such as interface status, security alerts, and environmental conditions.

- To view the current Syslog settings:

```
show -run syslog
```

- Type the following command to enter the configure Syslog mode:

```
config syslog
```

- To enable or disable Syslog (disabled by default):

```
set syslog <on | off>
```

- To enter the IP address of a Syslog server:

```
set destination <IP address>
```

Multiple Syslog addresses must be separated by spaces. Up to five Syslog servers can be defined.

- To set the severity of messages uploaded to the Syslog server, specify any combination of “c”, “e”, and “i”:

```
set severity <cei>
```

Do not include spaces between the letters. The letters indicate:

- **Critical.** Critical error messages about software or hardware malfunctions.

- **Error.** Error message, such as License expired.
- **Information.** Informational messages, such as reload requests and low-process stack messages.



**NOTE:** For a description of Syslog messages, refer to “SNMP Traps and Syslog Messages” on page 397.

6. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

### **configure top-talker**

Traffic data is collected continuously for the most active traffic flows, including the application name and protocol, the 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 Web console. Undefined application flows displayed in the Web console are flagged so that you can quickly populate application definitions with the correct addresses and ports.



**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 using the same protocol.

The Traffic utility maintains the 65,000 most active flows. You can view the top 50 flows in the Web console, but the complete list can be exported to a file in CSV format.

1. To view the current Top Talker settings:

```
show -run top-talker
```

2. Type the following to enter the config Top Talker mode:

```
config top-talker
```

3. To export the statistics file to CSV format:

```
export <ftp://<IP address>[:<user>:<pass>]/<path>
```

or

```
export <ftp://<IP address>/<path>
```

To delete the collected data:

```
delete
```

4. To send traffic data to a Cisco NetFlow server:

- a. Specify the IP address and UDP port number of the NetFlow server:

```
netflow set ip-address <IP address> udp-port <number>
```

- b. Enable NetFlow data collection (disabled by default);

netflow mode <on | off>

### **configure wan-performance-monitor**

You can enable WAN performance monitoring to measure the latency and loss between the current device and 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 (refer to “WAN Performance Statistics” on page 231). 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.

WAN traffic monitoring must be enabled to view the WAN reports (refer to “configure mon-apps” on page 325). Note that 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. The WAN performance settings will take effect if Multi-Path is disabled (refer to “configure multi-path” on page 326).

1. To view the current settings:

show -run wan-performance-monitor

2. Type the following command to enter the configure WAN performance monitor mode:

config wan-performance-monitor

3. To enable or disable WAN performance monitoring for ALL remote endpoints (disabled by default):

set mode <on | off>

To change the default loss and latency thresholds for ALL remote endpoints:

set latency-threshold <20-5000> probes-per-minute <1-60>  
 probes-above-latency <1-60> probes-lost <1-60> minutes-to-bad-latency <1-32>  
 minutes-to-bad-loss <1-32> minutes-to-good-latency <1-32>  
 minutes-to-good-loss <1-32>

Where:

- **latency-threshold <20-5000>** . Latency threshold in milliseconds (default is 5000).
- **probes-per-minute <1-60>** . Number of times per minute that each path is tested (default is 12).
- **probes-lost <1-60>** . Number of probes that must be lost per minute before the minute is marked as above the loss threshold (default is 2).

- **minutes-to-bad-latency <1-32>** . Number of consecutive minutes that the median latency must exceed the latency threshold before a WAN performance “latency failure” trap and Syslog entry are generated, and an event is shown on the WAN Performance report (default is 4).
- **minutes-to-bad-loss <1-32>** . Number of consecutive minutes that must exceed the loss threshold before a WAN performance “loss failure” trap and Syslog entry are generated, and an event is shown on the WAN Performance report (default is 4).
- **minutes-to-good-latency <1-32>** . Number of consecutive minutes of acceptable latency required before a WAN performance “active” trap and Syslog entry are generated, and an event is shown on the WAN Performance report (default is 4).
- **minutes-to-good-loss <1-32>** . Number of consecutive minutes of acceptable loss required before a WAN performance “active” trap and Syslog entry are generated, and an event is shown on the WAN Performance report (default is 4).

To restore the global default loss and latency thresholds for ALL remote endpoints:

**set-default**

4. To enable or disable WAN performance monitoring for a specific remote endpoint:

```
endpoint add ip-address <address> [latency-threshold <20-5000>]
[probes-per-minute <1-60>] [probes-above-latency <1-60>] [probes-lost <1-60>]
[minutes-to-bad-latency <1-32>] [minutes-to-bad-loss <1-32>]
[minutes-to-good-latency <1-32>] [minutes-to-good-loss <1-32>]
```

To change the settings for a specific remote endpoint:

```
endpoint set ip-address <address> [latency-threshold <20-5000>]
[probes-per-minute <1-60>] [probes-above-latency <1-60>] [probes-lost <1-60>]
[minutes-to-bad-latency <1-32>] [minutes-to-bad-loss <1-32>]
[minutes-to-good-latency <1-32>] [minutes-to-good-loss <1-32>]
```

To restore the default loss and latency thresholds for an endpoint:

**endpoint set-global ip-address <address>**

5. To commit the changes to the running configuration, exit configuration mode and type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

## Show Commands

---

### ***show aaa***

To view the current AAA settings:

```
show -run aaa
```

```
Console authentication methods: local
SSH authentication methods: local
Web authentication methods: local
Authorization mode: off
SSH login retries allowed before disconnect: 3
```

```
User Name  Privilege Level  Idle Timeout (seconds)
admin      read-write  1800
```

All users have read-write privileges and a 30 minute idle timeout because the AAA authorization mode is "off".

### ***show acceleration***

To view the current acceleration configuration and status, use the following commands:

```
show -run acceleration application <cifs | exchange | http> <configuration | status>
```

```
show -run acceleration cluster <configuration | status>>
```

```
show -run acceleration packet-flow <configuration | status>
```

Where:

- **application <cifs | exchange | http> <configuration | status>** . Lists the configuration or status for CIFS, Exchange, or HTTP application acceleration.
- The **CIFS** status shows the current number of active flows, passive flows, and number of files being tracked, along with several totals since the device was last reset, such as the total number of CIFS flows, the total reads and writes, and the number of reads and writes accelerated. Most active flows are accelerated; passive flows and flows for unsupported clients or servers are not. For example:

```
Active flows: 1
Passive flows: 6
Flows from unsupported clients: 2
Flows to unsupported servers: 2
Total flows: 32
Files currently tracked: 0
Accelerated writes: 0
Total writes: 0
Accelerated reads: 0
Total reads: 2
```



- The **Exchange** status shows the current number of active flows, and several totals since the device was last reset: the Packet Data Units (PDUs) reduced (cc) and decompressed (dc), the number of read, write, and “other” operations (“r/w/o”), and the number of reads and writes accelerated (“other” operations cannot be accelerated). For example:

```
Flows: Active:    5
PDUs : cc/dc:    367/456
      r/w/o:    723/543/342
Accel: Total:    612 (392 reads, 220 writes)
```

- The **HTTP** status shows the current cache usage for pre-fetched objects (items), cached static objects (datablocks), cookies, HTTP servers (hosts), and URLs (host-paths). For example:

```
***** Database Usage *****
      Total  Used
Items:      4096  0
Data Blocks: 8192  0
Cookies:    384   0
Hosts:      512   0
Host Paths: 16384  0
```

- **cluster < configuration | status >** . Lists the other WX devices in the same cluster (if any) or the last heartbeat sent and received by each device in the cluster. Clusters of devices can be defined for AFP if the outbound and return traffic does not always traverse the same two WX devices (asymmetric routing support).
- **packet-flow < configuration | status >** . Lists the global configuration settings for FCS, AFP, and FEC, or the configuration status for each remote endpoint.

## **show access-log**

The Access Control log files for the running configuration can be displayed in the CLI. To view the access control log file:

```
show access-log
```

```
CONSOLE: admin Session idle timeout 2005-12-16 10:02:51
CONSOLE: show uptime Login 2005-12-16 10:02:51
CONSOLE: admin Login 2005-12-16 14:56:39
CONSOLE: admin Session idle timeout 2005-12-16 15:26:41
CONSOLE: admin Login 2005-12-16 15:29:54
HTTPS: 172.23.8.148 admin Login 2005-12-16 15:36:23 HTTP/1.1 POST / 0
HTTPS: 10.84.26.92 admin Login 2005-12-16 15:49:11 HTTP/1.1 POST / 0
```

## **show all**

To view the entire current configuration:

```
show all
```

**show application**

To view the current definitions for all applications:

```
show -run application
```

To view the current definitions for one application:

```
show -run application <name>

CIFS: (Precedence = 6; Type = cifs)
rule 1:
  Source Port: 139,445
rule 2:
  Destination Port: 139,445
```

**show arp**

To view a list of static and dynamic ARP entries:

```
show -run arp
```

IP-address	Ethernet-addr	type	Interface
10.87.74.1	00:07:e9:1a:7f:ed	dynamic	REMOTE
10.87.74.12	00:0f:5a:00:00:00	dynamic	LOCAL
10.87.74.15	00:0e:0c:3b:92:60	dynamic	LOCAL

**show backup-sr**

To view the current backup configuration:

```
show -run backup-sr

Backup server mode: off
Configured primaries: none
```

**show clock**

To view the current clock settings:

```
show -run clock

Time: FRI DEC 16 15:59:25 2005
Timezone: 5 - (GMT -08:00) Pacific Time (US and Canada), Tijuana
Daylight saving: off
```

**show connection**

To view the connection status to the other devices in the community:

```
show connection
```

IP-Address	Role	Reduction	Assembly	RTT(ms)	Description
10.87.70.11	M	Disallowed	Failed		Disallowed by policy/No request received
10.87.71.10	M	Disallowed	Connected		Disallowed by policy
10.87.72.10	H	Connected	Connected	288	
10.87.73.11	M	Disallowed	Failed		Disallowed by policy/No request received

Role: H-Hub, S-Spoke, M-Mesh, B-Backup, A-Active Backup

**show console**

To view the current baud rate for the DB9 console port on the back of the WX device (the default is 9600):

```
show -run console
```

```
Baud rate: 9600
```

**show contact**

To view the contact information for your device administrator:

```
show contact
```

```
Contact: rclemens@company.com
```

**show dns**

To view the current DNS settings:

```
show -run dns
```

```
Default Domain Name:
```

```
DNS Server Addresses: None
```

DNS servers are used to resolve IP addresses on the Traffic report. When an IP address in the local domain is resolved by one of the DNS servers, the domain name is prepended to the host name shown on the Traffic report.

**show filter**

To view the applications and addresses whose traffic is not reduced, the protocols enabled for reduction, and whether fragmented packets are reduced:

```
show -run filter
```

```
Applications:
```

```
Groupwise
```

```
HTTPS
```

```
SNTP
```

```
SSH
```

```
Traceroute
```

```
Application mode: exclude
```

```
IP address pairs:
```

```
No address pairs in filter
```

```
Address pair mode: off
```

```
IP-protocol filter settings:
```

```
ip-protocol    status
```

```
1             enable
```

```
TCP           enable
```

```
UDP           enable
```

```
Reduce IP-fragments: on
```

**show flow-details**

For a specific traffic flow, you can view the number of bytes and packets sent and received, and whether reduction and acceleration were applied to the flow.

1. If necessary, run the Traffic report to identify the traffic flow you want to view (refer to “Traffic Statistics” on page 256).
2. To view the details of a specific traffic flow:

```
show flow-details src-ip <IP address> src-port <number> dst-ip <IP address>
dst-port <number> [proto <string>]
```

Where:

- **src-ip <IP address>** . Source IP address of the traffic flow.
- **src-port <number>** . Source port number of the traffic flow.
- **dst-ip <IP address>** . Destination IP address of the traffic flow.
- **dst-port <number>** . Destination port number of the traffic flow.
- **proto <string>** . Indicates the protocol is “tcp”, “udp”, or a protocol number (0 to 134). The default is TCP.

In the following example, the protocol defaults to TCP (protocol 6):

```
show flow-details src-ip 10.10.52.146 src-port 445 dst-ip 10.10.58.15 dst-port
1836
```

Retrieved flow details with the following parameters:

```
src-ip = 10.10.52.146, src-port = 445, dst-ip = 10.10.58.15, dst-port = 1836,
proto = 6
```

Flow Details:::

```
Bytes Sent: 29016709
Packets Sent: 19785
Bytes Received: 430250
Packets Received: 9756
Application Name: CIFS
Application Type: CIFS
Fast Connection configuration: off
Active Flow Pipeline configuration: on
Application Acceleration Configurations:
    Global CIFS acceleration mode: on
Traffic Type: Reduced, defined application
Fast Connection (FC) on this flow: Not applied because it is not enabled for FC
Active Flow Pipelining (AFP) on this flow: Not applied
Application acceleration (AAP) on this flow: Not applied
```

**show import-route-table**

To view information about the last time routes were imported from a router (the routes must be exported from the router and imported from a FTP or TFTP server).

```
show -run import-route-table
```

```
Date/Time of last import: 12/1/05
Number of routes      : 3000
Router IP address     : 10.20.30.1
```

**show interface**

To view the interface configuration and VLAN settings:

```
show -run interface
```

```
Settings for local interface
Link state: up
Speed/duplex: auto
Negotiated speed/duplex: 100-full
Hardware address: 00:0f:5a:00:00:00
Media type: copper
```

```
Settings for remote interface
Link state: up
Speed/duplex: auto
Negotiated speed/duplex: 100-full
Hardware address: 00:0f:5a:00:00:01
Media type: copper
```

```
Periodic passive interface test on:
Propagate failure - local to remote: off
Propagate failure - remote to local: off
Down time - local to remote: 15 sec
Down time - remote to local: 15 sec
802.1q VLAN Mode: off
Native VLAN ID: 1
VLAN ID: 1
Preserve VLAN ID on output packets: off
```

To include packet statistics on each interface:

```
show -run interface -verbose
```

```
Settings for local interface
Link state: up
Speed/duplex: auto
Negotiated speed/duplex: 100-full
Hardware address: 00:0f:5a:00:00:00
Media type: copper
  Octets Received      : 0
  Octets sent          : 0
  Packets Received     : 25932765
  Packets Sent         : 31764214
  Unicast Packets Received : 25932765
  Unicast Packets Sent   : 31764214
  Non-unicast Packets Received : 0
```

```

Non-unicast Packets Sent      : 0
Input Discards                : 0
Input Unknown Protocols      : 0
Input Errors                  : 0
Output Errors                 : 2

```

## ***show ip***

To view the current IP address, subnet mask, and gateway:

```

show -run ip

IP address: 10.87.74.12
Subnet mask: 255.255.255.0
Default gateway: 10.87.74.1

```

## ***show ipsec***

To view the current IPsec settings:

```
show -run ipsec [sa [<ip-address>]]
```

Where:

- **sa [ <ip-address> ]**. Displays the inbound and outbound security associations (SAs) for each endpoint or just the specified endpoint. Each SA specifies the algorithms and generated keys used to protect traffic in one direction. The SA information includes:
  - **SA Index**. Number that identifies each SA, also called the Security Parameter Index (SPI). To establish a secure connection, the outbound SA index on the sender must match an inbound SA index on the receiver.
  - **State**. Indicates whether an SA is “mature” (active) or “dying” (the key lifetime has expired). A new SA is negotiated when the key lifetime reaches 80% of the time limit or 50% of the data limit. After the first key expires, each endpoint has four SAs: two active (inbound and outbound) and two that are “dying.”
  - **Sequence #**. Indicates the sequence number of the last packet received. A packet is dropped if its sequence number is a duplicate or is not within 32 of the last received sequence number. Used for anti-replay protection.

## ***show license***

To view the current license and device serial number:

```

show -run license

Serial number: 0015000000
License key: Temporary license
License expiration: 12 days: 23 hours: 17 minutes
Licensed throughput: Unlimited
Licensed features: Packet Flow Acceleration, IPsec

```

**show location**

To view the specified physical location of the device:

```
show -run location
```

```
Location: Central data center
```

**show log**

The system log for the running configuration can be displayed in the CLI. To view the system log file:

```
show log
```

```
SR-10.87.74.12# show log
```

```
2005-12-16 00:00:46 01A88C90 I03 Finished daily rollup
```

```
2005-12-16 00:04:05 01A81550 I03 Saved monitoring stats to flash
```

```
.  
.
.
```

**show mon-apps**

To view the list of applications being monitored:

```
show -run mon-apps
```

**show multi-node**

To view whether multi-node is enabled:

```
show multi-node
```

```
Multi node: off
```

```
Master IP: 0.0.0.0
```

**show multi-node-status**

To view multi-node status and performance statistics on the master node:

```
show multi-node-status
```

```
Multi-node is not configured.
```

**show multi-path**

To view the current multi-path settings:

```
show -run multi-path
```

```
Multi-path mode: off
```

```
Secondary IP Address: 0.0.0.0
```

```
# multi-path endpoints: 0
```

```
# multi-path templates: 0
```

To view the last 32 events when traffic was switched between primary and secondary paths:

```
show multi-path events [<address>]
```

## **show ospf**

To view the current OSPF settings:

```
show -run ospf [all | neighbor [detail]]
```

The **all** option shows all configuration and neighbor information. The **neighbor detail** option shows details of the neighboring OSPF-enabled routers, such as the designated router (DR) and backup designated router (BDR). For example:

```
===== OSPF Neighbors =====
ID          Pri State Dead Time Address      Interface
13.13.13.1  1  2-Way 00:00:37 10.200.1.1   fei
14.14.14.2  1  Full  00:00:39 10.200.1.3   fei
15.15.15.2  1  2-Way 00:00:39 10.200.1.16  fei
11.11.11.2  1  2-Way 00:00:40 10.200.1.2   fei
16.16.16.2  1  Full  00:00:39 10.200.1.25  fei
===== OSPF Neighbors' Details =====
Neighbor 13.13.13.1, interface address 10.200.1.1

In the area 0 via interface fei
Neighbor Priority is 1, State is 2-Way, 2 state changes
DR is 10.200.1.25
BDR is 10.200.1.3
Options is DC N/P (0x15)
Dead timer due in 37 seconds
Authentication: none
```

## **show packet-capture**

To view the packet capture status and maximum trace size:

```
show packet-capture

Packet capture status: READY
Max trace size currently possible is 51998720 bytes
```

## **show packet-interception**

To view the current packet interception settings:

```
show -run packet-interception

Packet interception: Disabled
```



**show prime-time**

To view the current prime-time settings:

```
show -run prime-time
```

Prime-time enabled: off

Prime-time days: Sun, Mon, Tue, Wed, Thu, Fri, Sat

Prime-time hours: 0-24

**show profile-mode**

To view the current Profile Mode settings:

```
show -run profile-mode
```

Profile Mode: off

Remote SR devices: None

**show qos excl-filter**

To view the current LAN/WAN subnet pairs excluded from outbound QoS:

```
show -run qos excl-filter
```

Bandwidth Management filter : on

Bandwidth filter subnets:

Inbound Subnet	Inbound Mask	Outbound Subnet	Outbound Mask
*		10.87.74.0	255.255.255.0

**show qos inbound**

To view the current inbound QoS settings:

```
show -run qos inbound
```

Inbound Bandwidth management policy: on

Inbound Aggregate WAN Speed (Kbps): 10000

	Inbound	Bandwidth
Class	Max %	Queue Length
Default	100	Standard
Intranet	100	Standard
Reduced	100	Standard
Tcp	100	Standard

Intranet Class Subnets: Not Defined

**show qos outbound**

To view the current outbound QoS settings:

```
show -run qos outbound
```

Outbound QOS oversubscribed mode: on

Aggregate WAN speed: 10000 kbps

Outbound QOS mode: bw-weighted-fair-queueing

```

Outbound IP Precedence/DSCP mode: off
Restore original TOS/DSCP bits after assembly: on
WAN framing overhead: 14 bytes
Congestion control mode: off
Congestion control endpoint policy: all

```

```

# tunnels: 4, # templates: 0, # classes: 1
.
.
.

```

## **show radius**

To view the current RADIUS settings:

```
show -run radius
```

```

There are no Radius server groups defined.
There are no Radius servers defined.

```

## **show reduction**

To view the current reduction settings:

```
show -run reduction [all | network-sequence-mirroring | pre-sync status]
```

The “all” option includes reductions statistics since the last time the device was reset. For example:

```

===== Reduction Statistics =====
Packets: Total=79837075 - Accept=43450309
Overflow=0 FilterPassthru=75 Default Assembler=0 No Assembler=45723

Reject Protocol=57
Accept Protocol=58040
SR Traffic=1107
Local=36902
Mid Watermark packets=1351
Mid Watermark reached=5
Hi Watermark reached=1

```

The following table describes the reduction settings and statistics:

Keyword	Description
Total	Number of unreduced packets into the device.
Accept	Number of packets into the reduction engine.
Overflow	Packets not reduced because the reduction queue is full (the device is too busy or the WAN link is too slow).
FilterPassthru	Packets not reduced due to application or address filter settings.
Default Assembler	Packets reduced and sent to the default assembler.
No Assembler	Packets not reduced because of no remote WX device.
<b>The following statistics are shown only if they are non-zero.</b>	
Reject Protocol	Packets for IP protocols that are not reduced.

Keyword	Description
Accept Protocol	Packets for additional IP protocols that are enabled for reduction (does not include TCP and UDP packets, which are reduced by default).
Exclude Address	Packets not reduced due to source/destination filter settings.
TTL Expired	Packets not reduced because the Time to Live value was zero.
Accept Fragmented	Fragmented packets reduced (fragment reduction is enabled by default).
Reject Fragmented	Fragmented packets not reduced (fragment reduction is disabled).
Malformed	Malformed packets not reduced.
SR Traffic	Management packets sent to other WX devices (not reduced).
Local	Packets destined for the local subnet (not reduced).
Mid Watermark packets	Packets that received less reduction processing because the reduction queue exceeded the optimum level (the device is busy or the WAN link is slow).
Mid Watermark reached	Number of times that the reduction queue exceeded the optimum level.
Hi Watermark reached	Number of times the reduction queue became full. Packets received while the queue is full are counted as overflow (not reduced).

### ***show reduction-subnet***

To view the current reduction subnets and subnet settings:

```
show -run reduction-subnet
```

```
Mode: include
```

```
Wan-reduction-subnet Mode: off
```

```

Destination      Netmask      Cost  Enabled  Interface
192.168.0.0      255.255.255.0  1    no      Local

```

The Enabled column indicates whether the subnet is advertised.

### ***show reg-detail***

On a registration server, enter the following command to view the details for all registered devices, a specific device, or just the reducers or assemblers:

```
show -run reg-detail [<IP address> | -assemblers | -reducers]
```

```
Number of registered nodes: 4
```

```
Number of reducers: 4
```

```
Number of assemblers: 4
```

```
Node list:
```

```

IP-Address  Type  Duty Proto SW-Ver  Errors Last-Register-Time  Name
192.168.52.22 SA/SR  0    4    0 JAN 07 13:07:30 2005 52/22-SR20
192.168.53.22 SA/SR  0    7    0 JAN 07 09:56:43 2005 53/22-SR80
192.168.54.22 SA/SR  0    6    0 JAN 07 14:41:21 2005 54/22-SR15
192.168.55.22 SA/SR  R    0    7    0 JAN 07 09:51:42 2005 55/22-SR100

```

```
Key for 'Duty': H=Hub R=RegServer S=SecondaryRegServer
```

```
Key for 'Type': SA=Sequence Assembler SR=Sequence Reducer
```

The **Proto** and **SW-Ver** columns identify the registration protocol for each device (internal use only). The **Errors** indicate the number of times that the server failed to propagate registration updates to a device.

To reset all the error counts to zero:

```
config reg-server clear-error-count
```



**NOTE:** Each device obtains all the latest registration information, including any missed updates, when it checks in with the registration server (every eight hours).

### ***show reg-server***

To view the current registration server settings (the communities are shown only if the device is a registration server):

```
show -run reg-server
```

```
Registration server: 192.168.55.22
Secondary registration server: not set
This system is currently the registration server
Connection timeout (seconds): 2
Connection retry count: 1
Self registration frequency: 24-hours
```

```
2 Communities
Community "default-192.168.55.22" has 0 entries:
Community "Main" has 4 entries:
192.168.52.22 192.168.53.22 192.168.54.22 192.168.55.22
```

### ***show reg-summary***

On a registration server, enter the following command to list the registered WX devices:

```
show -run reg-summary
```

```
Number of registered nodes: 3
Node list:
10.87.52.22
10.87.53.22
10.87.54.22
```

### ***show remote-routes***

Remote routes are the reduction subnets advertised by other WX devices in the community. To view the current remote routes and settings:

```
show -run remote-routes
```

```
Validation status: off
Validation frequency: 3600 seconds
```

```
Destination Netmask Cost Assembler-IP On Type Last-Validation
10.87.52.0 255.255.255.0 1 10.87.52.22 Y dynamic SAT DEC 17
11:04:44 2005
```

```

10.87.53.0 255.255.255.0 1 10.87.53.22 Y dynamic SAT DEC 17
11:04:45 2005
10.87.54.0 255.255.255.0 1 10.87.54.22 Y dynamic SAT DEC 17
11:04:45 2005
10.87.56.0 255.255.255.0 1 10.87.56.22 Y dynamic SAT DEC 17
11:04:45 2005

```

## ***show rip***

To view the current RIP configuration:

```

show -run rip

RIP receive mode: off
RIP send mode (for packet interception): off
RIP version: 2
RIP ageout interval: 300 seconds
RIP auth Type: none
RIP password: Not set
RIP is currently not running.

```

## ***show route***

To view the current routes and route settings (all routes are shown by default):

```

show -run route [protocol <ospf | rip | static>] [subnet <subnet/mask>]

Precedence: dynamic
Router load balancing policy for equal-cost routes: off
ToS marking for router-based load balancing: off
ToS marking policy for router-based load balancing: per-destination
ICMP Redirect Ageout Interval: 10

```

Destination	Netmask	Gateway	Type
0.0.0.0	0.0.0.0	10.87.74.1	static
10.87.74.0	255.255.255.0	10.87.74.12	dynamic
127.0.0.1	0.0.0.0	127.0.0.1	dynamic

## ***show route-poll***

The WX device can obtain dynamic routes by periodically polling a Cisco router. To view the current route poll settings:

```

config -run route-poll

Remote-host: Not set
Remote Port: 514
Secondary Remote-host: Not set
Secondary Remote Port: 514
Local-user: Not set
Remote-user: Not set
Mode: none
Frequency poll: 5
Allow BGP routes: off

```

**show security**

To view the current security settings:

```
show -run security

Web status: on
SSH status: on
Front-panel status: on
No allow list defined.
No deny list defined.
```

**show snmp**

To view the current SNMP settings:

```
show -run snmp

SNMP status: on
Read community string: Set but not displayed for security reasons.
Write community string: Set but not displayed for security reasons.
Trap status: off
Trap community string: Set but not displayed for security reasons.
Trap destination: No entries
Authentication-failure trap status: off
```

**show sntp**

WX devices support the Simple Network Time Protocol (SNTP). An SNTP server provides a common time base for devices within your network. To view the current SNTP settings:

```
show -run sntp

SNTP status: off
SNTP server: Not set
SNTP secondary server: Not set
Update interval: 1440
```

**show stack-group**

The WX 100 can act as a server to distribute the processing load to a “stack” of up to six client WX devices. To view the stack configuration on a WX 100:

```
show -run stack-group

Client mode: off

STACK CONFIGURATION:

Port  Status      Model  Disk Status  # Tunnels
      Model                OUT  IN
Master SR-100                0    0
  1 Active   WXC-500 OK      3    3
  2 Active   WXC-500 OK      3    3
host-session: clients-only
sequence-mirror-server: on
```

**show syslog**

WX devices can send Syslog messages to one or more Syslog servers. A Syslog server allows you to centrally log and analyze configuration events and system error messages such as interface status, security alerts, and environmental conditions.

To view the current Syslog settings:

```
show -run syslog

Syslog status: off
Severity: CE
Destination: Not set
```

**show system**

To view general system information, such as the device name, location, and model number:

```
show -run system

System name: SR-10.87.72.10
Location:
Contact:

Software version: 5.2.0.14
Model No.: SR-100 - v2.0
Link-switch: HW version 2, SW version 22

System started at: FRI DEC 02 16:41:07 2005
System up for: 2 days 21 hours 4 minutes

Info for flash file system /ata1
File system size: 510648320 bytes
Free space: 442998784 bytes
File system block size: 8192

Total physical memory: 4160221184 (0xF7F7F000) bytes
```

**show system-name**

To view the system name:

```
show -run system-name

System name: SR-10.87.72.10
```

**show top-talker**

Traffic data is collected continuously for the most active traffic flows, including the application name and protocol, the 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 Web console. The Traffic utility maintains the 65,000 most active flows. You can view the top 50 flows in the Web console, but the complete list can be exported to a file in CSV format.

To view the current Top Talker settings:

```
show -run top-talker
```

**Top-Talker Parameters**

Ip top-talk: on

Top Talker collection period: continuous

Export to Cisco NetFlow collector: disabled

Cisco NetFlow IP address: 0.0.0.0, UDP port: 0

**show uptime**

To view the length of time the device has been running:

```
show -run uptime
```

System started at: FRI DEC 02 16:41:07 2005

System up for: 2 days 21 hours 17 minutes

**show version**

To view the device's model number and hardware and software versions:

```
show version
```

Software version: 5.2.0.14

Model No.: SR-100 - v2.0

Link-switch: HW version 2, SW version 22

**show wan-performance-mon**

You can enable WAN performance monitoring to measure the latency and loss between the current device and one or more remote WX devices. To view the current settings:

```
show -run wan-performance-monitor
```

Wan-performance-monitoring mode: off

	Minutes	Minutes	Minutes	Minutes	Probes	
Latency	To Bad	To Bad	To Good	To Good	Probes	Per
Endpoint	Threshold	Latency	Loss	Latency	Loss	Lost Minute
Global Values	5000	4	4	4	2	12



## Appendix A

# WX Device Specifications

This appendix lists the technical specifications for the following WX devices:

- “WX 15 Specifications” in the next section
- “WX 20 Specifications” on page 387
- “WX 50 and WX 60 Specifications” on page 388
- “WX 100 Specifications” on page 390
- “WXC 250 Specifications” on page 391
- “WXC 500 Specifications” on page 393
- “DB9 Console Port Pin-Outs” on page 395

### WX 15 Specifications

**Table 11: WX 15 Specifications**

Product Features	WX 15
Protocols supported	Any IP-based traffic (TCP, UDP, GRE, etc.)
Operator-defined passthrough filter	Yes, by application or address. Passes native (non-reduced) data at wire speed.
Installation	In-Line between aggregation switch and edge router or off WAN router
<b>System</b>	
Network interfaces	Two auto-sensing 10/100 BaseT Ethernet interfaces
On-board Flash storage	Yes, no spinning media
<b>Performance</b>	
Total reduction throughput speed	Supports WAN speeds up to 1 Mbps
Maximum connections	Up to 2 WX/WXC connections, 2 IPSec tunnels, 2 virtual endpoints, and 5 WAN monitoring endpoints
Application definitions	Up to 100 application definitions
Routes	Up to 1000 routes
<b>Quality of Service (QoS)</b>	
Honor, preserve and/or set ToS/DSCP	Yes, retain settings or prioritize using ToS/DiffServ values by application
Bandwidth allocation	Yes, create traffic classes for bandwidth allocation with time of day option

Product Features		WX 15
Traffic Acceleration		
Packet Flow Acceleration (PFA)	Active Flow Pipelining, Fast Connection Setup, and Forward Error Correction	
Application Flow Acceleration	Microsoft CIFS, Microsoft Exchange, and HTTP	
Management		
SNMP, Syslog	Yes, SNMPv2c, MIB II, WX Enterprise MIB and syslog	
Secure remote access	SSH and HTTPS (SSL)	
Monitoring		
Reduction statistics	Per device, per application, and per destination; both real-time and historical	
QoS, bandwidth management	Per destination, per traffic class, real-time and historical	
Acceleration	TCP session time and throughput; both real-time and historical	
Data export	Yes, CSV format and NetFlow version 5 records	
Application reporting	Detail by IP addresses, and/or port numbers, and/or IP protocol, and/or DSCP/ToS value, with greater detail by URL element or application type	
Network upgradeable	Yes, via FTP, HTTP and TFTP	
Fault tolerant non-stop operation	Yes, 10/100 BaseT auto switch-to-wire on any power, hardware, or software failure condition	
High availability	Yes, a backup device can support multiple primary devices	
Power	AC power 100-240v, 50-60Hz, 50 Watts Max or 170 BTU/hr	
Dimensions and Weight		
Height	44 mm (1.75 in.): 1 rack unit	
Width	390 mm (15.3 in.)	
Depth	230 mm (9.1 in.)	
Weight	1.8 kg (4 lbs.)	
Operating Environment		
Temperature	5° C to 40° C	
Relative Humidity	10 % to 85 %, non-condensing at 35° C	
Maximum Altitude	10,000 feet (12,192 meters)	
Non-Operating Environment		
Temperature	-40° C to 70° C	
Relative Humidity	5 % to 95 %, non-condensing at 35° C	
Maximum Altitude	40,000 feet (3048 meters)	
Regulations		
Emissions	FCC Class A, EN 55022 Class A, EN 55024 Immunity	
Safety	CSA C22.2 No. 950 M95, UL 1950 3 Edition, EN 60950	
Acoustic Noise	Maximum noise level is less than 70dB.  Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 70 dB(A) oder weniger gemäss EN ISO 7779.	

## WX 20 Specifications

**Table 12: WX 20 Specifications**

Product Features	WX 20
Protocols supported	Any IP-based traffic (TCP, UDP, GRE, etc.)
Operator-defined passthrough filter	Yes, by application or address. Passes native (non-reduced) data at wire speed.
Installation	In-Line between aggregation switch and edge router or off WAN router
<b>System</b>	
Network interfaces	Two auto-sensing 10/100 BaseT Ethernet interfaces
On-board Flash storage	Yes, no spinning media
<b>Performance</b>	
Total reduction throughput speed	Supports WAN speeds up to 2 Mbps
Maximum connections	Up to 15 WX/WXC connections, 15 IPSec tunnels, 5 virtual endpoints, and 20 WAN monitoring endpoints
Application definitions	Up to 256 application definitions
Routes	Up to 8K routes
<b>Quality of Service (QoS)</b>	
Honor, preserve and/or set ToS/DSCP	Yes, retain settings or prioritize using ToS/DiffServ values by application
Bandwidth allocation	Yes, create traffic classes for bandwidth allocation with time of day option
<b>Traffic Acceleration</b>	
Packet Flow Acceleration (PFA)	Active Flow Pipelining, Fast Connection Setup, and Forward Error Correction
Application Flow Acceleration	Microsoft CIFS, Microsoft Exchange, and HTTP
<b>Management</b>	
SNMP, Syslog	Yes, SNMPv2c, MIB II, WX Enterprise MIB and syslog
Secure remote access	SSH and HTTPS (SSL)
<b>Monitoring</b>	
Reduction statistics	Per device, per application, and per destination; both real-time and historical
QoS, bandwidth management	Per destination, per traffic class, real-time and historical
Acceleration	TCP session time and throughput; both real-time and historical
Data export	Yes, CSV format and NetFlow version 5 records
Application reporting	Detail by IP addresses, and/or port numbers, and/or IP protocol, and/or DSCP/ToS value, with greater detail by URL element or application type
Network upgradeable	Yes, via FTP, HTTP and TFTP
Fault tolerant non-stop operation	Yes, 10/100 BaseT auto switch-to-wire on any power, hardware, or software failure condition
High availability	Yes, a backup device can support multiple primary devices
Power	AC power 110-230v, 47-63Hz, 150 Watts Max or 510 BTU/hr
<b>Dimensions and Weight</b>	
Height	44 mm (1.8 in.): 1 rack unit
Width	435 mm (17.1 in.)
Depth	363 mm (14.3 in.)

Product Features	WX 20
Weight	8.6 kg (19 lbs.)
<b>Operating Environment</b>	
Temperature	5° C to 40° C
Relative Humidity	10 % to 85 %, non-condensing at 35° C
Maximum Altitude	10,000 feet (12,192 meters)
<b>Non-Operating Environment</b>	
Temperature	-40° C to 70° C
Relative Humidity	5 % to 95 %, non-condensing at 35° C
Maximum Altitude	40,000 feet (3048 meters)
<b>Regulations</b>	
Emissions	FCC Class A, EN 55022 Class A, EN 55024 Immunity
Safety	CSA C22.2 No. 950 M95, UL 1950 3 Edition, EN 60950
Acoustic Noise	Maximum noise level is less than 70dB. Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 70 dB(A) oder weniger gemäss EN ISO 7779.

## WX 50 and WX 60 Specifications

**Table 13: WX 50 and WX 60 Specifications**

Product Features	WX 50 and WX 60
Protocols supported	Any IP-based traffic (TCP, UDP, GRE, etc.)
Operator-defined passthrough filter	Yes, by application or address. Passes native (non-reduced) data at wire speed.
Installation	In-Line between aggregation switch and edge router or off WAN router
<b>System</b>	
Network interfaces	WX 50: Two auto-sensing 10/100 BaseT Ethernet interfaces WX 60: Two auto-sensing 10/100/1000 BaseT Ethernet ports
On-board Flash storage	Yes, no spinning media
<b>Performance</b>	
Total reduction throughput speed	Supports WAN speeds up to 20 Mbps
Maximum connections	WX 50: Up to 120 WX/WXC connections, 100 IPSec tunnels, 120 virtual endpoints, and 240 WAN monitoring endpoints WX 60: Up to 150 WX/WXC connections, 100 IPSec tunnels, 120 virtual endpoints, and 240 WAN monitoring endpoints
Application definitions	Up to 256 application definitions
Routes	Up to 8K routes
<b>Quality of Service (QoS)</b>	
Honor, preserve and/or set ToS/DSCP	Yes, retain settings or prioritize using ToS/DiffServ values by application
Bandwidth allocation	Yes, create traffic classes for bandwidth allocation with time of day option

Product Features	WX 50 and WX 60
<b>Traffic Acceleration</b>	
Packet Flow Acceleration (PFA)	Active Flow Pipelining, Fast Connection Setup, and Forward Error Correction
Application Flow Acceleration	Microsoft CIFS, Microsoft Exchange, and HTTP
<b>Management</b>	
SNMP, Syslog	Yes, SNMPv2c, MIB II, WX Enterprise MIB and syslog
Secure remote access	SSH and HTTPS (SSL)
<b>Monitoring</b>	
Reduction statistics	Per device, per application, and per destination; both real-time and historical
QoS, bandwidth management	Per destination, per traffic class, real-time and historical
Acceleration	TCP session time and throughput; both real-time and historical
Data export	Yes, CSV format and NetFlow version 5 records
Application reporting	Detail by IP addresses, and/or port numbers, and/or IP protocol, and/or DSCP/ToS value, with greater detail by URL element or application type
Network upgradeable	Yes, via FTP, HTTP and TFTP
Fault tolerant non-stop operation	Yes, 10/100/1000 BaseT auto switch-to-wire on any power, hardware, or software failure condition
High availability	Yes, a backup device can support multiple primary devices
Power	AC power 100-240v, 50-60Hz, 150 Watts Max or 510 BTU/hr
<b>Dimensions and Weight</b>	
Height	88 mm (3.44 in.): 2 rack units
Width	435 mm (17.1 in.)
Depth	425 mm (16.7 in.)
Weight	9.2 kg (20.2 lbs.)
<b>Operating Environment</b>	
Temperature	5° C to 40° C
Relative Humidity	10 % to 85 %, non-condensing at 35° C
Maximum Altitude	10,000 feet (12,192 meters)
<b>Non-Operating Environment</b>	
Temperature	-40° C to 70° C
Relative Humidity	5 % to 95 %, non-condensing at 35° C
Maximum Altitude	40,000 feet (3048 meters)
<b>Regulations</b>	
Emissions	FCC Class A, EN 55022 Class A, EN 55024 Immunity
Safety	CSA C22.2 No. 950 M95, UL 1950 3 Edition, EN 60950
Acoustic Noise	Maximum noise level is less than 70dB. Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 70 dB(A) oder weniger gemäss EN ISO 7779.

## WX 100 Specifications

**Table 14: WX 100 Specifications**

Product Features	WX 100
Protocols supported	Any IP-based traffic (TCP, UDP, GRE, etc.)
Operator-defined passthrough filter	Yes, by application or address. Passes native (non-reduced) data at wire speed.
Installation	In-Line between aggregation switch and edge router or off WAN router
<b>System</b>	
Network interfaces	Two auto-sensing 10/100/1000 BaseT Ethernet interfaces or Two 1000 Base-SX multimode fiber Ethernet interfaces
Client interfaces	Up to six WX 60s or WXC 500s can be connected to the WX 100
On-board Flash storage	Yes, no spinning media
<b>Performance</b>	
Total reduction throughput speed	Supports WAN speeds up to 155 Mbps
Maximum connections	Up to 320 WX/WXC connections, 100 IPSec tunnels, 320 virtual endpoints, and 440 WAN monitoring endpoints
Application definitions	Up to 256 application definitions
Routes	Up to 8K routes
<b>Quality of Service (QoS)</b>	
Honor, preserve and/or set ToS/DSCP	Yes, retain settings or prioritize using ToS/DiffServ values by application
Bandwidth allocation	Yes, create traffic classes for bandwidth allocation with time of day option
<b>Traffic Acceleration</b>	
Packet Flow Acceleration (PFA)	Active Flow Pipelining, Fast Connection Setup, and Forward Error Correction
Application Flow Acceleration	Microsoft CIFS, Microsoft Exchange, and HTTP
<b>Management</b>	
SNMP, Syslog	Yes, SNMPv2c, MIB II, WX Enterprise MIB and syslog
Secure remote access	SSH and HTTPS (SSL)
<b>Monitoring</b>	
Reduction statistics	Per device, per application, and per destination; both real-time and historical
QoS, bandwidth management	Per destination, per traffic class, real-time and historical
Acceleration	TCP session time and throughput; both real-time and historical
Data export	Yes, CSV format and NetFlow version 5 records
Application reporting	Detail by IP addresses, and/or port numbers, and/or IP protocol, and/or DSCP/ToS value, with greater detail by URL element or application type
Network upgradeable	Yes, via FTP, HTTP and TFTP
Fault tolerant non-stop operation	Yes, 10/100/1000 BaseT auto switch-to-wire on any power, hardware, or software failure condition
High availability	Yes, a backup device can support multiple primary devices
Power	Dual AC power 110-240v, 50-60Hz, 300 Watts Max or 1025 BTU/hr. The appliance is designed to work with IT power systems.

Product Features	WX 100
<b>Dimensions and Weight</b>	
Height	88 mm (3.44 in.): 2 rack units
Width	435 mm (17.1 in.)
Depth	524 mm (20.6in.)
Weight	13.6 kg (30 lbs.)
<b>Operating Environment</b>	
Temperature	5° C to 40° C
Relative Humidity	10 % to 85 %, non-condensing at 35° C
Maximum Altitude	10,000 feet (12,192 meters)
<b>Non-Operating Environment</b>	
Temperature	-40° C to 70° C
Relative Humidity	5 % to 95 %, non-condensing at 35° C
Maximum Altitude	40,000 feet (3048 meters)
<b>Regulations</b>	
Emissions	FCC Class A, EN 55022 Class A, EN 55024 Immunity
Safety	CSA C22.2 No. 60950, UL 60950, EN 60950
Acoustic Noise	Maximum noise level is less than 70dB. Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 70 dB(A) oder weniger gemäss EN ISO 7779.

## WXC 250 Specifications

**Table 15: WXC 250 Specifications**

Product Features	WXC 250
Protocols supported	Any IP-based traffic (TCP, UDP, GRE, etc.)
Operator-defined passthrough filter	Yes, by application or address. Passes native (non-reduced) data at wire speed.
Installation	In-Line between aggregation switch and edge router or off WAN router
<b>System</b>	
Network interfaces	Two auto-sensing 10/100 BaseT Ethernet ports with MDI/MDI-X switch
On-board Flash storage	Yes
Disk drives	One 40 GB hard disk drive
<b>Performance</b>	
Total reduction throughput speed	Supports WAN speeds up to 2 Mbps
Maximum connections	Up to 15 WX/WXC connections, 15 IPSec tunnels, 5 virtual endpoints, and 20 WAN monitoring endpoints
Application definitions	Up to 256 application definitions
Routes	Up to 8K routes

<b>Product Features</b>		<b>WXC 250</b>
<b>Quality of Service (QoS)</b>		
Honor, preserve and/or set ToS/DSCP		Yes, retain settings or prioritize using ToS/DiffServ values by application
Bandwidth allocation		Yes, create traffic classes for bandwidth allocation with time of day option
<b>Traffic Acceleration</b>		
Packet Flow Acceleration (PFA)		Active Flow Pipelining, Fast Connection Setup, and Forward Error Correction
Application Flow Acceleration		Microsoft CIFS, Microsoft Exchange, and HTTP
<b>Management</b>		
SNMP, Syslog		Yes, SNMPv2c, MIB II, WX Enterprise MIB and syslog
Secure remote access		SSH and HTTPS (SSL)
<b>Monitoring</b>		
Reduction statistics		Per device, per application, and per destination; both real-time and historical
QoS, bandwidth management		Per destination, per traffic class, real-time and historical
Acceleration		TCP session time and throughput; both real-time and historical
Data export		Yes, CSV format and NetFlow version 5 records
Application reporting		Detail by IP addresses, and/or port numbers, and/or IP protocol, and/or DSCP/ToS value, with greater detail by URL element or application type
Network upgradeable		Yes, via FTP, HTTP and TFTP
Fault tolerant non-stop operation		Yes, 10/100 BaseT auto switch-to-wire on any power, hardware, or software failure condition
High availability		Yes, a backup device can support multiple primary devices
Power		AC power 100-240v, 50-60Hz, 150 Watts Max or 510 BTU/hr
<b>Dimensions and Weight</b>		
Height		44 mm (1.8 in.): 1 rack unit
Width		435 mm (17.1 in.)
Depth		363 mm (14.3 in.)
Weight		9.5 kg (21 lbs.)
<b>Operating Environment</b>		
Temperature		5° C to 40° C
Relative Humidity		10 % to 85 %, non-condensing at 35° C
Maximum Altitude		10,000 feet (12,192 meters)
<b>Non-Operating Environment</b>		
Temperature		-40° C to 70° C
Relative Humidity		5 % to 95 %, non-condensing at 35° C
Maximum Altitude		40,000 feet (3048 meters)
<b>Regulations</b>		
Emissions		FCC Class A, EN 55022 Class A, EN 55024 Immunity
Safety		CSA C22.2 No. 950 M95, UL 1950 3 Edition, EN 60950
Acoustic Noise		Maximum noise level is less than 70dB. Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 70 dB(A) oder weniger gemäss EN ISO 7779.



## WXC 500 Specifications

**Table 16: WXC 500 Specifications**

Product Features	WXC 500
Protocols supported	Any IP-based traffic (TCP, UDP, GRE, etc.)
Operator-defined passthrough filter	Yes, by application or address. Passes native (non-reduced) data at wire speed.
Installation	In-Line between aggregation switch and edge router or off WAN router
<b>System</b>	
Network interfaces	Two auto-sensing 10/100/1000 BaseT Ethernet ports
On-board Flash storage	Yes
Disk drives	Two 250 GB hard disk drives
<b>Performance</b>	
Total reduction throughput speed	Supports WAN speeds up to 20 Mbps
Maximum connections	Up to 60 connections to other WXC's, 60 IPSec/WXC tunnels, 60 virtual endpoints, and 120 WAN monitoring endpoints
Application definitions	Up to 256 application definitions
Routes	Up to 8K routes
<b>Quality of Service (QoS)</b>	
Honor, preserve and/or set ToS/DSCP	Yes, retain settings or prioritize using ToS/DiffServ values by application
Bandwidth allocation	Yes, create traffic classes for bandwidth allocation with time of day option
<b>Traffic Acceleration</b>	
Packet Flow Acceleration (PFA)	Active Flow Pipelining, Fast Connection Setup, and Forward Error Correction
Application Flow Acceleration	Microsoft CIFS, Microsoft Exchange, and HTTP
<b>Management</b>	
SNMP, Syslog	Yes, SNMPv2c, MIB II, WX Enterprise MIB and syslog
Secure remote access	SSH and HTTPS (SSL)
<b>Monitoring</b>	
Reduction statistics	Per device, per application, and per destination; both real-time and historical
QoS, bandwidth management	Per destination, per traffic class, real-time and historical
Acceleration	TCP session time and throughput; both real-time and historical
Data export	Yes, CSV format and NetFlow version 5 records
Application reporting	Detail by IP addresses, and/or port numbers, and/or IP protocol, and/or DSCP/ToS value, with greater detail by URL element or application type
Network upgradeable	Yes, via FTP, HTTP and TFTP
Fault tolerant non-stop operation	Yes, 10/100/1000 BaseT auto switch-to-wire on any power, hardware, or software failure condition
High availability	Yes, a backup device can support multiple primary devices
Power	AC power 100-240v, 50-60Hz, 150 Watts Max or 510 BTU/hr
<b>Dimensions and Weight</b>	
Height	88 mm (3.44 in.): 2 rack units
Width	435 mm (17.1 in.)

<b>Product Features</b>		<b>WXC 500</b>
Depth		425 mm (16.7 in.)
Weight		11.3 kg (25 lbs.)
<b>Operating Environment</b>		
Temperature		5° C to 40° C
Relative Humidity		10 % to 85 %, non-condensing at 35° C
Maximum Altitude		10,000 feet (12,192 meters)
<b>Non-Operating Environment</b>		
Temperature		-40° C to 70° C
Relative Humidity		5 % to 95 %, non-condensing at 35° C
Maximum Altitude		40,000 feet (3048 meters)
<b>Regulations</b>		
Emissions		FCC Class A, EN 55022 Class A, EN 55024 Immunity
Safety		CSA C22.2 No. 950 M95, UL 1950 3 Edition, EN 60950
Acoustic Noise		Maximum noise level is less than 70dB. Maschinenlärminformations-Verordnung - 3. GPSGV, der höchste Schalldruckpegel beträgt 70 dB(A) oder weniger gemäss EN ISO 7779.

## DB9 Console Port Pin-Outs

The following tables list the pin-outs for a null-modem cable used to connect the DB9 console port to a DB9 or DB25 terminal port. Applies to all WX devices.

**Table 17: DB9 to DB9 Cable**

Console Port	DB9	DB9	Terminal Port
Receive Data	2	3	Transmit Data
Transmit Data	3	2	Receive Data
Data Terminal Ready	4	6 + 1	Data Set Ready + Carrier Detect
System Ground	5	5	System Ground
Data Set Ready + Carrier Detect	6 + 1	4	Data Terminal Ready
Request to Send	7	8	Clear to Send
Clear to Send	8	7	Request to Send

**Table 18: DB9 to DB25 Cable**

Console Port	DB9	DB25	Terminal Port
Receive Data	2	2	Transmit Data
Transmit Data	3	3	Receive Data
Data Terminal Ready	4	6 + 8	Data Set Ready + Carrier Detect
System Ground	5	7	System Ground
Data Set Ready + Carrier Detect	6 + 1	20	Data Terminal Ready
Request to Send	7	5	Clear to Send
Clear to Send	8	4	Request to Send



## Appendix B

# SNMP Traps and Syslog Messages

This appendix describes the SNMP Traps generated by WX devices, and describes the messages that are sent to a Syslog server (if configured).

## SNMP Traps

Table 19 lists the generic traps supported.

**Table 19: Generic SNMP Traps**

Trap	Description
Cold Start	The device was restarted.
LAN Link Up	Indicates the Local interface link has been established.
LAN Link Down	Indicates the Local interface link has failed. Verify that the link state change was not due to a network error.
WAN Link Up	Indicates the Remote interface link has been established.
WAN Link Down	Indicates the Remote interface link has failed. Verify that the link state change was not due to a network error.

Table 20 lists the enterprise specific traps supported.

**Table 20: Enterprise-specific SNMP Traps Table continued on next page**

Event	Description/Recommended Action
pnCommonEventInFailSafeMode	Indicates that the device was restarted in Safe Mode. Safe Mode operation keeps the device powered on, but all traffic is passed through without reduction.
pnCommonEventPowerSupplyFailure	One or more sources of power to the system has failed. A redundant power-supply has presumably taken over.
pnCommonEventPowerSupplyOk	One or more previously failed sources of power is now working normally. The transition to normal condition happened without the system having to be restarted.
pnCommonEventLicenseExpired	Software license has expired. Data reduction/assembly has been disabled. Please contact Juniper Networks to obtain a permanent license
pnCommonEventThruputLimitExceeded	Exceeded licensed throughput. Please contact Juniper Networks to obtain a new license with a higher speed.
pnCommonEventLicenseWillExpire	Software license will expire soon. If you are using an evaluation license, contact Juniper Networks to obtain a permanent license.

Event	Description/Recommended Action
pnCommonEventLoginFailure	Verify the user is authorized to administer the device. Any unauthorized access should be treated as a serious problem.
pnCommonEventFaultTolerantPassThrough	An anomalous health condition was detected. It would have subsequently triggered hardware pass through mode followed by a reboot.
pnCommonEventFanFailure	A cooling fan inside the device has failed. The 'pnCommonEventDescr' object has the name of the fan that failed.
pnCommonEventFanSpeedVariation	The speed of a cooling fan inside the device is either too low or too high. The 'pnCommonEventDescr' object has the name of the fan that has the problem.
pnCommonEventFanOk	A cooling fan inside the device that had previously failed or whose speed variation was high is now working properly. The 'pnCommonEventDescr' object has the name of the fan that has recovered.
pnCommonEventInterfaceSpeedMismatch	A mismatch is detected between the local and remote interface settings. This can happen due to a mismatch in the local and remote interface speed or mode.
pnCommonEventInterfaceSpeedOk	A mismatch previously detected between the local and remote interface settings is now resolved. The local and remote interface speed and mode are matched.
pnCommonEventInterfaceDuplexMismatch	A possible mismatch was detected between the duplex setting of either the local or remote interface and that of the device attached to that interface. The interface (local or remote) is identified by the 'pnCommonEventDescr' object. This notification is different from InterfaceSpeedMismatch event, which compares the local and remote interfaces on the same WX device.
pnCommonEventIpsecSecurityAssociationAdded	An IPsec security association has been negotiated and added to security association database.
pnCommonEventIpsecSecurityAssociationDeleted	An IPsec security association has been deleted from the security association database.
pnCommonEventIpsecSecurityAssociationExpired	An IPsec security association has expired.
pnSrEventRipAuthFailure	Indicates that a RIP packet received from a device could not be authenticated. Check the authentication information on the WX device and the sending device.
pnSrEventReducerBufferOverflow	The reduction input buffer is approaching full capacity.
pnSrEventReducerSessionClosed	The reducer session to the device described in pnCommonEventDescr was terminated.
pnSrEventAssemblerSessionClosed	The assembler session to the device described in pnCommonEventDescr was terminated.
pnSrEventReducerSessionOpened	The reducer session to the device described in pnCommonEventDescr was opened.
pnSrEventAssemblerSessionOpened	The assembler session to the device described in pnCommonEventDescr was opened.
pnSrEventPrimaryRegServerUnreachable	The primary registration server is currently unreachable.
pnSrEventSecondaryRegServerUnreachable	The secondary registration server is currently unreachable.
pnSrEventMultiNodeMasterUp	The system designated as the 'master' of a multi-node configuration came up. This notification is generated by the system that's designated as the 'master' of the multi-node. Note that the corresponding Down notification is generated by the designated 'master-backup' of the same multi-node.

Event	Description/Recommended Action
pnSrEventMultiNodeMasterDown	The system designated as the 'master' of a multi-node configuration is currently down. This notification is generated by the system that's designated as the 'master-backup' of the same multi-node. Note that the corresponding Up notification is generated by the designated 'master' of the same multi-node.
pnSrEventMultiNodeLastUp	The system designated as the 'last-node' of a multi-node came up. This notification is generated by the system that's designated as the 'last-node' of the multi-node. Note that the corresponding Down notification is generated by the designated 'master' of the same multi-node.
pnSrEventMultiNodeLastDown	The system designated as the 'last-node' of a multi-node is currently down. This notification is generated by the system that's designated as the 'master' of the same multi-node. Note that the corresponding Up notification is generated by the designated 'last-node' of the same multi-node.
pnSrEventMultiPathStatusChange	The primary or secondary path to another multipath-enabled system became inactive or failed. This may have caused traffic designated to flow over this path to be switched to or from this path.
pnSrEventPrimaryDownBackupEngaged	The system designated as the 'primary' is currently unreachable. This notification is generated by the system that's designated as the 'backup' device. The backup device is engaged for the primary device.
pnSrEventPrimaryDownBackupEngageFailed	The system designated as the 'primary' is currently unreachable. This notification is generated by the system that's designated as the 'backup' device. The backup device failed to engage for the primary device.
pnSrEventPrimaryUpBackupDisengaged	The system designated as the 'primary' is currently reachable. This notification is generated by the system that's designated as the 'backup' device. The backup device has disengaged.
pnSrEventWanPerfStatusChange	The status of the Path on which WAN Performance Monitoring is enabled has changed. The performance of the path has changed either from acceptable to unacceptable or vice versa.

## Syslog Messages

Table 21 lists the Syslog messages generated by WX devices.

**Table 21: Syslog Messages**

Message ID	101: PN_LIC_LICENSE_WILL_EXPIRE_SOON_ID
Message	License will expire on < date >
Message Type	Informational
Recommended Action	Once the license expires, please contact Juniper Networks to obtain a new license.
Message ID	102: PN_LIC_SPEED_THRESHOLD_EXCEEDED_ID
Message	Exceeded licensed throughput
Message Type	Error
Recommended Action	Contact Juniper Networks to obtain a new license with speed configured to a higher value
Message ID	103: PN_LIC_LICENSE_EXPIRED_ID
Message	License expired, Data reduction/assembly has been disabled
Message Type	Error
Recommended Action	Contact Juniper Networks to obtain a new license

Message ID	602: PN_ROUTING_RIP_AUTH_FAIL
Message	"RIP Authentication failed from <IPAddr>", where <IPAddr> is the address of the machine for which we could not authenticate the packet
Message Type	Informational
Recommended Action	Check the RIP authentication settings on the WX device and the <IPAddr> machine.
Message ID	902: PN_REDUCER_PASSTHRU_INFO_ID
Message	SR: Connection state set to pass through for ip = <ip address> .
Message Type	Informational.
Recommended Action	Heartbeats are missed for device <ip address> .
Message ID	903: PN_REDUCER_END_SESSION_INFO_ID
Message	SR: Session closed - ip = <ip address> sesid = <id> .
Message Type	Informational
Recommended Action	Reducer session to device <ip address> has ended. If this is not user triggered action such as policy change or reboot, then check network connectivity to the device. The log file on the system provides additional information.
Message ID	904: PN_REDUCER_OVERFLOW_INFO_IND
Message	SR: Reducer buffer is reaching full capacity
Message Type	Informational
Recommended Action	If the situation persists, reduce the traffic entering the reducer. Appropriate traffic filter may also be used to reduce the amount of packets to be processed by reducer.
Message ID	1002: PN_ASSEMBLER_END_SESSION_INFO_ID
Message	SA: Session closed - ip = <ip address> sesid = <id> .
Message Type	Informational
Recommended Action	Assembler session to device <ip address> has ended. If this is not user triggered action such as policy change or reboot, then check network connectivity to the device. The log file on the system provides additional information
Message ID	1102: PN_REGISTER_PRIMARY_SELFREG_ERROR_ID
Message	REG: Self registration failed. IP = <ip address> .
Message Type	Error
Recommended Action	Check the network connectivity to primary registration server <ip address> .
Message ID	1103: PN_REGISTER_SEC_SELFREG_ERROR_ID
Message	REG: Self registration failed for secondary registration server. IP = <ip address> .
Message Type	Error
Recommended Action	Check the network connectivity to secondary registration server <ip address> .
Message ID	1104: PN_REGISTER_PRIMARY_SELFREG_INFO_ID
Message	REG: Self registration completed. IP = <ip address> .
Message Type	Informational
Recommended Action	None
Message ID	1105: PN_REGISTER_SEC_SELFREG_INFO_ID
Message	REG: Self registration completed for secondary registration server. IP = <ip address> .
Message Type	Informational
Recommended Action	None



Message ID	1106: PN_REGISTER_PASSWORD_MISMATCH_ERROR_ID
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 console.
Message ID	1202: PN_BRIDGE_GENERIC_HARDENING_ERROR_ID
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 contact Technical Support to further analyze the anomaly.
Message ID	1203: PN_BRIDGE_LOCAL_LINK_UP_INFO_ID
Message	Local interface: Link Up, < speed > , < duplex mode >
Message Type	Informational
Recommended Action	None
Message ID	1204: PN_BRIDGE_LOCAL_LINK_DOWN_INFO_ID
Message	Local interface: Link Down
Message Type	Informational
Recommended Action	Verify that the link state change was not due to a network error.
Message ID	1205: PN_BRIDGE_REMOTE_LINK_UP_INFO_ID
Message	Remote interface: Link Up, < speed > , < duplex mode >
Message Type	Informational
Recommended Action	None
Message ID	1206: PN_BRIDGE_REMOTE_LINK_DOWN_INFO_ID
Message	Remote interface: Link Down
Message Type	Informational
Recommended Action	Verify that the link state change was not due to a network error.
Message ID	1402: PN_CTRL_BACKUP_PRIMARY_DOWN_ENGAGED_ERROR_ID
Message	BACKUP: No response from Primary device IP = < ip address > . Backup is engaged.
Message Type	Error
Recommended Action	Heartbeats missed from Primary device. Please check the health of the primary.
Message ID	1403: PN_CTRL_BACKUP_PRIMARY_DOWN_NOTENGAGED_ERROR_ID
Message	BACKUP: No response from Primary device IP = < ip address > . Failed to engage backup.
Message Type	Error
Recommended Action	Heartbeats missed from Primary device. Please check the health of the primary device. The log file on the system provides additional information on the failure to engage the backup device (startup configuration file not available etc.).
Message ID	1404: PN_CTRL_BACKUP_PRIMARY_UP_DISENGAGED_INFO_ID
Message	BACKUP: Response received from Primary device IP = < ip address > . Backup is disengaged.
Message Type	Informational
Recommended Action	None. The connectivity to primary device is restored.

Message ID	1702: PN_MGMT_STARTUP_CONFIG_SAVED_ID
Message	SaveStartupConfig: Saved successfully
Message Type	Informational
Recommended Action	Verify that someone authorized to configure the system saved the configuration.
Message ID	1703: PN_MGMT_CONFIG_SAVE_FAILURE_ID
Message	SaveConfig: Cannot save < module > settings: status = < status >
Message Type	Error
Recommended Action	Contact Technical Support.
Message ID	1802: PN_INIT_IN_SAFE_MODE_ID
Message	Safe-mode suspend: case 2
Message Type	Critical Error*
Recommended Action	Contact Technical Support.  Note that this message is also sent if you explicitly reboot the system into Safe Mode from the Web console or the Command Line Interface (CLI).
Message ID	1803: PN_INIT_COLD_START_ID
Message	Cold Start
Message Type	Informational
Recommended Action	If the WX device restarted unexpectedly, please investigate the reason. Contact Technical Support if there seems to be a problem.
Message ID	1902: PN_SECURITY_LOGIN_FAILURE_ID
Message	Login failed: access = < method > user = < name > IP = < ip-addr >
Message Type	Error
Recommended Action	The message has the access method (CONSOLE, SSH, or WEB) 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 ID	1903: PN_SECURITY_LOGIN_SUCCESS_ID
Message	Login ok: access = < method > user = < name > IP = < ip-addr >
Message Type	Informational
Recommended Action	Please verify that the person who logged in was someone authorized to configure the system.
Message ID	2002: PN_FAN_FAILURE_ERROR
Message	Fan Error (CPU or Chassis fan not operational).
Message Type	Error
Recommended Action	CPU or chassis fan may not be working, and may need to be replaced.
Message ID	2003: PN_FAN_SPEED_VARIATION_ERROR
Message	Fan Speed Error (Cpu or Chassis speed variation).
Message Type	Error
Recommended Action	The fan may need to be replaced.

Message ID	2004: PN_FAN_OK
Message	Fan OK (CPU or Chassis). Fan OK (CPU or Chassis fan speed normal)
Message Type	Informational
Recommended Action	The CPU or chassis fan has recovered from a previous failure, such as fan speed variation or fan failure error. No action is required.
Message ID	2102: PN_MULTINODE_MASTER_NODE_UP_ID
Message	SR: Multi-Node Master Node is Up
Message Type	Informational
Recommended Action	This message indicates that the master node is up. No action.
Message ID	2103: PN_MULTINODE_MASTER_NODE_DOWN_ID
Message	SR: Multi-Node Master Node is Down
Message Type	Error
Recommended Action	If this master node was not taken down intentionally, check the running configuration and the network connectivity.
Message ID	2104: PN_MULTINODE_LAST_NODE_UP_ID
Message	SR: Multi-Node Last Node is Up
Message Type	Informational
Recommended Action	This message indicates that the last node is up. No action is required.
Message ID	2105 (PN_MULTINODE_LAST_NODE_DOWN_ID)
Message	SR: Multi-Node Last Node is Down
Message Type	Error
Recommended Action	If this last node was not taken down intentionally, check the running configuration and the network connectivity.
Message ID	2501: PN_IPSEC_GENERIC_ERROR_ID
Message	One of the following: <ul style="list-style-type: none"> <li>■ IPsec Added SA &lt; source IP address &gt; -&gt; &lt; destination IP address &gt; s : SPI &lt; SPI number &gt; &lt; encryption algorithm &gt; &lt; authentication algorithm &gt; Hours Remaining 24.00 MBytes Remaining 100.00</li> <li>■ IPsec Expired SA &lt; source IP address &gt; -&gt; &lt; destination IP address &gt; s : SPI &lt; SPI number &gt; due to &lt; "life time" or "data life time" &gt;</li> <li>■ IPsec Deleted SA &lt; source IP address &gt; -&gt; &lt; destination IP address &gt; s : SPI &lt; SPI number &gt;</li> </ul>
Message Type	Informational
Recommended Action	These messages indicate when IPsec security associations are added, expired, and deleted. No action is required.
Message ID	2601: PN_MISC_DISK_FAILURE_ID
Message	Disk < /ata2 or /ata3 > failed initialization! WARNING: Disk disabled - performance will degrade
Message Type	Informational
Recommended Action	NSC continues without this disk, but with some loss in the percentage of data reduction. If both disks fail, NSC reverts to MSR. No action is required.

Message ID	2801: PN_WP_GENERIC_ERROR_ID
Message	One of the following: <ul style="list-style-type: none"><li>■ WP: ***** Unacceptable Performance detected due to LOSS on Path, Path Ip = &lt; remote IP address &gt; *****</li><li>■ WP: ***** Unacceptable Performance detected due to LATENCY on Path, Path Ip = &lt; remote IP address &gt; *****</li><li>■ WP: ***** Acceptable Performance detected on Path, Path Ip = &lt; remote IP address &gt; *****</li></ul>
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 NetFlow packets and performance data that can be exported by a WX device, and covers the following sections:

- NetFlow Version 5 Export on page 405
- Performance Statistics Export on page 406
- Top Traffic Export on page 414

## NetFlow Version 5 Export

Traffic data can be exported to a Cisco NetFlow server in Version 5 format (refer to “Traffic Statistics” on page 256).

Table 22 describes the NetFlow packet header.

**Table 22: 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 WX 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 23 describes each traffic flow entry in a NetFlow packet (up to 30 entries per packet).

**Table 23: 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 exported in CSV format (refer to “Exporting Performance Data” on page 279).

- “General Device Information” in the next section
- “Data Section Information” on page 407
- “System Session Statistics” on page 408
- “Reduction Session Statistics” on page 410
- “Application Session Statistics” on page 410
- “Bandwidth Management Statistics” on page 412
- “Inbound Traffic By Port Statistics” on page 413

## General Device Information

Table 24 describes the exported general device information.

**Table 24: General Device Information**

Parameter	Description
Device IP	IP address of the WX device.
Software version	Version of WXOS software that was running when the statistics were exported.
Serial number	Serial number of the WX device that exported the statistics.
License speed	Licensed speed of the WX 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 25 describes the data section information that precedes the set of statistic tables for each exported time range.

**Table 25: Data Section Information**

Parameter	Description
< time > data section	Indicates the time range for the statistics tables that follow: <ul style="list-style-type: none"> <li>■ This hour</li> <li>■ Last hour</li> <li>■ Today</li> <li>■ Yesterday</li> <li>■ This week</li> <li>■ Last week</li> </ul>
ip =	IP address of the WX device.
device local time =	Local date and time of the export.
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 26 describes the exported system session statistics.

**Table 26: 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 (WX device).
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.
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 WX 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 WX 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 WX 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 WX 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.



Parameter	Description
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. Data input is the number of IP bytes into the WX device from the Local port.
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 WX device, as follows: <ul style="list-style-type: none"> <li>■ PkIn1 Less than 64 bytes</li> <li>■ PkIn2 64 to 127</li> <li>■ PkIn3 128 to 255</li> <li>■ PkIn4 256 to 511</li> <li>■ PkIn5 512 to 1023</li> <li>■ PkIn6 More than 1023 bytes</li> </ul>
PkOut1 to PkOut6	Six fields that show the number of packets in each of six packet-size ranges for traffic out of the WX device, as follows: <ul style="list-style-type: none"> <li>■ PkOut1 Less than 64 bytes</li> <li>■ PkOut2 64 to 127</li> <li>■ PkOut3 128 to 255</li> <li>■ PkOut4 256 to 511</li> <li>■ PkOut5 512 to 1023</li> <li>■ PkOut6 More than 1023 bytes</li> </ul>

## Reduction Session Statistics

Table 27 describes the reduction session CSV exported statistics.

**Table 27: 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 WX 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 Ipcsec	Number of packets that were identified for encryption and intended for the destination IP address.
Packets Out of Ipcsec	Number of encrypted packets sent to the destination IP address.
Packets Dropped by Ipcsec	Number of packets intended for the destination IP address that were dropped according to the default IPsec policy.
Ipcsec Overhead	Number of bytes added by IPsec processing.

## Application Session Statistics

Table 28 describes the exported application session statistics.

**Table 28: 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 WX 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 WX 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 WX 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.

Parameter	Description
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 NSC (WXC devices only), and sent to the WX device listed with the destination IP address and application ID.

## WAN Statistics

Table 29 describes the exported WAN statistics.

**Table 29: 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 WX device that sent or received data from this device.
Bytes From WAN	Number of bytes received from the WAN for the remote WX device and application.
Bytes To WAN	Number of bytes sent to the WAN for the remote WX device and application.

## Application Flow Acceleration Statistics

Table 30 describes the exported Application Flow Acceleration statistics.

**Table 30: 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 WX 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 31 describes the bandwidth management statistics collected per application class for each reduction tunnel.

**Table 31: Bandwidth Management Statistics**

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Tunnel	<b>Outbound bandwidth management:</b> The IP address of the destination assembler or the default allocation.
	<b>Inbound bandwidth management:</b> The parameter is Inbound.
Class	<b>Outbound bandwidth management:</b> The bandwidth class ID, which is a collection of applications that a user has mapped to the class.
	<b>Inbound bandwidth management:</b> One of the four pre-defined classes (i.e., Reduced, Intranet, TCP or Default).
Bytes In	<b>Outbound bandwidth management:</b> The total number of application bytes into the WX device.
	<b>Inbound bandwidth management:</b> The total number of bytes into the Remote interface of the WX device by class.
Bytes Out	<b>Outbound bandwidth management:</b> The total number of application bytes out of outbound bandwidth management.
	<b>Inbound bandwidth management:</b> the total number of bytes out of inbound bandwidth management.
Bytes Dropped	<b>Outbound bandwidth management:</b> The total number of application bytes dropped by the bandwidth management feature.
	<b>Inbound bandwidth management:</b> The total number of bytes dropped by the bandwidth management feature.
Packets In	<b>Outbound bandwidth management:</b> The total number of application packets into the WX device.
	<b>inbound bandwidth management:</b> The total number of packets passed into the device by inbound bandwidth management.
Packets Out	<b>Outbound bandwidth management:</b> The total number of application packets transmitted by the device. (The total number does not include meta packetization.)
	<b>Inbound bandwidth management:</b> The total number of packets out of inbound bandwidth management.
Packets Dropped	<b>Outbound bandwidth management:</b> The total number of application packets dropped by the bandwidth management feature.
	<b>Inbound bandwidth management:</b> The total number of packets dropped by the bandwidth management feature.

## WAN Performance Statistics

Table 32 describes the exported WAN performance statistics.

**Table 32: 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 WX 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 WX 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.
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 33 describes the Inbound traffic by port statistics.

**Table 33: 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 33 describes the Traffic statistics exported to the *ip-flow.csv* file.

**Table 34: Top Traffic 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 WX 15, 32K for WX 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.
BytesRcvd	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"> <li>■ <b>RA</b>. Reduced application (matched an application definition)</li> <li>■ <b>RO</b>. Reduced undefined application</li> <li>■ <b>PT</b>. Passed through due to policy setting</li> <li>■ <b>U</b>. Unknown passthrough traffic, such as non-TCP/UDP traffic</li> </ul>
StartTime	Start date and time of traffic flow.
EndTime	End date and time of traffic flow.

## Appendix D

# Common Application Port Numbers

The following table lists common application port numbers, as listed by the Internet Assigned Numbers Authority (IANA, <http://www.iana.org/assignments/port-numbers>).



**NOTE:** WX devices reserve port numbers 3577 and 3578 for TCP and UDP data transmission.

**Table 35: 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

Keyword	Port Number	Protocol	Description
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



## Appendix E

# Safety and EMC Certifications

The following table lists the safety and EMC certifications for each type of WX device.

**Table 0-1 Safety and EMC Certifications for WX Devices**

Description	WX 15	WX 20	WX 50/60	WX 100	WXC 250	WXC 500
<b>Conformity for EMC</b>						
EN 55022:1998 + A1 + A2	X	X	X		X	X
EN 55024:1998 + A1 + A2	X	X	X		X	X
FCC (Class A) Part 15 Subpart J		X	X			
FCC (Class A) Part 15 Subpart B:2003	X			X	X	X
EN 61000-3-2:2000		X		X	X	X
EN 61000-3-3:1995 + A1		X	X	X	X	X
<b>Safety Standard (CB Scheme)</b>						
IEC 60950-1:2001	X	X	X	X	X	X
<b>cTUVus Canadian/US safety</b>						
UL 60950:2000				X	X	X
UL 60950-1:2003			X	X		X
CAN/CSA-C22.2 No 60950-1-03			X	X		X
<b>TUV-GS (German safety)</b>						
EN 60950:1992 + A1 + A2 + A3 + A4 + A11		X	X			
EN 60950:2000				X	X	X
EN 60825-1:1994 + A1 + A2 (lasers)				X		
UL 60950-1:2001 + A11			X			X
<b>Gost</b>		X	X	X		


## Product Reclamation and Recycling Program

---

Juniper Networks is committed to environmentally responsible behavior. As part of this commitment, we work to comply with environmental standards such as the European Union's *Waste Electrical and Electronic Equipment* (WEEE) Directive and *Restriction of Hazardous Substances* (RoHS) Directive.

These directives and other similar regulations from countries outside the European Union regulate electronic waste management and the reduction or elimination of specific hazardous materials in electronic products. The WEEE Directive requires electrical and electronics manufacturers to provide mechanisms for the recycling and reuse of their products. The RoHS Directive restricts the use of certain substances that are commonly found in electronic products today. Restricted substances include heavy metals, including lead, and polybrominated materials. The RoHS Directive, with some exemptions, applies to all electrical and electronic equipment.

In accordance with Article 11(2) of Directive 2002/96/EC (WEEE), products put on the market after 13 August 2005 are marked with the following symbol or include it in their documentation: a crossed-out wheeled waste bin with a bar beneath.



Juniper Networks provides recycling support for our equipment worldwide to comply with the WEEE Directive. For recycling information, send e-mail to [recycling@juniper.net](mailto:recycling@juniper.net) indicating the type of Juniper Networks equipment that you wish to dispose of and the country where it is currently located, or contact your Juniper Networks account representative.

Products returned through our reclamation process are recycled, recovered, or disposed of in a responsible manner. Our packaging is designed to be recycled and should be handled in accordance with your local recycling policies.

## Appendix F

# Safety Recommendations and Warnings



### Power Cable Warning (Japanese)

---



#### 注意

附属の電源コードセットはこの製品専用です。  
他の電気機器には使用しないでください。

---

The preceding translates as follows:

#### Warning

The attached power cable is only for this product. Do not use the cable for another product.

### VCCI Compliance

The following VCCI compliance information applies to the WX/WXC product that meets VCCI Class A limits.

#### クラスA情報技術装置

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

The preceding translates as follows:

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

## ***Lightning Activity Warning***



**CAUTION:** Do not work on the system or connect or disconnect cables during periods of lightning activity.

---

## ***Jewelry Removal Warning***



**CAUTION:** Before working on equipment that is connected to power lines, remove jewelry (including rings, necklaces, and watches). Metal objects will heat up when connected to power and ground and can cause serious burns or weld the metal object to the terminals.

---

## ***Installation Warning***



**CAUTION:** Read the installation instructions before you apply power to the system.

---

## ***IT Power Statement***

The device is designed to work with IT power systems.

## ***SELV Circuit Warning***



**CAUTION:** The ports labeled "Ethernet," "Local/Remote," "Console," and "USB" are safety extra-low voltage (SELV) circuits. SELV circuits should only be connected to other SELV circuits. Avoid connecting the SELV circuit to the telephone network voltage (TNV) circuits.

---

## ***Circuit Breaker (15A) Warning***



**CAUTION:** This product relies on the building's installation for short-circuit (overcurrent) protection. Ensure that a fuse or circuit breaker no larger than 120 VAC, 15A U.S. (240 VAC, 10A international) is used on the phase conductors (all current-carrying conductors).

---

### ***Grounded Equipment Warning***



**CAUTION:** This equipment is intended to be grounded. Ensure that the host is connected to earth ground during normal use.

---



**VARNING!** Apparaten skall anslutas till jordat uttag när den ansluts till ett nätverk.

---



**ADVARSEL** Apparatet må tilkoples jordet stikkontakt.

---



**VAROITUS** Laite on liitettävä suojamaadoituskoskettimilla varustettuun pistorasiaan.

---

### ***Class 1 Laser Product Warning***



**CAUTION:** Class 1 laser product.

---

### ***Laser Beam Warning***



**CAUTION:** Do not stare into the beam or view it directly with optical instruments.

---

### ***Battery Warning***



**CAUTION:** WX and WXC devices have no user serviceable parts. Opening the device voids the warranty. As a safety caution, note that opening the chassis exposes a lithium battery. If you attempt to remove or replace the lithium cell, do not use a conductive instrument, as a short-circuit may cause the cell to explode. A replacement cell must be of the same type (CR2032). Dispose of a spent cell promptly—do not recharge, disassemble, or incinerate. Keep cells away from children.

---

## Rack Mounting of Systems

Observe the following precautions for rack stability and safety. Also, refer to the rack installation documentation provided with the rack and the WX/WXC device for specific caution statements and procedures.

Systems are considered to be components in a rack. Thus, "component" refers to any system as well as to various peripherals or supporting hardware.



**CAUTION:** Installing systems in a rack without the front and side stabilizers installed could cause the rack to tip over, potentially resulting in bodily injury. Therefore, always install the stabilizers before installing components in the rack.

After installing systems/components in a rack, never pull more than one component out of the rack on its slide assembly at one time. The weight of more than one extended component could cause the rack to tip over and may result in serious injury.

---



**NOTE:** Your system is safety-certified as a free-standing unit and as a component for use in a standard rack cabinet using the customer rack kit. It is your responsibility to have the final combination of system and rack kit in a rack cabinet evaluated for suitability by a certified safety agency. Juniper Networks disclaims all liability and warranties in connection with such combinations.

---

System rack kits are intended to be installed in a rack by trained service technicians. If you install the kit in any other rack, be sure that the rack meets the specifications of a standard 19" rack.

## Anti-static Precautions



**CAUTION:** This product contains static-sensitive components and should be handled with care. It is recommended that the product be handled in a Special Handling Area as defined in EN100015-1:1992. Such an area has working surfaces, floor coverings, and chairs connected to a common earth reference point. A grounded wrist strap should be worn during handling. Failure to employ adequate anti-static measures can cause irreparable damage to components on the memory board, such as the processor and memory modules.

---

# Glossary

<b>access control list</b>	List of IP addresses from which an administrator can login to a WX device.
<b>assembly</b>	Process by which a WX device re-assembles reduced traffic into its original form.
<b>auto-negotiation</b>	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.
<b>bandwidth</b>	The amount of data that can be sent through a network connection, measured in bits per second (bps).
<b>bridge</b>	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.
<b>CLI</b>	See <i>command line interface</i> .
<b>command line interface (CLI)</b>	A method of configuring the WX device by typing in commands via the local serial interface or remote SSH session.
<b>community</b>	Two or more WX devices that can reduce and assemble data for each other. Initially, all WX devices belong to the Default community. Each WX device contacts the registration server to identify the other devices in the same community.
<b>endpoint</b>	WX device. When you install a WX device in your network, the device's auto-discovery function locates all other devices in the community and exchanges network information with each device.
<b>filter</b>	Operator defined IP addresses or TCP port numbers that determine valid addresses or applications for reduction processing. A single filter or a list of filters can be defined for each system.
<b>full-duplex</b>	A mode of operation that enables a pair of systems connected by a link to transmit frames to one another at the same time.
<b>gateway</b>	A device that connects and forwards packets between computers or different networks. See also, <i>router</i> .
<b>half-duplex</b>	A mode of operation that allows only a single station to successfully transmit a frame at a given time.
<b>hardware passthrough</b>	Hardware-driven process by which all traffic is passed through the WX device at wire-speed. It is invoked automatically upon disruption.

<b>HTTP</b>	Hyper Text Transfer Protocol. The protocol most often used to transfer information from World Wide Web servers to browsers.
<b>ICMP</b>	Internet Control Message Protocol. An Internet Protocol used to communicate between devices on a network to manage errors and generate control messages.
<b>Interior Gateway Protocol (IGP)</b>	A group of protocols that provide routing information to the routers within an autonomous network.
<b>Internet Protocol (IP)</b>	The protocol that is used to route a data packet from its source to its destination over the Internet.
<b>IP address</b>	A numeric address, such as 10.10.187.22, assigned to every device on the network.
<b>IP subnet mask</b>	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.
<b>IP subnet</b>	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.
<b>latency</b>	The time necessary for a packet of data to travel from a source to a destination across a network.
<b>local port</b>	Ethernet port on the back of the WX device. Use to connect to a LAN aggregating switch. <i>See also, remote port.</i>
<b>log</b>	A record of device activity. Logs are recorded for system information, performance, backup, and recovery.
<b>MIB</b>	Management Information Base. A database containing ongoing configuration information and statistics of a device in a network. MIBs are used with SNMP.
<b>MTU</b>	Maximum Transmission Unit. The largest size packet that can be transmitted by a device on a network.
<b>netmap</b>	Reduction subnets advertised by each WX device. Each WX device dynamically adjusts its advertised subnets to exclude unreachable addresses. In this case, multiple remote routes are advertised for the same subnet to exclude unresponsive addresses.
<b>operator interface</b>	The front-panel keypad and LCD, a local terminal via the serial interface, a remote terminal via the web console, or a remote terminal via the ssh.
<b>OSPF</b>	Open Shortest Path First. An interior gateway protocol that routes messages according to the least expensive path.
<b>packet</b>	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.
<b>passthrough mode</b>	A function of the WX device where data passes through at wire-speed upon device disruption or overflow.



<b>ping</b>	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.
<b>reduction rate</b>	The rate of data reduction in percentage of a WX device.
<b>reduction subnets</b>	The subnets for which a WX device can assemble reduced data. Each WX device advertises its reduction subnets to the other devices in the community.
<b>registration server</b>	The WX device that stores the network information for the WX devices in each community. Each device periodically contacts the registration server to identify the other devices in the same community.
<b>remote port</b>	Ethernet port on the back of the WX device. Used to connect to WAN router Ethernet port. <i>See also, local port.</i>
<b>response time</b>	The time it takes for a host to respond to a user command.
<b>RIP</b>	<i>See Routing Information Protocol.</i>
<b>round-trip time (RTT)</b>	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.
<b>router</b>	Specialized computer that forwards data packets between networks. Routers can exchange information about their network connectivity (or accessibility) with neighboring network routers using standard routing protocols. This information is used by the router to determine an optimal path for a packet being forwarded.
<b>Routing Information Protocol (RIP)</b>	An interior gateway protocol used in IP networks.
<b>Secure Shell</b>	A program used for secure remote login to a WX device.
<b>Simple Network Management Protocol (SNMP)</b>	The Internet standard protocol for network management software.
<b>Simple Network Time Protocol (SNTP)</b>	A protocol that can synchronize clocks on local computers with time clocks on the Internet.
<b>software passthrough</b>	Software-driven process by which a WX device transparently passes packets through the system in lieu of processing (reducing).
<b>SSH</b>	see Secure Shell.
<b>static IP address</b>	A permanent IP address for a client, server, or other network device.
<b>Switch</b>	A networking device that sends packets directly to a port associated with a given network address.
<b>TCP</b>	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.
<b>tunneling</b>	Encapsulating one type of packet inside the data field of another packet.

<b>User Datagram Protocol (UDP)</b>	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.
<b>warm reboot</b>	A reboot of the WX device without powering off the unit.
<b>Web Console</b>	A method for configuring and monitoring the statistics of the WX device using a Web browser.

# Index

## Numerics

- 3DES encryption for IPSec .....217, 223, 322
- 802.1q VLAN support.....61, 319

## A

- AAA settings .....79, 300
- acceleration
  - Active Flow Pipelining .....196, 197, 302
  - CIFS traffic .....204, 302
  - Exchange traffic .....207, 302
  - Fast Connection Setup.....196, 198, 302
  - Forward Error Correction.....196, 302
  - HTTP traffic.....209, 302
  - reports .....248
- access control lists.....361
- access control log file .....278, 369
- Active Flow Pipelining
  - clusters .....307
  - configuring .....196, 197, 302
  - feature/topology setting.....344
  - report .....248
- active FTP .....93
- advertising reduction subnets .....131, 350
- AES encryption for IPSec.....217, 223, 322
- ageout time, device .....354
- aggregate local WAN speed .....156
- Application Flow Acceleration
  - about.....200
  - CIFS and Exchange reports.....253
  - CIFS traffic .....204, 302
  - Exchange traffic .....207, 302
  - HTTP reports.....254
  - HTTP traffic.....209, 302
- applications
  - about application definitions .....89
  - accelerating
    - Active Flow Pipelining .....197, 306
    - CIFS traffic .....204, 302
    - Exchange traffic .....207, 302
    - Fast Connection Setup.....198, 306
    - HTTP traffic.....209, 302
  - assigning to traffic classes.....95
  - common port numbers .....415

- defining
  - manually .....92, 309
  - using the Traffic report.....257
- defining gateways for .....349
- monitoring .....97, 135, 325
- reducing.....135, 317
- summary statistics
  - all traffic .....240
  - WAN traffic .....230
  - visibility in tunnels .....146, 347
- ARP, configuring.....103, 312
- assemblers
  - default.....141, 346
  - preferred.....144, 347
- asymmetric routing support for AFP.....307
- authentication methods, selecting.....80, 300
- automatic installation .....30

## B

- backup devices, configuring .....312
- bandwidth management
  - inbound .....184, 335
  - outbound, see "outbound QoS"
- baud rate
  - default .....41
  - setting .....315
- BGP routes, polling from a router .....360
- boot images
  - activating .....270, 295
  - loading .....267, 299
- browser support .....45
- buttons
  - Bypass/Disable.....37
  - front panel.....35
- bypass condition, Multi-Path .....119, 328
- Bypass/Disable button .....37
- bypass/disable command .....43, 291
- bytes graph .....239

## C

- caching, for HTTP acceleration
  - about .....202
  - configuring .....307

carving out unreachable addresses		route-poll .....	359
and outbound QoS .....	162	security .....	360
enable/disable .....	352	snmp .....	361
certifications .....	417	sntp .....	362
CIFS acceleration		stack-group .....	363
about .....	201	syslog .....	364
configuring .....	204, 302	top-talker .....	365
reports .....	253	wan-performance-monitor .....	366
circuit speeds		file management	
and router overhead .....	156	copy .....	290
configuring .....	165, 178, 339	list .....	292
Citrix names, in application definitions .....	94	remove .....	296
classes, traffic		show commands	
inbound .....	184, 335	show access-log .....	369
outbound .....	167, 173, 338	show flow-details .....	372
outbound QoS and Multi-Path .....	95	show log .....	375
CLI commands		top level .....	290
about		commit .....	290
basics of using .....	285	embed .....	291
command modes .....	286	flow-details .....	372
command summary .....	287	import-route-table .....	291
entering commands		load-config .....	292
from a file .....	299	packet-capture .....	293
from a terminal or SSH program .....	283	ping .....	294
from the Web console .....	275	reboot .....	295
configuration		rollback .....	296
aaa .....	300	save-config .....	297
acceleration .....	302	set .....	298
application .....	309	shutdown .....	298
arp .....	312	source .....	299
backup .....	312	support .....	299
clock .....	315	traceroute .....	300
console .....	315	upgrade .....	299
dns .....	316	client devices	
filter .....	316	client-mode command .....	363
interface .....	318	connecting .....	43
ip .....	320	disconnecting .....	44
ipsec .....	320	clusters, AFP .....	307
license .....	324	CMS	
mon-apps .....	325	about .....	28
multi-path .....	326	exclusive access to WX devices .....	360
ospf .....	330	command modes .....	286
packet-interception .....	331	communities, defining .....	75, 352
prime-time .....	333	community topology .....	99, 344
profile-mode .....	334	configuration file	
qos inbound .....	335	displaying .....	265, 297
qos outbound .....	337	loading .....	266, 292
radius .....	342	saving .....	263, 297
reduction .....	343	setting to the factory default .....	268, 292
reduction-subnet .....	350	congestion control .....	166, 180, 339
reg-server .....	352	connect timeout, registration server .....	354
remote-routes .....	355	console port	
rip .....	356	baud rate, setting .....	315
route .....	356	DB9 cable pin-outs .....	395

default settings ..... 41, 284  
 copying files ..... 290  
 CSV, interpreting results ..... 405

## D

data packets, Forward Error Correction ..... 197, 305  
 data reduction statistics  
   bytes graph ..... 239  
   peak data reduction ..... 238, 261  
   viewing ..... 237  
 dead interval, OSPF ..... 331  
 dead-time interval, RADIUS ..... 83, 342  
 dedicated WANs ..... 157  
 default assemblers ..... 141, 346  
 default gateway, configuring  
   in CMS ..... 30  
   in front panel ..... 35  
   in Web/CLI ..... 57, 320  
 default IPsec policy ..... 320  
 Default traffic class  
   inbound QoS ..... 184, 335  
   outbound QoS ..... 173, 338  
   outbound QoS and Multi-Path ..... 95  
 default user name and password ..... 41  
 deployment, examples ..... 22  
 device configuration  
   displaying ..... 265, 297  
   loading ..... 266, 292  
   saving ..... 263, 297  
   setting to the factory default ..... 268, 292  
 device names ..... 58, 298  
 diagnostic files, generating ..... 280, 299  
 disk access policy, NSC ..... 346  
 disk icons ..... 134  
 diversion settings, Multi-Path ..... 121, 329  
 DNS servers, configuring for the Traffic report ..... 59  
 domain names in the Traffic report  
   configuring ..... 59, 316  
   viewing ..... 258  
 downgrading to a previous release ..... 267, 296  
 DSCP values, see "ToS/DSCP values"  
 dynamic resource allocation (DRA), configuring ..... 348  
 dynamic routes  
   importing from a file ..... 71, 291  
   polling from a router ..... 70, 359  
   using OSPF ..... 69, 330  
   using RIP ..... 69, 356

## E

EMC certifications ..... 417  
 encryption, see "IPsec"  
 endpoints  
   IPsec ..... 216, 322  
   Multi-Path ..... 120, 328

NSC ..... 134, 346  
 outbound QoS ..... 176, 339  
 Packet Flow Acceleration ..... 194, 305  
 reduction ..... 129, 345  
 summary report ..... 258  
 WAN performance monitoring ..... 123, 366  
 erasing the disks ..... 269, 293  
 Ethernet ports, connecting the cables ..... 32, 38  
 Exchange acceleration  
   about ..... 201  
   configuring ..... 207, 302  
   reports ..... 253  
 Executing report ..... 260  
 exporting data  
   device performance statistics ..... 279  
   interpreting performance results ..... 405  
   packet capture ..... 273, 293  
   secondary registration server database ..... 354  
   traffic statistics ..... 256, 365  
 external routing for packet interception ..... 106, 331

## F

factory default configuration ..... 268, 292  
 Fast Connection Setup  
   configuring ..... 196, 198, 302  
   report ..... 250  
 fast reduction tunnels ..... 348  
 features/topology, configuring ..... 344  
 filters, reduction  
   configuring application ..... 135, 316  
   source/destination ..... 101, 316  
 firewall requirements ..... 30  
 flow details, viewing ..... 372  
 Forward Error Correction  
   configuring ..... 196, 302  
   report ..... 252  
 fragments, reducing ..... 318  
 front panel  
   securing ..... 87, 360  
   using the buttons ..... 35  
 FTP application type ..... 93, 310  
 FTP servers, using  
   to copy a packet capture ..... 294  
   to copy system files ..... 290  
   to export diagnostic files ..... 299  
   to import routes ..... 71, 291  
   to load a boot image ..... 267, 299  
   to load configuration files ..... 266, 292  
   to pre-sync files with NSC ..... 346  
   to pre-sync files with NSM ..... 147  
   to roll back a boot image ..... 296  
   to save configuration files ..... 263, 297

**G**

gateways, configuring	
application .....	349
default	
in CMS .....	30
in front panel .....	35
in Web/CLI .....	57, 320
in Multi-Path configurations .....	117, 327
guaranteed bandwidths	
configuring .....	169, 175, 338
overriding .....	172

**H**

hardware passthrough, disabling .....	37, 43, 291
heartbeat packets	
for all reduction tunnels .....	347
for high-loss tunnels .....	308
hello interval, OSPF .....	331
high-availability support .....	59
HMAC/SHA-1 authentication for IPSec .....	217, 223, 322
HTTP acceleration	
about .....	202
configuring .....	209, 302
reports .....	254
Hub and Spoke topology .....	344

**I**

IANA port map .....	258, 325
ICMP redirect age-out setting .....	357
icons	
disk .....	134
endpoint and tunnel .....	131
IPSec status .....	221
Multi-Path status .....	121
on EndPoints Summary report .....	259
idle user timeout .....	85, 300
importing routes	
by polling a router .....	70, 359
from a file .....	71, 291
inbound QoS .....	336
configuring .....	184, 335
report .....	246
inbound speed .....	186
inline deployment .....	30
installation	
automatic .....	30
inline and off-path .....	30
post-install tasks .....	53
pre-install tasks .....	29
WX 100 .....	31, 37
interface	
link failure propagation .....	60, 318
manual mode test .....	61, 319
periodic mode test .....	319

## settings, configuring

in front panel .....	36
in Web/CLI .....	59, 318
statistics .....	318, 373
Intranet traffic class .....	184, 335
IP address, configuring	
in CMS .....	30
in front panel .....	35
in Web/CLI .....	57, 320
secondary address, Multi-Path .....	116, 327
IP compression, meta-packet	
configuring .....	145, 347
firewall requirements .....	30
IPSec	
configuration procedure .....	215, 320
defining endpoints .....	216, 322
defining templates .....	222, 321
using the Setup Wizard .....	215

**J**

JVM support .....	45
-------------------	----

**K**

keep-alive packets	
for all reduction tunnels .....	347
for high-loss tunnels .....	308
key lifetimes .....	223, 322
keyboard shortcuts, CLI .....	285
keys	
IPSec .....	223
OSPF .....	69, 330
RADIUS .....	83, 342

**L**

LAN-WAN routing check .....	348
latency threshold	
Multi-Path .....	121, 329
WAN performance monitoring .....	124, 366
Layer 2 multicast traffic .....	243
LEDs, checking	
WX 100 .....	36, 42
license keys, entering .....	62, 324
lifetimes, IPSec key .....	223, 322
link failure propagation .....	60, 318
load balancing	
across routers	
route-based .....	73, 357
router-based using ToS .....	358
across WX devices .....	139, 347
loading a boot image .....	267, 299
local domain name .....	59, 316
local routes	
about .....	66
adding static .....	68, 356

- from OSPF ..... 69, 330
- from RIP ..... 69, 356
- importing from a file ..... 71, 291
- polling from a router ..... 70, 359
- local users, defining ..... 84, 300
- log files
  - access control ..... 278, 369
  - system ..... 277, 375
- logging in
  - CLI ..... 284
  - Web console ..... 55
- login retries, SSH ..... 81, 302
- loss threshold
  - Multi-Path ..... 329
  - WAN performance monitoring ..... 366

## M

- MAC addresses
  - in ARP entries ..... 103, 312
  - in assembled packets ..... 350
- management traffic, encrypting ..... 322
- manual and automatic installation ..... 30
- marking methods, Multi-Path ..... 116, 329
- maximum bandwidths
  - inbound ..... 186, 336
  - outbound
    - configuring ..... 169, 175, 338
    - overriding ..... 172
- MD5
  - for IPSec ..... 217, 223, 322
  - for OSPF ..... 330
- Mesh topology ..... 344
- meta packets
  - configuring size and wait time ..... 349
  - disabling multi-packet ..... 349
  - IP compression ..... 145, 347
  - wait time ..... 349
- minimum WAN speed ..... 166, 180, 339
- Molecular Sequence Reduction, see "MSR"
- monitor settings ..... 45
- monitoring
  - applications ..... 97, 135, 325
  - virtual endpoints ..... 228, 230
  - WAN performance
    - configuring ..... 123, 366
    - viewing reports ..... 231
- MSR
  - about ..... 19
  - symbol size ..... 349
- MSS override for AFP ..... 308
- multi-flow emulation ..... 146, 347
- Multi-Path configurations
  - about ..... 114
  - defining endpoints ..... 120, 328

- defining templates ..... 118, 328
- router configuration ..... 122
- viewing reports ..... 231
- multiple tunnels on the WX 100 ..... 363

## N

- names, special characters in ..... 57
- NetFlow records, generating ..... 274, 365
- network
  - cables, connecting ..... 32, 38
  - interfaces, configuring
    - in front panel ..... 36
    - in Web/CLI ..... 59, 318
  - settings, configuring
    - in front panel ..... 35
    - in Web/CLI ..... 57, 320
- Network Sequence Caching, see "NSC"
- non-WX endpoints
  - outbound QoS ..... 179, 340
- NSC
  - defining applications ..... 135, 346
  - defining endpoints ..... 134, 346
  - disk access policy ..... 346
  - file pre-synchronization ..... 146, 346
  - overflow mode ..... 346
- NTP, configuring ..... 62, 362

## O

- off-path deployment
  - configuring ..... 105, 331
  - installing ..... 30
- operator access, securing ..... 86, 361
- OSPF, configuring ..... 69, 330
- outbound QoS
  - about ..... 154
  - and Packet Flow Acceleration ..... 194
  - configuration procedure ..... 163
  - congestion control ..... 166, 180, 339
  - dedicated and oversubscribed WANs ..... 157
  - defining endpoints ..... 165, 176, 339
  - defining settings by endpoint ..... 171, 219
  - defining templates ..... 164, 174, 338
  - defining traffic classes ..... 95, 167, 173, 338
  - excluding LAN/WAN addresses ..... 180
  - non-WX endpoints ..... 179, 340
  - outbound speed
    - about ..... 156
    - defining ..... 165, 176, 177, 338
  - report ..... 244
  - starting and stopping ..... 183
  - ToS/DSCP values ..... 181, 184, 341
  - using the Setup Wizard ..... 164
- outbound speed
  - about ..... 156

- overflow mode, NSC..... 346
- overflow, traffic volume ..... 242, 343, 378
- oversubscribed WANs ..... 157
- P**
- packet age-out setting ..... 338
- packet capture
  - changing the password ..... 88, 360
  - using ..... 273, 293
- Packet Flow Acceleration
  - Active Flow Pipelining ..... 196, 197, 302
  - Fast Connection Setup ..... 196, 198, 302
  - Forward Error Correction ..... 196, 302
  - reports ..... 248
- packet fragments, reducing ..... 318
- packet interception ..... 105, 331
- packet size distribution statistics ..... 243
- pass-phrase, IPsec ..... 217, 322
- passthrough statistics ..... 242
- passwords
  - default ..... 41
  - defining ..... 84, 300
  - OSPF ..... 69, 330
  - packet capture ..... 88, 360
  - registration server ..... 76, 354
  - RIP ..... 70, 107, 356
- peak data reduction ..... 238, 261
- performance data, exporting ..... 279
- performance monitoring, WAN
  - configuring ..... 123, 366
  - viewing reports ..... 231
- periodic interface mode test ..... 319
- permanent license keys ..... 62, 324
- ping utility ..... 271, 294
- point-to-multipoint configuration ..... 23
- policy routes, defining gateways by application ..... 349
- policy-based routing for packet interception ... 106, 331
- port numbers
  - common application ..... 415
  - in application definitions ..... 94, 310, 372
  - RADIUS server ..... 83, 342
  - required for TCP and UDP ..... 30
  - viewing on Traffic by Port report ..... 325
- post-installation tasks ..... 53
- preferred assemblers ..... 144, 347
- preferred path ..... 119, 328
- pre-fetch, for HTTP acceleration
  - about ..... 202
  - configuring ..... 307
- pre-installation tasks ..... 29
- pre-synchronization, file ..... 146, 346
- primary boot image ..... 295
- prime time
  - defining ..... 104, 333
  - viewing on reports ..... 228
- privilege level, user ..... 85, 301
- Profile Mode
  - defining remote subnets ..... 334
  - virtual devices ..... 334
- protocols, in application definitions ..... 94, 310, 372
- Q**
- QoS
  - inbound
    - configuring ..... 184, 335
    - report ..... 246
  - outbound, see "outbound QoS"
- queue lengths
  - inbound QoS ..... 336
  - outbound QoS ..... 338
- queue processing by ToS/DSCP values ..... 184
- R**
- rack-mount installation
  - WX 100 ..... 32, 38
- RADIUS servers and server groups, defining ..... 82, 342
- rebooting the device ..... 270, 295
- recompression and tunnel switching
  - about ..... 148
  - enabling ..... 347
- recovery image ..... 295
- recovery packets, Forward Error Correction ..... 197, 305
- Reduced traffic class ..... 184, 335
- reduction subnets
  - configuring ..... 131, 350
  - filtering source/destination ..... 101
- reduction tradeoff for speed ..... 350
- reduction tunnels
  - dynamic resource allocation ..... 348
  - enabling endpoints ..... 129, 345
  - fast ..... 348
  - heartbeat packets ..... 347
  - meta packets ..... 349
  - MSR symbol size ..... 349
  - source MAC addresses ..... 350
  - statistics ..... 343, 378
  - tunnel switching ..... 347
- registration servers
  - configuring ..... 75, 352
  - configuring in Quick Setup ..... 46
- remote circuit speeds
  - and router overhead ..... 156
  - configuring ..... 165, 178, 339
- remote routes, viewing ..... 138, 355



- reports
  - about.....227
  - acceleration.....248
  - Active Flow Pipelining.....248
  - Application Summary
    - all traffic.....240
    - WAN traffic.....230
  - CIFS and Exchange acceleration.....253
  - Data Reduction.....237
  - Endpoints Summary.....258
  - Executive.....260
  - Fast Connection Setup.....250
  - Forward Error Correction.....252
  - HTTP acceleration.....254
  - Inbound Bandwidth.....246
  - Outbound Bandwidth.....244
  - Packet Size Distribution.....243
  - Passthrough Data.....242
  - throughput
    - all traffic.....235
    - WAN traffic.....228
  - Traffic.....256
  - WAN/Multi-Path performance.....231
- requirements
  - browser and JVM versions.....45
  - SSH version.....284
- retransmissions, RADIUS.....83, 342
- retries, SSH login.....81, 302
- RIP
  - for dynamic routes.....69, 356
  - for packet interception.....105, 331
- rolling back to a previous release.....267, 296
- route injection.....105, 331
- router balancing
  - route-based.....73, 357
  - router-based using ToS.....358
- router configuration
  - for packet interception.....108
  - Multi-Path.....122
- routes
  - adding static.....68, 356
  - configuring local.....66
  - from OSPF.....69, 330
  - from RIP.....69, 356
  - importing from a file.....71, 291
  - LAN-WAN check.....348
  - polling from a router.....70, 359
  - remote, viewing.....138, 355
- RTT
  - and meta-packet wait times.....349
  - reported by ping.....271, 294
  - reported by traceroute.....272, 300
- S**
  - Safe Mode.....270, 295
  - safety and EMC certifications.....417
  - sample topologies.....22
  - secondary boot image.....295
  - secondary IP address, Multi-Path.....116, 327
  - secondary registration server.....75, 352
  - secret key, RADIUS.....83, 342
  - Secure Shell (SSH), supported version.....284
  - secure wipe.....269, 293
  - security associations.....321, 374
  - security features.....79
    - changing the packet capture password.....88, 360
    - controlling operator access.....86, 361
    - defining local users.....84, 300
    - defining RADIUS servers and server groups.....82
    - securing front panel access.....87, 360
    - selecting authentication methods.....80, 300
  - serial port
    - baud rate, setting.....315
    - default settings.....41
  - Server/Client Summary.....281
  - servers
    - DNS.....59
    - NetFlow.....274, 365
    - NTP.....62, 362
    - RADIUS.....82, 343
    - registration.....75, 352
    - Syslog.....65, 364
    - WX 100
      - connecting client devices.....43
      - disconnecting client devices.....44
      - installing.....31, 37
    - Server/Client Summary.....281
  - Setup Wizard
    - IPSec.....215
    - outbound QoS.....164
  - SMB signing, disabling.....206
  - SNMP
    - configuring.....64, 361
    - list of traps.....397
  - SNTP, configuring.....62, 362
  - software passthrough.....37, 43
  - source address in RADIUS packets.....343
  - source MAC address, changing.....350
  - source/destination subnets.....101
  - special characters.....57
  - specifications, device.....385
  - Spoke topology.....344
  - SSH interface
    - downloading SSH applications.....284
    - enabling and disabling.....361
    - standalone mode, WX 100.....364
    - static routes, adding.....68, 356

statistics	
acceleration .....	248
Active Flow Pipelining .....	248
application	
all traffic .....	240
WAN traffic .....	230
data reduction .....	237
executive summary .....	260
exporting .....	279
Fast Connection Setup .....	250
Forward Error Correction .....	252
inbound bandwidth .....	246
interface .....	318, 373
interpreting CSV .....	405
outbound bandwidth .....	244
packet size distribution .....	243
passthrough traffic .....	242
throughput	
all traffic .....	235
WAN traffic .....	228
traffic .....	256
WAN/Multi-Path performance .....	231
straight-through cable .....	32, 38
subnet mask, configuring	
in CMS .....	30
in front panel .....	35
in Web/CLI .....	57, 320
subnets	
advertising for reduction .....	131, 350
defining whether encryption is required... ..	218, 224, 323
discovering .....	66
excluding from default assemblers .....	143, 346
excluding from outbound QoS .....	180, 341
excluding from reduction .....	101, 316
filtering the Traffic report .....	257
unadvertised subnets and outbound QoS .....	162
support	
browser and JVM .....	45
generating diagnostic files .....	280, 299
SSH version .....	284
technical .....	18
switch-to-wire .....	21
symbol size, MSR .....	349
Syslog	
configuring .....	65, 364
list of messages .....	397
system log file .....	277, 375
system software, upgrading .....	267, 299
technical support .....	18
templates	
IPSec, defining .....	222, 321
Multi-Path, defining .....	118, 328
outbound QoS	
defining .....	174, 338
names of .....	173, 337
terminal emulation program .....	41, 284
thresholds, loss and latency	
Multi-Path .....	121, 329
WAN performance monitoring .....	124, 366
throughput statistics	
all traffic .....	235
WAN traffic .....	228
time settings	
manual .....	62, 315
NTP server .....	62, 362
timeout	
idle user .....	85, 300
RADIUS server .....	83, 342
registration server .....	354
topology	
sample .....	22
setting .....	99, 344
ToS/DSCP values	
defining by QoS traffic class .....	181, 341
in application definitions .....	94, 311
in Multi-Path configurations .....	117, 327
in UDP heartbeat packets .....	348
processing queues by .....	184
traceroute utility .....	272, 300
tradeoff, reduction for speed .....	350
traffic classes	
inbound .....	184, 335
outbound .....	167, 173, 338
outbound QoS and Multi-Path .....	95
traffic flows	
exporting to CSV file .....	256, 365
sending to NetFlow server .....	274, 365
viewing details of one flow .....	372
viewing top flows .....	256
Traffic report .....	256
traps	
configuring .....	64
list of .....	397
tunnel mode .....	145, 347
tunnel switching	
about .....	148
enabling .....	347
tunnels, reduction .....	129, 345
types of applications .....	93, 309
<b>T</b>	
TCP	
required ports .....	30
traffic class .....	184, 335

**U**

## UDP

- and heartbeat packets ..... 308, 347
- and meta packets ..... 145, 347
- required ports ..... 30
- unadvertised subnets and outbound QoS ..... 162
- undefined applications, defining
  - manually ..... 92, 309
  - using the Traffic report ..... 257
- upgrading the SRS software ..... 267, 299
- URLs, in application definitions ..... 94, 311
- user names and passwords
  - default ..... 41
  - defining ..... 84, 300

**V**

- validating remote routes ..... 139, 355
- virtual endpoints
  - for outbound QoS ..... 179, 340
  - in Profile Mode ..... 334
  - monitoring ..... 228, 230
- VLAN 802.1q support ..... 61, 319

**W**

- WAN circuit speeds
  - and router overhead ..... 156
  - configuring ..... 165, 178, 339
  - congestion control ..... 166, 180, 339
- WAN performance monitoring
  - configuring ..... 123, 366
  - viewing reports ..... 231
- WAN reduction subnet
  - CLI option ..... 351
  - for off-path devices ..... 133
- WAN statistics ..... 228
- WCCP for packet interception ..... 106, 331
- Web console
  - about ..... 56
  - allowing access ..... 86
  - allowing access by address ..... 361
  - browser and JVM support ..... 45
  - enabling and disabling ..... 360
  - logging in ..... 55
  - monitor settings ..... 45
- Weighted Fair Queuing ..... 183, 337
- Weighted Strict Priority ..... 184, 337
- wiping the disks ..... 269, 293
- Wizard
  - IPSec ..... 215
  - outbound QoS ..... 164
- WX 100 clients
  - and multiple tunnels ..... 363
  - client-mode command ..... 363
  - connecting ..... 43

- disconnecting ..... 44
- Server/Client Summary ..... 281
- WX 100 installation ..... 31, 37
- WX 100 server, standalone mode ..... 364
- WXC devices
  - enabling NSC ..... 134, 346
  - enabling NSC for applications ..... 135, 346
  - file pre-synchronization ..... 146, 346
  - rebooting ..... 270, 295
  - wiping the disks ..... 269, 293



# Copyrights

## Traceroute Copyright License

---

Copyright (c) 1990, 1993

The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Van Jacobson. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## OpenSSL Copyright License

---

Copyright (c) 1998-2000 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

“This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)”

4. The names “OpenSSL Toolkit” and “OpenSSL Project” must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [openssl-core@openssl.org](mailto:openssl-core@openssl.org).
5. Products derived from this software may not be called “OpenSSL” nor may “OpenSSL” appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

“This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)”

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Original SSLeay License

Copyright (C) 1995-1998 Eric Young ([ey@cryptsoft.com](mailto:ey@cryptsoft.com))

All rights reserved.

This package is an SSL implementation written by Eric Young ([ey@cryptsoft.com](mailto:ey@cryptsoft.com)).

The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used.

This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

All advertising materials mentioning features or use of this software must display the following acknowledgement:

“This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)”

The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

“This product includes software written by Tim Hudson (tjh@cryptsoft.com)”

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

## GNU GENERAL PUBLIC LICENSE

---

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies

of this license document, but changing it is not allowed.

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.



## TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

**0.** This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The “Program”, below, refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification”.) Each licensee is addressed as “you”.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

**1.** You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

**2.** You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

**3.** You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

**4.** You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

**5.** You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

**6.** Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

**7.** If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

**8.** If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

**9.** The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

**10.** If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

**11.** BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

**12.** IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### END OF TERMS AND CONDITIONS

#### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

one line to give the program's name and an idea of what it does.

Copyright (C) yyyy name of author

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA. Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author

Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program  
`Gnomovision' (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

## KAME Copyright License

---

This product contains a modified version of the IPsec software developed by the KAME Project.

Copyright (C) 1995, 1996, 1997, and 1998 WIDE Project.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PROJECT OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.