



**WXOS Software for
Application Acceleration Platforms**

WX/WXC Operator's Guide

*Release 5.5
June 2008*

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
USA
408-745-2000
www.juniper.net

Part Number: 530-017951-01

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 487.

This product includes a modified version of OpenSSL. © 2001-2007 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 488.

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 488.

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 496.

This installation includes a modified version of ospfd. © 2001-2007 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.: 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, TCP Acceleration, 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 © 2008, Juniper Networks, Inc. All rights reserved. Printed in USA.

WX/WXC Operator's Guide, Release 5.5

Revision History
November 2007 —Rev 1
June 2008—Rev 2

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.

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

	About This Guide	15
	Audience	15
	Package Contents	15
	Operator's Guide Contents	16
	Document Conventions	17
	Commonly Used Terms	17
	Typographical Conventions	17
	Technical Support	18
	Obtaining Additional Product Information	18
Chapter 1	Introduction	19
	About the WX and WXC Devices	19
	Features and Benefits	20
	What's New in Version 5.5	21
	Sample Topologies	22
	Typical Inline Deployment	22
	Off-Path Deployment	23
	Point-to-Multipoint Topology	23
	Virtual Private Network (VPN) Topology	24
	Basic Concepts	24
	Communities and Registration Servers	25
	Service Tunnels	25
	Local Routes and Compression Subnets	26
	Remote Routes	26
	Community Topologies	27
	High Availability Support	27
	Demo Mode	27
	WX Central Management System (CMS)	28
	Where to Go Next	28
Chapter 2	Installation	29
	Before You Begin	29
	Battery Warning	30
	Manual and Automatic Installations	30
	Inline and Off-Path Installations	31
	Interface Speeds and Modes	31
	Installing the WX 15, WX 20, and WXC 250	32
	Hardware Installation	32
	Configuring Network Settings	33
	Installing the WX 60 and WXC 500	36
	Hardware Installation	36
	Configuring Network Settings	38
	Installing the WXC 590	40

Hardware Installation	40
Disconnecting Power from the WXC 590	43
Replacing the Disk Drives.....	43
Configuring Network Settings	44
Installing the WX 100	46
Hardware Installation	46
Copper-Wire Interfaces	46
Fiber-Optic Interfaces.....	48
Disconnecting Power from the WX 100.....	49
Configuring Network Settings	49
Connecting Client Devices to the Server	52
Distributing Tunnels Across Client Devices	53
Disconnecting Client Devices from the Server	54
Running Quick Setup through the Web Console.....	54
Post-Installation Tasks.....	59
Where to Go Next	59

Chapter 3 Configuring Basic Setup Policies 61

Using the Web Console	61
Logging In.....	61
Understanding the WXOS Web Console Interface	62
Using Special Characters.....	63
Configuring Basic Setup Policies.....	63
Configuring Device Address and Contact Information	63
Configuring the Interface Settings.....	65
Configuring 802.1Q VLAN Support	67
Configuring Time Settings	68
Obtaining a Permanent License.....	69
Enabling SNMP	71
Enabling Syslog Reporting	72
Configuring Local Routes	73
Adding Static Routes	75
Enabling RIP and OSPF Support.....	76
Enabling Route Polling.....	77
Importing a Routing Table.....	78
Enabling Route-Based Router Balancing.....	80
Configuring Registration Servers and Communities.....	82
Defining Registration Servers and Passwords.....	82
Defining Communities	84
Configuring AAA	86
Selecting Authentication Methods.....	87
Enabling Authorization Checking.....	88
Defining RADIUS Servers and Server Groups	89
Defining TACACS+ Servers	91
Defining Local Users.....	92
Securing Operator Access.....	94
Securing Front Panel Access.....	95
Managing Applications.....	95
About Application Definitions.....	96
Viewing the Application Overview	98
Configuring Application Definitions	99
Testing New Application Definitions.....	103
Assigning Applications to Traffic Classes	104
Monitoring Applications.....	105

Chapter 4	Configuring Advanced Setup Policies	107
	Configuring Topology Settings.....	108
	Selecting a Topology.....	108
	Partial Mesh Example	108
	Tiered Network Example.....	109
	Selecting the Community Size	110
	Using Source/Destination Filters.....	112
	Configuring the ARP Table	114
	Defining the Prime Time	115
	Configuring Packet Interception	116
	Methods of Packet Interception	116
	Route Injection.....	116
	WCCP	117
	External	118
	Configuring Packet Interception for Off-Path Devices	118
	RIP Router/Switch Configuration Commands.....	120
	Single Layer 3 Switch	120
	Dual Off-Path Devices on Two Layer 3 Switches	122
	WCCP Router Configuration Commands	123
	Unicast Example	124
	Multicast Example for a Cisco Branch Router	125
	Multicast Example for the Catalyst 6509.....	126
	External Policy-Based Router Commands.....	127
	Alternatives to Packet Interception	127
	Layer 2 Switch Sandwich	127
	Layer 3 Switch Sandwich	128
	Configuring Policy-Based Multi-Path.....	129
	Procedure for Configuring Multi-Path.....	130
	Enabling Multi-Path and Defining Marking Methods	131
	Defining Multi-Path Templates.....	133
	Defining Multi-Path Endpoints	135
	Configuring Routers to Support Multi-Path.....	137
	Configuring WAN Performance Monitoring.....	138
	Configuring Events.....	140
	Configuring Multiple Tunnels Between WX 100 Servers.....	144
Chapter 5	Configuring Compression Policies	145
	Configuring Basic Compression Policies.....	145
	Configuring Endpoints for Compression	145
	Advertising Compression Subnets	148
	Configuring Network Sequence Caching.....	150
	Compressing Traffic by Application	151
	Configuring Advanced Compression Policies.....	153
	Viewing and Fetching Remote Routes	154
	Configuring Tunnel Load Balancing Policies.....	155
	Defining Default Decompressors	157
	Defining Preferred Decompressors.....	159
	Configuring Tunnel Mode Settings.....	160
	Configuring Pre-Synchronization for Network Sequence Caching	161
	Configuring Tunnel Switching	163
	Tunnel Switching Between Communities	163
	Procedure for Tunnel Switching Between Communities.....	165
	Tunnel Switching Between Hub and Spoke Devices	165

Chapter 6	Applying Quality of Service (QoS) Policies	167
	Using Outbound QoS to Enhance Performance.....	167
	Understanding Outbound Bandwidth Management	168
	Traffic Classes and Bandwidths	169
	QoS Templates and Endpoints.....	170
	WAN Circuit Speeds and Router Overhead	170
	Dedicated and Oversubscribed WANs	171
	Bandwidth Detection	172
	Direct Setup Versus Wizard Configuration Results	174
	Class Priorities and Excess Bandwidth Allocation	176
	ToS/DSCP Values	177
	Unadvertised Subnets	177
	Configuring Outbound QoS Policies	177
	Procedure for Configuring Outbound QoS Policies.....	178
	Using the Outbound QoS Setup Wizard	179
	Defining Outbound QoS Settings by Endpoint	186
	Defining Traffic Classes	188
	Defining Outbound QoS Templates	189
	Defining Outbound QoS Endpoints.....	191
	Changing Outbound ToS/DSCP Values.....	196
	Starting and Stopping Outbound QoS	198
	Processing Queues Based on Incoming ToS/DSCP Values	199
	Configuring Inbound QoS Policies.....	199
	Summary of Key Terms	202
Chapter 7	Accelerating WAN Traffic	203
	Packet Flow Acceleration	203
	Overview of Packet Flow Acceleration.....	203
	TCP Acceleration.....	204
	Forward Error Correction	206
	Fast Connection Setup	206
	Requirements for Using Packet Flow Acceleration.....	207
	Enabling Packet Flow Acceleration by Endpoint	208
	Enabling TCP Acceleration by Application	212
	Enabling Fast Connection Setup by Application.....	213
	Application Flow Acceleration.....	214
	Overview of Application Flow Acceleration	214
	Microsoft CIFS and Microsoft Exchange Acceleration	215
	HTTP Acceleration	216
	Enabling Microsoft CIFS Acceleration	218
	Enabling Microsoft Exchange Acceleration	221
	Enabling HTTP Acceleration	223
Chapter 8	Configuring IP Security (IPSec) and SSL Optimization	227
	Configuring IP Security (IPSec)	227
	Overview of IPSec	227
	Default IPSec Policy	228
	IPSec Implementation Details	228
	Procedure for Configuring IPSec Policies	229
	Using the IPSec Setup Wizard.....	229
	Defining IPSec Settings by Endpoint.....	234
	Defining IPSec Templates.....	236
	Defining the Default IPSec Policy.....	238

Defining the IPSec Application Filter	239
Optimizing SSL Traffic	240
Overview of SSL Optimization	240
Importing SSL Certificates	241
Enabling Applications for SSL Optimization	243
Chapter 9 Monitoring and Reporting	245
Viewing and Printing Reports	245
WAN Statistics	246
WAN Throughput Statistics	246
WAN Application Summary	248
WAN Performance Statistics	249
Compression Statistics	253
Device Throughput Statistics	253
Data Compression Statistics	255
Application Summary Statistics	258
Passthrough Statistics	260
Packet Size Distribution Statistics	261
Outbound Bandwidth Statistics	262
Inbound Bandwidth Statistics	264
Acceleration Statistics	266
TCP Acceleration Statistics	266
Fast Connection Setup Statistics	268
Forward Error Correction Statistics	269
CIFS and Exchange Acceleration Statistics	271
HTTP Acceleration Statistics	272
Traffic Statistics	274
Endpoints Summary	277
Executive Summary	278
Events Summary	280
Chapter 10 Maintaining WX Devices	283
Maintaining Configurations and Software	283
Saving the Device Configuration	283
Displaying the Running Configuration	285
Loading a Device Configuration File	286
Loading a Boot Image	287
Clearing Application Monitoring Statistics	288
Setting the Device to the Factory Default Configuration	288
Rebooting the Device	290
Using Maintenance Tools	291
Pinging a Network Device	291
Running a Traceroute to a Network Device	292
Running a Packet Capture	293
Generating NetFlow Records	294
Entering CLI Commands from the Web Console	295
Viewing and Saving System Logs	296
Viewing and Saving the Access Control Log	297
Exporting Performance Data	298
Creating a Diagnostic File	299
Viewing Flow Diagnostics	300
Viewing the WX 100 Server/Client Summary	303

Chapter 11	Using the Command Line Interface (CLI)	305
Accessing the CLI		305
Using a Secure Shell Program from a Remote Workstation		306
Using a Terminal Connected to the Serial Port.....		306
Logging In Using the CLI		307
CLI Basics.....		307
Command Modes.....		308
CLI Command Summary.....		309
System-Level Commands		313
activate		313
commit		313
configure		313
copy		314
embed		314
flow-reset		314
import-route-table		315
list		315
load-config		316
packet-capture		317
ping		318
reboot		319
remove		320
reset		320
rollback		320
save-config		321
set		322
shutdown		322
source		323
support		323
upgrade		323
traceroute		324
Configuration Commands		324
configure aaa		324
configure acceleration		326
configure application		332
configure arp		335
configure backup		336
configure clock		338
configure console		339
configure dns		339
configure event		340
configure filter		342
configure flow-reset		343
configure interface		344
configure ip		345
configure ipsec		346
configure license		349
configure log		350
configure mon-apps		351
configure multi-path		352
configure ospf		356
configure packet-interception		357
configure path-mtu-discovery		361
configure prime-time		361

configure profile-mode	362
configure qos inbound	363
configure qos outbound	365
configure radius	370
configure reduction	371
configure reduction-subnet	378
configure reg-server	380
configure remote-routes	382
configure rip	383
configure route	384
configure route-poll	387
configure security	388
configure snmp	389
configure sntp	390
configure ssl certificate	391
configure ssl optimization	392
configure stack-group	392
configure syslog	394
configure system	394
configure tacplus	395
configure top-talker	396
configure wan-performance-monitor	397
Show Commands.....	399
show aaa	399
show acceleration.....	399
show access-log.....	400
show all	400
show application	401
show arp	401
show backup-sr	401
show clock.....	401
show connection	402
show console.....	402
show contact	402
show dns	402
show event	403
show filter	403
show flow-details	404
show flow-reset	405
show import-route-table	405
show interface.....	405
show ip.....	406
show ipsec.....	406
show license	407
show location	407
show log.....	407
show mon-apps	407
show multi-path.....	408
show ospf	408
show packet-capture.....	408
show packet-interception.....	409
show path-mtu-discovery	410
show prime-time	410
show profile-mode.....	410

	show qos excl-filter	411
	show qos inbound	411
	show qos outbound	411
	show radius	411
	show reduction	412
	show reduction-subnet	413
	show reg-detail	413
	show reg-server	414
	show reg-summary	414
	show remote-routes	414
	show rip	415
	show route	415
	show route-poll	415
	show security	416
	show snmp	416
	show snmp	416
	show ssl certificate	416
	show ssl optimization	416
	show stack-group	417
	show syslog	417
	show system	418
	show system-name	418
	show tacplus	418
	show top-talker	419
	show uptime	419
	show version	419
	show wan-performance-mon	419
Appendix A	WX Device Specifications	421
	General Specifications—All Platforms	422
	WX Family Specifications	424
	WXC Family Specifications	425
	DB9 Console Port Pin-Outs	426
Appendix B	SNMP Traps and Syslog Messages	427
	Severity Levels	427
	System Events and SNMP Traps	428
	Syslog Messages	432
Appendix C	Understanding Exported Data Results	437
	NetFlow Version 5 Export	437
	Performance Statistics Export	438
	General Device Information	439
	Data Section Information	439
	System Session Statistics	440
	Compression Session Statistics	442
	Application Session Statistics	442
	WAN Statistics	443
	Application Flow Acceleration Statistics	443
	Bandwidth Management Statistics	444
	WAN Performance Statistics	445
	Inbound Traffic By Port Statistics	445
	Top Traffic Export	446

	Flow Diagnostics Export.....	447
Appendix D	Common Application Port Numbers	451
Appendix E	Demo Mode	453
	About Demo Mode.....	453
	Purpose and Benefits.....	453
	Sample Topology.....	454
	Security.....	454
	Return on Investment.....	454
	Pre-Installation Tasks.....	455
	Installing a WX 15, WX 20, or WXC 250 in Demo Mode.....	456
	Hardware Installation.....	456
	Configuring Network Settings.....	457
	Installing Other Platforms in Demo Mode.....	459
	Hardware Installation.....	459
	Configuring Network Settings.....	460
	Configuring Demo Mode through the Web Console.....	461
	Running Quick Setup.....	461
	Defining Virtual Devices in Demo Mode.....	461
	Excluding Traffic to the Local Subnet.....	463
	Viewing Performance Reports.....	464
	Exporting Performance Data.....	466
	Converting from Demo Mode to Active Mode.....	466
Appendix F	Safety and EMC Certifications	467
	Product Reclamation and Recycling Program.....	468
Appendix G	Safety Recommendations and Warnings	469
	Power Cable Warning (Japanese).....	469
	VCCI Compliance.....	469
	Lightning Activity Warning.....	470
	Jewelry Removal Warning.....	470
	Installation Warning.....	470
	IT Power Statement.....	470
	SELV Circuit Warning.....	470
	Circuit Breaker (15A) Warning.....	470
	Grounded Equipment Warning.....	471
	Class 1 Laser Product Warning.....	471
	Laser Beam Warning.....	471
	Battery Warning.....	471
	Rack Mounting of Systems.....	472
	Anti-Static Precautions.....	472
	Glossary	473
	Index	477
	Copyrights	487
	Traceroute Copyright License.....	487
	OpenSSL Copyright License.....	488
	GNU GENERAL PUBLIC LICENSE.....	490

KAME Copyright License496

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.

This section describes 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 and WXC devices are shipped with the following:

- 1 WX 15, WX 20, WX 60, WX 100, WXC 250, WXC 500, or WXC 590
- 1 female/female DB-9 crossover cable (WX 15, WX 20, WX 100, WXC 250, and WXC 590)
- 2 rack-mount flanges for rack mount installation (already assembled on WXC 590)
- 6 screws for the rack-mount flanges (WX 60, and WX 100)
- 4 screws for the rack-mount flanges (WX 15, WX 20, and WXC 250)
- 4 rubber feet for desktop placement (already assembled on WX 15)
- 1 power Cord (two for WX 100 and WXC 590)
- 1 quick start card
- 1 documentation/utilities CD
- 1 release notes document



CAUTION: Special packaging material is provided to protect the WXC systems during shipping. Retain the packing material in case the unit needs to be shipped again for any reason. Shipping the unit without the original packaging material will void the warranty.

Operator's Guide Contents

- Chapter 1, “Introduction” on page 19, introduces the WX and WXC devices, describes the new features, and provides sample topologies for deployment.
- Chapter 2, “Installation” on page 29, describes how to install and initially configure WX and WXC devices.
- Chapter 3, “Configuring Basic Setup Policies” on page 61, 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 107, describes how to configure advanced policies, such as topology parameters, packet interception, and Policy-Based Multi-Path (PBM).
- Chapter 5, “Configuring Compression Policies” on page 145, describes how to configure policy settings for data compression, and the communication links with other devices in the community.
- Chapter 6, “Applying Quality of Service (QoS) Policies” on page 167, 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 203, describes how to configure Packet Flow Acceleration for TCP applications, and Application Flow Acceleration for CIFS, Exchange, and HTTP traffic.
- Chapter 8, “Configuring IP Security (IPSec) and SSL Optimization” on page 227, describes how to configure IPSec to encrypt the traffic between two WX and WXC devices.
- Chapter 9, “Monitoring and Reporting” on page 245, describes the detailed graphs and reports that you use to monitor network performance.
- Chapter 10, “Maintaining WX Devices” on page 283, 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 305, 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 421, lists the specifications for each type of WX and WXC device.

- Appendix B, “SNMP Traps and Syslog Messages” on page 427, describes SNMP trap and syslog messages generated by the WX and WXC devices.
- Appendix C, “Understanding Exported Data Results” on page 437, describes the details of exported data results. After exporting the compression statistics to a comma-separated values file, use this appendix to interpret the data.
- Appendix D, “Common Application Port Numbers” on page 451, provides a listing of common application port numbers that you can use when defining new applications.
- Appendix E, “Demo Mode” on page 453, describes how to configure Demo Mode to test the performance of WX and WXC devices without affecting the network traffic.
- Appendix F, “Safety and EMC Certifications” on page 467, lists the safety and EMC certifications for each type of WX and WXC device.
- Appendix H, “Safety Recommendations and Warnings” on page 469, lists hardware safety recommendations and warnings.
- “Glossary” on page 473, 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 for which you must substitute another value.	<code>set ip < IP address ></code>
Square brackets	Encloses optional parameters.	<code>show log [< n >]</code>
Curved brackets	Encloses related parameters.	<code>set mode {on off}</code>
Italics	Names of manuals, directories, or files.	For more information about WX CMS, refer to the <i>WX Central Management System (CMS) Administrator's Guide</i> .

Technical Support

For technical support, use the following methods:

- Go to <http://www.juniper.net/support> and open a case online using the Case Manager.
- Send email to support@juniper.net.
- To call from the United States, Canada, or Mexico, dial + 1-888-314-JTAC. To call from other locations, check the list of local support centers at http://www.juniper.net/support/support_contacts.html or dial + 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 5.5” on page 21
- “Sample Topologies” on page 22
- “Basic Concepts” on page 24
- “WX 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 decompresses (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 service 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.** While MSR effectively increases available bandwidth, Packet Flow Acceleration provides several methods to improve TCP application performance in networks where the use of available bandwidth is constrained by network latency.
- **Application Flow Acceleration.** 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 for specific applications and network 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 421.

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 email, 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), and 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 5.5

WXOS 5.5 includes the following new features:

- **SSL optimization.** WXOS 5.5 can selectively optimize application traffic that uses the Secure Socket Layer (SSL) for encryption. SSL traffic is decrypted, optimized with compression, acceleration, and QoS, and then re-encrypted with IPsec before transmission. On the receiving WX, the traffic, is decrypted, re-encrypted using SSL, and then forwarded to its destination (refer to “Optimizing SSL Traffic” on page 240).
- **IPsec enhancements.** IPsec encryption can now be applied to selected applications, rather than to all applications. For each application, you can specify whether IPsec is required, used if available, or never used (refer to “Defining the IPsec Application Filter” on page 239).
- **SMB Signing.** CIFS traffic flows to servers that require Server Message Block (SMB) signing can now be accelerated by configuring the WX device to log in to the Windows server and create the signature applied to the traffic flow (refer to “Enabling Microsoft CIFS Acceleration” on page 218).
- **Packet capture filtering.** Packet captures can now be limited to a specific interface, IP protocol, source and destination address and port, and specific TCP flags. You can start a packet capture from the Flow Diagnostics page by clicking a traffic flow to populate the filter. Also, the packet capture capability can be enabled by user account (refer to “Running a Packet Capture” on page 293).

- **TACACS + support.** WXOS users can now be authenticated and authorized remotely by a TACACS + server (refer to “Defining TACACS + Servers” on page 91).
- **SSH version requirement option.** For better security, SSHv2 can be required over SSHv1 for device management (refer to “configure security” on page 388).
- **New daylight savings time support.** The new start and end dates for Daylight Savings Time in the U.S. and Canada starting in 2007 are supported in WXOS 5.5.
- **New tunnel mode default.** The default tunnel mode has changed from IPComp to UDP.

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 23
- “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). When two or more devices are installed in the same community, a service tunnel is formed between them.

Figure 1: Typical Inline Deployment




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. Service tunnels, which are automatically established and managed by the WX and WXC devices at the various corporate sites, continuously process and compress the data traveling through these tunnels thereby reducing traffic on the WAN circuits and creating more bandwidth.

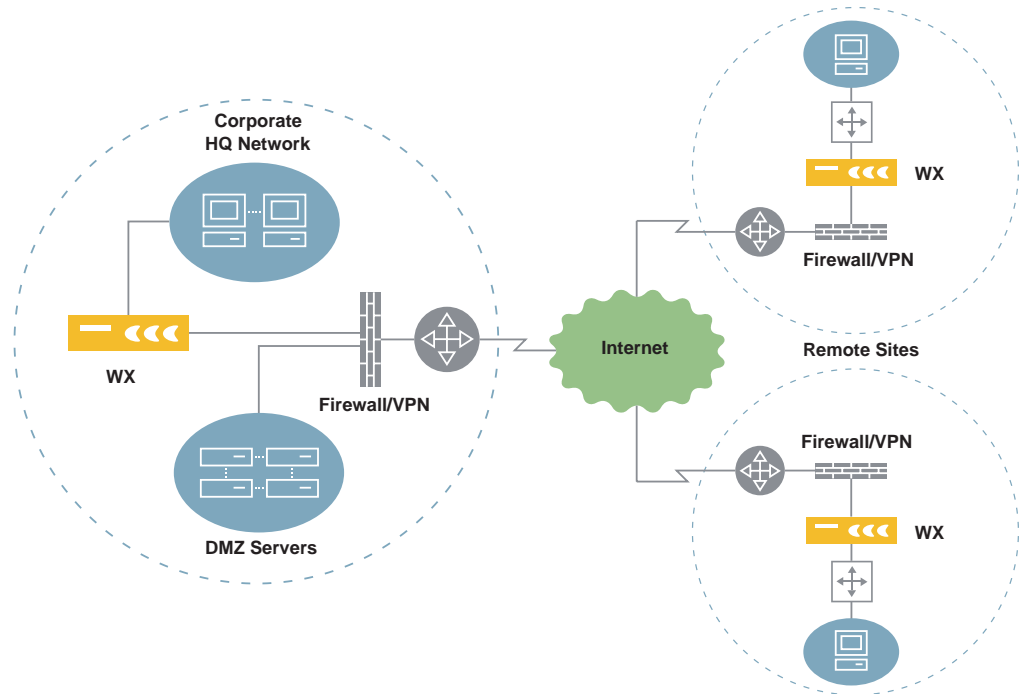
Note that it is not necessary to deploy a WX or WXC device for every remote site that links into the corporate headquarters network. In such instances, data from the hub is passed through without compression.

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 compressing 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
- “Service Tunnels” on page 25
- “Local Routes and Compression 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 compression (one on each side of the WAN). Two or more devices that can compress and decompress data for each other are said to be in the same community. You can selectively enable or disable data compression 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 compression occurs only between devices in the same community, in large deployments you can limit the number of devices in each community. To send compressed 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.

Service Tunnels

When you install a new WX or WXC device and specify a registration server, the device attempts to form a service 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 compressed data to remote devices, and IN tunnels that convey compressed data to be decompressed. At any time, you can disable data compression from all other devices and/or compress data only for specific devices in the community.

Local Routes and Compression Subnets

Local routes are the routes defined in the device's routing table. When you 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

Compression subnets are the LAN subnets for which the local device can decompress the data compressed by other devices. Static routes and routes discovered dynamically on the Local interface are added to the list of compression 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 compression subnet option through the CLI so that routes discovered on the Remote interface are included on the list of compression subnets.



NOTE: For off-path devices, where only the Local interface is connected to the network, all routes are listed as compression 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 compression subnets advertised by the other WX and WXC devices in the community. Each device can compress 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 one of the following community topologies. The community topology setting ensures that each device's resources are allocated efficiently.

- **Mesh.** Multiple devices are interconnected and each one can compress and decompress data for all the others.
- **Hub and spoke.** A central device (Hub) can compress and decompress data for all other devices in the community. By default, the spoke devices compress data only for the hub. A community can have multiple hubs. Each device attempts to form a service tunnel with a hub before creating tunnels to other devices.
- **Point-to-point.** Two devices of the same type communicating exclusively with each other. Typically used to connect two data centers.

For hub and mesh devices, you can specify the maximum number of devices so that sufficient resources are allocated for the potential number of service 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.

You can also disable 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.

Demo Mode

Demo Mode lets you see how a WX or WXC device performs in your network without affecting network traffic (supported on all devices except the WX 15). In Demo Mode, the device passively calculates potential data compression 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 compression 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 245).

In Demo 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 Demo Mode, refer to “Demo Mode” on page 453.

WX Central Management System (CMS)

The WX 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 selectively apply configuration changes and software upgrades. 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 61 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 31
- “Running Quick Setup through the Web Console” on page 54
- “Post-Installation Tasks” on page 59

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 15A or 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 a power cord (two for the WX 100 and WXC 590) to the back of the chassis, 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. For a WX 15 and WXC 590, provide three inches of clearance on both sides of the chassis to allow cooling air to be drawn through the side panels. 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.
- Ensure that paper materials or heavy equipment are not stacked on top of a device.
- For rack-mount installations, reserve space for a 1U form factor device (WX 15, WX 20, and WXC 250) or a 2U form factor device (WX 60, WX 100, WXC 500, and WXC 590).

- Identify a 10/100 Ethernet LAN port (for a WX 15, WX 20 or WXC 250) or a 10/100/1000 LAN port (for a WX 60, WX 100, WXC 500, or WXC 590) where you can connect the WX or WXC. 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 or WXC 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.
- Reserve an IP address and identify the default gateway for the WX or WXC. 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 applied to the device. For more information, refer to the CMS administrator's guide.

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 116.

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

- “Interface Speeds and Modes” on page 31
- “Installing the WX 15, WX 20, and WXC 250” on page 32
- “Installing the WX 60 and WXC 500” on page 36
- “Installing the WXC 590” on page 40
- “Installing the WX 100” on page 46

Interface Speeds and Modes

Interface speed and duplex settings should be the same across all devices: the switch, the WX 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 the fiber-optic version of the WX 100 because fiber does not support hardware bypass.

Installing the WX 15, WX 20, and WXC 250

This section describes the installation process for the WX 15, WX 20, and WXC 250.

- “Hardware Installation” on page 32
- “Configuring Network Settings” on page 33

Hardware Installation

To install the WX 15, WX 20, or WXC 250 in your network:

1. Set up the chassis.
 - To install the device in a 19-inch device rack, install the supplied brackets (front panel forward) to the sides of the device with the countersunk screws provided. Next, install the chassis in your network device rack.
 - To install the WX 15 on a desktop, place the chassis on a desktop or on top of another device so that all four rubber feet are secure on the flat surface. For a WX 20 or WXC 250, you must first install the supplied rubber feet in the marked areas on the bottom of the chassis.
2. Connect the network cables to the device.




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

The WX 15, WX 20, and WXC 250 have two 10/100 Ethernet interfaces. These RJ-45 ports are labeled REMOTE and LOCAL on the back of the chassis (Figure 6).

To connect the network cables:

- a. Locate the cable that connects the switch (or other aggregating device) to the router.
- b. Disconnect this cable from the router and connect it to the LOCAL port on the WX or WXC device.
- c. Connect a straight-through cable (not provided) from the router port to the REMOTE port on the WX or WXC device.

Figure 6: WX 15 Ethernet Ports

3. Use one of the following methods to verify connectivity across the WX or WXC when the power is off. This step ensures that the correct cables are used and that traffic will pass through the device in the event of a power loss.
 - Ping a host on the remote side of the WX or WXC from a host on the local side of the device.
 - 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 device, plug in the supplied power cord to the back of the chassis, and then connect the power cord to the local power source.



NOTE: The maximum power usage is 50 Watts Max or 170 BTU/hour for the WX 15, and 150 Watts Max or 510 BTU/hour for the WX 20 and WXC 250.

Now that the device is installed and powered on, continue to the next section to configure network settings for the device.

Configuring Network Settings

If you have the CMS management system, a full device configuration can be downloaded automatically when you apply power to the device (refer to “Manual and Automatic Installations” on page 30). To manually configure the network settings for the WX 15, WX 20, or WXC 250, 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 426.

To set IP parameters for the device using a terminal emulation program:

1. Connect an ANSI-compatible terminal to the serial port on the back of the device (Figure 7).

Figure 7: Connecting the WX 15 to an ANSI-Compatible Terminal



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 **juniper** for the password.
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.

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 4) for the speed and mode of the Local interface.
 - 0 - 10-full
 - 1 - 10-half
 - 2 - 100-full
 - 3 - 100-half
 - 4 - 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 above the Ethernet ports. Figure 8 shows the LEDs for the WX 15. Figure 9 shows the LEDs for the WX 20 and WXC 250.

Figure 8: Checking the WX 15 Link LEDs



- The link LEDs indicate the device is properly connected. If the link LEDs do not light, toggle the MDI/MDI-X buttons (WX 20 and WXC 250 only).
- The 100/10 LEDs for the REMOTE and LOCAL ports indicate 100 Mbps connectivity when the light is on.

Figure 9: Checking the WX 20 and WXC 250 Link LEDs



- If you install the WX 15 in a high-availability environment, you can press the **Bypass Disable** button to 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. By default, all traffic passes through the device during a power failure.

9. Check the LEDs on the front panel:

Front Panel LED	Model	Description
POWER	All	Indicates that power is on.
BYPASS	WX 20, WXC 250	Indicates traffic is passing through without any processing (hardware passthrough). Occurs during a reboot or system failure.
STATUS	WX 15	Indicates the device has contacted the registration server in the last 24 hours. After 24 hours of no contact with the registration server, the light is turned off, and traffic is passed through without any processing (software passthrough).
OPERATIONAL	WX 15	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	WX 15	Indicates that a power failure will block all traffic through the device (hardware passthrough disabled). To enable or disable hardware passthrough, press the Bypass Disable 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 54.

Installing the WX 60 and WXC 500

This section describes the installation process for the WX 60 and WXC 500.

- “Hardware Installation” on page 36
- “Configuring Network Settings” on page 38.

Hardware Installation

To install the WX 60 or WXC 500 in your network:

1. Set up the chassis.
 - If you plan to install the device in a 19-inch device rack, install the supplied brackets (front panel forward) to the sides of the device with the countersunk screws provided in the kit. Next, install the chassis in your network device rack.
 - To install the device on a desktop, place the chassis upside down on a smooth, flat surface. Next, install the supplied rubber feet in the marked areas on the bottom of the chassis. Finally, place the chassis on a desktop or on top of another device so that all four rubber feet are secure on the flat surface.

2. Connect the network cables.




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

The WX 60 and WXC 500 have two 10/100/1000 auto-sensing Ethernet interfaces. These RJ-45 ports are labeled REMOTE and LOCAL on the back of the chassis (Figure 10 on page 37).

To connect the network cables:

- a. Locate the cable that connects the switch (or other aggregating device) to the router.
- b. Disconnect this cable from the router and connect it to the LOCAL port on the WX or WXC device.
- c. Connect a straight-through cable (not provided) from the router port to the REMOTE port on the device.

Figure 10: Ethernet Ports



3. Use one of the following methods to verify connectivity across the device when the power is off. This step ensures that the correct cables are used and that traffic will pass through the device in the event of a power loss.
 - Ping a host on the remote side of the device from a host on the local side of the device.
 - 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 device, plug in the supplied power cord to the back of the chassis, and then connect the power cord to the local power source.



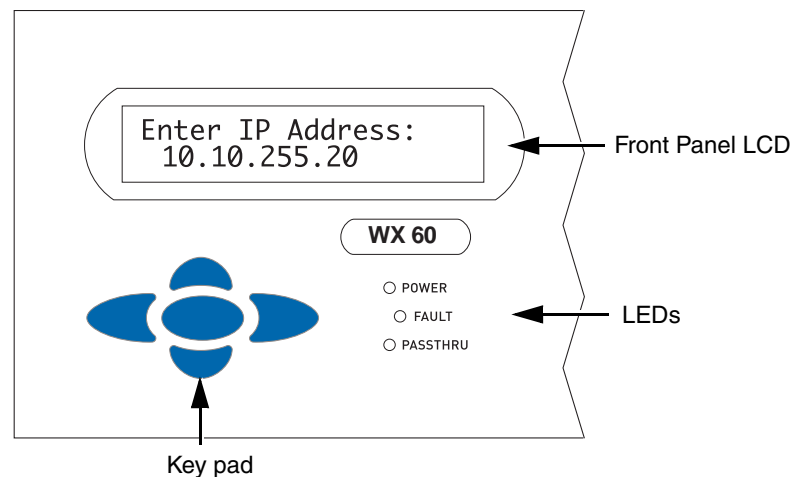
NOTE: The WX 60 and WXC 500 maximum power usage is 150 Watts Max or 510 BTU/hour.

Now that the device is installed and powered on, continue to the next section to configure network settings.

Configuring Network Settings

The configuration can be downloaded automatically from CMS when you 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 11 shows the front panel keypad and LCD of the WX 60.

Figure 11: WX 60 Front Panel Keypad and LCD



1. Press the Enter button (center button).
2. At the “Select Setup Network_” prompt in the LCD, press Enter.
3. Use the front-panel keypad to assign an IP address, the subnet mask, and the default gateway as follows:
 - 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.
 - Press Enter (the center button) after each setting.
 - After you enter the gateway address, use the left arrow to select “Save & Reboot” and press Enter.



NOTE: The default gateway is typically the next hop on the Remote side of the device. You can change this later if you designate the device as a default decompressor (refer to “Defining Default Decompressors” on page 157).

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.

To configure the interfaces from the front panel:


- a. Press Enter to display the “Select Setup Network_” prompt in the LCD.
- b. Use the down arrow to show the “Local If Settings” menu option, and press Enter.
- c. Use the left arrow to select **y**, and press Enter.
- d. Use the down arrow to show the desired speed and duplex setting, and press Enter. The options are 10/Half, 10/Full, 100/Half, 100/Full, 1000/Full, Auto-Negotiate.
- e. Use the left arrow to select **Commit&Save**, and press Enter. Repeat this process for the Remote interface.



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 12: Checking the WX 60 and WXC 500 Link LEDs



- The link LEDs indicate the device is properly connected
- The speed LEDs indicate 100 or 1000 Mbps connectivity. To verify the interface speed, use the front panel or the CLI.
- In a high-availability environment, you can press the **Bypass Disable** button to 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. By default, all traffic passes through the device during a power failure.

6. Check the LEDs on the front panel:

Front Panel LED	Description
POWER	Indicates that power is on.
PASSTHRU	Indicates traffic is passing through without any processing (hardware passthrough). Occurs during a reboot or system failure.
FAULT	Indicates a system failure (hardware passthrough).

The installation is complete. You can now run Quick Setup, as described in “Running Quick Setup through the Web Console” on page 54.

Installing the WXC 590

The following sections describe the installation process for the WXC 590.

- “Hardware Installation” on page 40
- “Disconnecting Power from the WXC 590” on page 43
- “Replacing the Disk Drives” on page 43
- “Configuring Network Settings” on page 44

Hardware Installation


To install the WXC 590 in your network:

1. Set up the chassis.
 - Install the chassis in your network device rack.
 - To install the WXC 590 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 secure on the flat surface.
2. Connect the network cables and verify connectivity.

The standard WXC 590 has two 10/100/1000 auto-sensing Ethernet interfaces (Figure 13).



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

Figure 13: WXC 590 Ethernet Ports

To connect the network cables to the WXC 590:

- 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 LOCAL port on the WXC 590.
- c. Connect a cross-over cable (not provided) from the router port to the REMOTE port on the WX 590.



CAUTION: The WXC 590 and WX 100 require a cross-over cable to allow traffic to pass through during a reboot or system failure (all other WX platforms use a straight-through cable). Always test connectivity with the power off (see the next step).

3. Use one of the following methods to verify connectivity across the WXC 590 when the power is off. This step ensures that the correct cables are used and that traffic will pass through the WXC 590 in the event of a power loss.
 - Ping a host on the remote side of the WXC 590 from a host on the local side of the WX 590.
 - 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 WXC 590, 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.



NOTE: The lights on the hard drives do not light up during operation.

The WXC 590 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.

Varoit Kojie 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 WXC 590 is installed and powered on, configure the network settings, as described in “Configuring Network Settings” on page 49.

Disconnecting Power from the WXC 590



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

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.

Replacing the Disk Drives

The hard drives on the WXC 590 should not be removed while the power is on. Doing so will set the drive status to non-operational. If a drive is removed while the system is running, enter the following CLI command to reactivate the drive:

```
reset disk status
```

You must also enter this command after replacing a failed drive. For more information about replacing the hard drives, refer to *WXC 590 Field-Replaceable Units Removal and Installation*.

Configuring Network Settings

The configuration can be downloaded automatically from CMS when you apply power to the device (refer to “Manual and Automatic Installations” on page 30). To manually configure the network settings for the WXC 590, 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 426.

1. Connect an ANSI-compatible terminal to the serial port on the front of the WXC 590 (Figure 13 on page 41).
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, 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 **juniper** for the password.
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 WXC 590. You may want to change the default gateway if you designate the device as a Default Decompressor. After installing the WXC 590, refer to “Defining Default Decompressors” on page 157 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.
 - 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 14).

Figure 14: Checking the Link LEDs on the WXC 590



- 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 WXC 590 in a high-availability environment, you can disable hardware passthrough (refer to the **embed** CLI command on page 314), 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 54.

Installing the WX 100

The WX 100 can be used as a standalone device or as a “stack” 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, WXC 500s, or WXC 590s (all clients must be the same device type).

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

- “Hardware Installation” on page 46
- “Disconnecting Power from the WX 100” on page 49
- “Configuring Network Settings” on page 49
- “Connecting Client Devices to the Server” on page 52
- “Distributing Tunnels Across Client Devices” on page 53
- “Disconnecting Client Devices from the Server” on page 54

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 secure on 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 15).



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

Figure 15: 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. Connect a cross-over cable (not provided) from the router port to the Remote port on the WX 100.



CAUTION: The WXC 590 and WX 100 require a cross-over cable to allow traffic to pass through during a reboot or system failure (all other WX platforms use a straight-through cable). Always test connectivity with the power off (see the next step).

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.

Varoit 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.

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 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 426.

1. Connect an ANSI-compatible terminal to the serial port on the front of the WX 100 (Figure 15 on page 47).
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 **juniper** for the password.
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 Decompressor. After installing the WX 100, refer to “Defining Default Decompressors” on page 157 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 16). Note that the fiber-optic version has one LED for each interface to indicate 1 Gbps.

Figure 16: 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 314), 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 54.



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 52). To run the WX 100 in standalone mode (no clients), do not roll back to WXOS 5.3 (WXOS 5.3 supports the WX 100 only in stack mode (one or more clients)).

Connecting Client Devices to the Server

After you install the WX 100 and run Quick Setup, you can add more processing capacity by connecting up to six WX 55s, WX 60s, WX 80s, WXC 500s, or WXC 590s as client devices (all clients must be the same device type).

The following table shows the number of WX 60, WXC 500, and WXC 590 clients that are recommended for the licensed speed of the WX 100.

WX 100 Speed	WX 60	WXC 500	WXC 590
Up to 45 Mbps	2	2	1
Up to 60 Mbps	3	3	2
Up to 80 Mbps	4	4	2
Up to 100 Mbps	5	5	2
Up to 125 Mbps	6	6	3
Up to 155 Mbps	6	6	3

To connect one or more client devices to the server:

1. If necessary, upgrade WX client devices to WXOS 5.0 or later, or WXC clients to WXOS 5.1 or later.
2. 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 WX 100 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 15 on page 47). 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 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 150).
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 compression and decompression on the server to ensure that tunnels are distributed across all clients:
 - a. Click Compression, clear the two check boxes at the top of the Endpoints page, click **Submit**.
 - b. Reselect the two check boxes, and click **Submit**.

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 stack server, the number of client devices is shown in the banner of the Web console (unless hidden by a license expiration warning). The compression Endpoints page on the WX 100 indicates which tunnels are handled by each client (refer to “Configuring Endpoints for Compression” on page 145).

Distributing Tunnels Across Client Devices

The tunnels on a WX 100 stack server might not be distributed evenly across all the clients. After you add client devices to a WX 100, use the following procedure to rebalance the tunnels assigned to each client. You may want to perform this procedure occasionally to maintain an even distribution of the tunnels.

1. Disable compression and decompression on the WX 100. Click Compression, clear the two check boxes at the top of the Endpoints page, and click **Submit**.
2. Click the browser Refresh button and verify that all tunnels are down.
3. When all the tunnels are down, enter the following CLI commands (skip the second command if the server does not have WXC clients):

```
reset stack-group tunnel-lb-info
reset stack-group tunnel-pref-info
set stack-group tunnel-lb-pkt-count off
commit
```

4. On the Endpoints page, select the Enable this device to COMPRESS check box, and click **Submit**.
 5. Verify that all tunnels are established.
 6. Select the Enable this device to DECOMPRESS check box, and click **Submit**.
 7. Enter the following commands to re-enable load balancing by packet count:

```
set stack-group tunnel-lb-pkt-count on
commit
```

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 stack 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 or WXC 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.

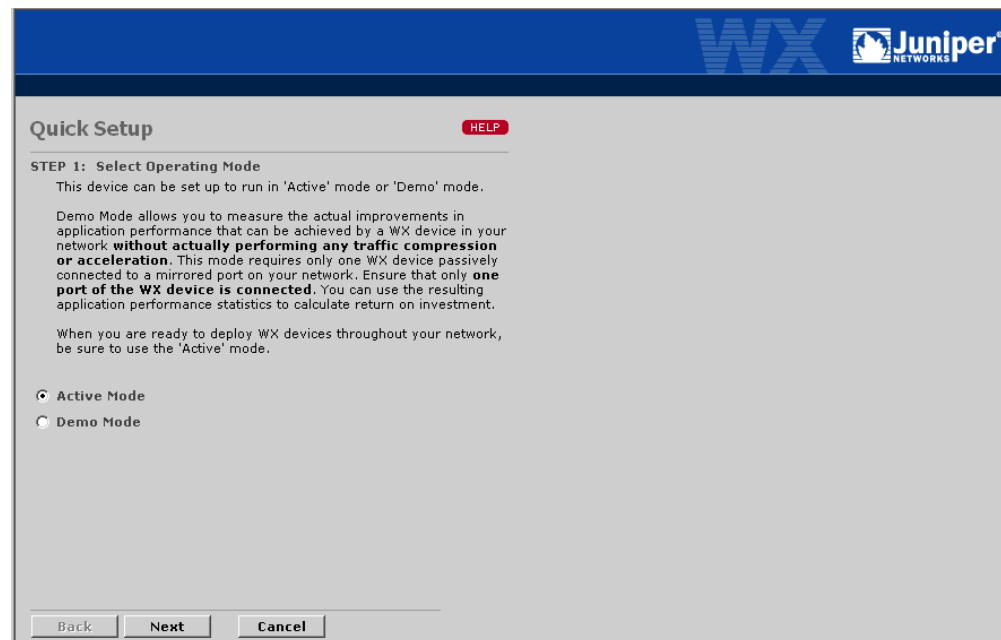
To run Quick Setup from the Web console:

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 WX device>
3. If the Security Alert dialog box appears, click **Yes** to proceed.
4. In the Login page, type **admin** for the user name and **juniper** for the password, and click **Login**.

Figure 17: Select Active or Demo Mode



5. Select the operating mode of the device:

Active Mode	Active operation where the WX can compress data, accelerate TCP applications, and manage WAN bandwidth.
Demo Mode	<p>Passive operation where the WX can calculate potential compression 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>Note: To use Demo 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 Demo Mode if the device is installed in the data path.</p> <p>For more information about setting up and using a device in Demo Mode, refer to "Demo Mode" on page 453.</p>



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

If you select **Demo Mode**, click **Next**, and then click **Finish**.

6. If you select **Active Mode**, click **Next** to open the Set Time page.

Figure 18: Set the Time

Do the following:

- a. Enter the current date and time or select **Use NTP Server** and enter the IP address of your NTP server in the **Primary** field. A secondary NTP server is optional.
- b. Select the local time zone.

7. Click **Next** to open the Registration Server Setup page.

Figure 19: Registration Server Setup

Quick Setup HELP

STEP 3: Registration Server Setup

WX devices exchange information with other WX devices in the community through a designated Registration Server. You can make this device the Registration Server or you can direct it to another device that has been designated as a Registration Server.

☒ Direct this device to an existing Registration Server
☐ Make this device the Registration Server

Enter the Registration Server IP address and password below.

If the password does not match the password set on the Registration Server, this device will not be permitted to communicate with the Registration Server. As a result, this device will be unable to form tunnels with other devices in the community.

Reg. Server IP Address

Reg. Server Password

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 compress and accelerate traffic 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 82.

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

- Enter the IP address and password of the current (or future) registration server. If the password is incorrect, this WX will be unable to form tunnels with the other devices in the community.

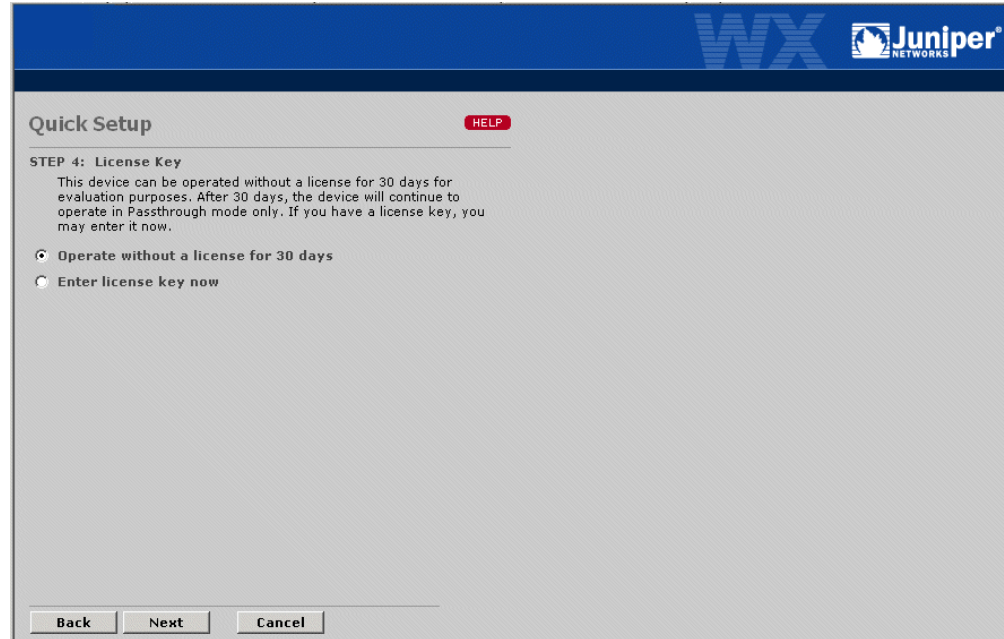


NOTE: If the registration server has not been installed, you can enter its IP address and password in advance.

- Select **Make this device the Registration Server**, and enter the registration server password in both fields. The password is used to authenticate each device, and should be different from the administrator password.

8. Click **Next** to open the License Key page.

Figure 20: Entering a License Key



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 without any processing.

If you have a license key, select **Enter license key now**, and enter the following:

Serial Number	If the serial number is not shown, get the “S/N” number from the back panel.
License Key	Enter your permanent license key.

9. Click **Next**, and then click **Finish**.

The 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 61.
- To use CLI, refer to “Using the Command Line Interface (CLI)” on page 305.

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

- To change the default password for the **admin** account, refer to “Defining Local Users” on page 92.
- Change the default topology setting from mesh to hub or spoke, if appropriate (refer to “Configuring Topology Settings” on page 108). You can also change the default community size from large to small.
- Configure the local routes for the device, as described in “Configuring Local Routes” on page 73. 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 compression, as described in “Advertising Compression Subnets” on page 148.
- Review the available security features, such as limiting operator access to specific IP addresses or subnets, as described in “Configuring AAA” on page 86.
- Review the application definitions provided and add any new ones needed for your network, as described in “Managing Applications” on page 95.
- Configure inbound and outbound bandwidth management, as described in “Applying Quality of Service (QoS) Policies” on page 167.
- Enable traffic acceleration, as described in “Accelerating WAN Traffic” on page 203.

Where to Go Next

After installing a 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 61.
- “Using the Command Line Interface (CLI)” on page 305.

Chapter 3

Configuring Basic Setup Policies

The following topics describe the basic setup procedures:

- “Using the Web Console” on page 61
- “Configuring Basic Setup Policies” on page 63
- “Configuring AAA” on page 86
- “Managing Applications” on page 95



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 305 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 compression, 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. If a Security Alert dialog box appears, click **Yes** to proceed.
3. In the Enter Network Password dialog box, type your user name and password.



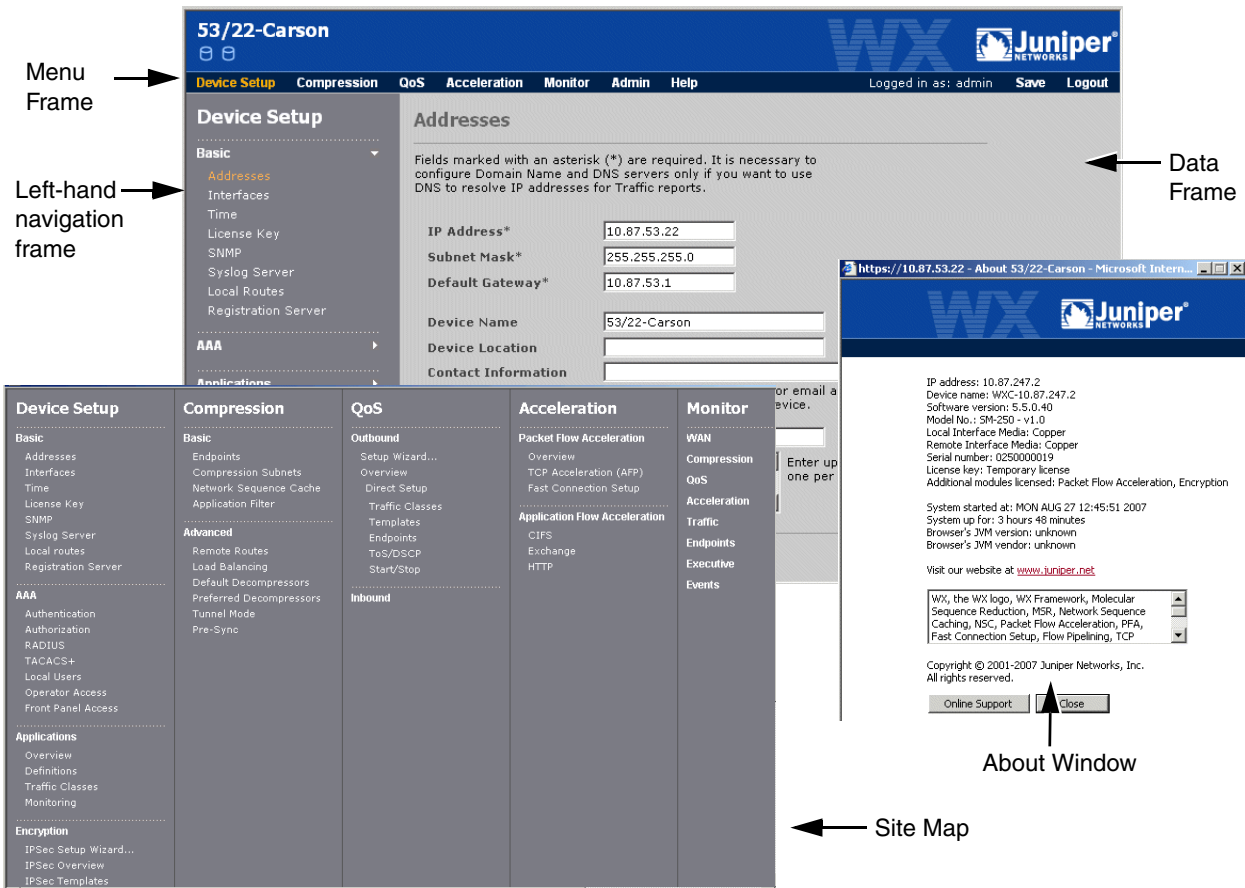
NOTE: When a new device is accessed for the first time, use **admin** and **juniper** for the user name and password, and run Quick Setup (refer to “Running Quick Setup through the Web Console” on page 54).

4. 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 21: 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. Do not use asterisks (*), question marks (?), or angle brackets (< >) in names.

Configuring Basic Setup Policies

The following topics describe the basic configuration procedures:

- “Configuring Device Address and Contact Information” on page 63
- “Configuring the Interface Settings” on page 65
- “Configuring 802.1Q VLAN Support” on page 67
- “Configuring Time Settings” on page 68
- “Obtaining a Permanent License” on page 69
- “Enabling SNMP” on page 71
- “Enabling Syslog Reporting” on page 72
- “Configuring Local Routes” on page 73
- “Configuring Registration Servers and Communities” on page 82

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 **Device Setup** in the menu frame.

Figure 22: Configuring Network Address and Contact Information

53/22-Carson

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

Basic

Addresses

Interfaces

Time

License Key

SNMP

Syslog Server

Local Routes

Registration Server

AAA

Applications

Encryption

Advanced

Addresses

Fields marked with an asterisk (*) are required. It is necessary to configure Domain Name and DNS servers only if you want to use DNS to resolve IP addresses for Traffic reports.

IP Address* 10.87.53.22

Subnet Mask* 255.255.255.0

Default Gateway* 10.87.53.1

Device Name 53/22-Carson

Device Location

Contact Information

Enter name, phone number or email address of the person who will be supporting the device.

Domain Name

DNS Servers

Enter up to 3 IP addresses, one per line.

Submit Reset

2. Specify the following information:

IP address	Enter the IP address of the device. NOTE: If you change the IP address or subnet mask, you must reboot the device. 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 82).
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, which 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	<p>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.</p> <p>When an IP address in the local domain is resolved by one of the specified DNS servers, the local domain name is prepended to the host name shown on the Traffic report.</p> <p>If this field is left blank, only the host names are shown for resolved IP addresses in the local domain. Resolved addresses outside the local domain include the domain name returned by the DNS server.</p>
DNS servers	Enter the IP addresses of up to three DNS servers (one per line) that can be used to resolve IP addresses on the Traffic report (refer to “Traffic Statistics” on page 274).

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 15, WX 20, and WXC 250 have two 10/100 NICs. The WX 60, WX 100, WXC 500, and WXC 590 have 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 344). 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 23).

Figure 23: Using the High Availability Support Feature

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

You can also disable the hardware passthrough so that the router detects the loss of traffic if the WX or WXC device fails (refer to the “embed” CLI command on page 314).

To view or configure the interface settings:

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

Figure 24: 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 compression of VLAN traffic that conforms to the IEEE 802.1Q specification. However, VLAN traffic that uses ISL encapsulation is passed through without compression.

To enable compression of 802.1Q VLAN traffic:

1. Click **Device Setup** in the menu frame, and then click **Interfaces** in the left-hand navigation frame.
2. 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 compress traffic for all the VLANs.
 - **Preserve VLAN ID on output packets.** Select the check box to preserve the VLAN ID in the header of compressed 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 request 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 Compression Subnets page and cannot be advertised for compression.

To include WAN-side routes on the list of compression 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 73), verify that the appropriate subnets are discovered and advertised (refer to “Advertising Compression Subnets” on page 148). Since both LAN and WAN-side subnets will be shown on the Compression Subnets page, be sure to advertise only the true LAN-side subnets.

3. 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 Device Setup page, click **Time** in the left-hand navigation frame.

Figure 25: Configuring the Time Settings for a Device

The screenshot shows the 'Time' configuration page in the Juniper WX CMS. The left-hand navigation frame is expanded to 'Time'. The main content area has a title 'Time' and two radio buttons: 'Use NTP Server' and 'Enter Local Time'. The 'Use NTP Server' option is selected. Below it, there are fields for 'Primary' and 'Secondary' IP addresses. The 'Enter Local Time' option is also visible, with fields for 'Time' (HH:MM) and 'Date' (MM/DD/YYYY). Below these, there is a 'Time Zone' dropdown menu showing '(GMT -08:00) Pacific Time (US and Canada), Tijuana'. There is also a 'Daylight Saving' checkbox labeled 'Automatically adjust time for daylight saving', which is checked. At the bottom, there are 'Submit' and 'Reset' buttons. A note at the bottom right states: 'If you want to preserve the changes of NTP Server, Time Zone, and Daylight Saving you must save the configuration to flash memory using the "SAVE" tab after submit.'

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. Note that the WX CMS server also can be used as an NTP server.
 - 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 saving** 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 compression.

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).
- If you purchased license upgrades, you will need the Authorization Code Certificate that was emailed to you in electronic PDF format. If you operate the platform at its base speed, only the serial number is needed to generate a permanent license.
- User ID and password to access the License Key server at:

http://www.juniper.net/generate_license

If you lose the license key, you can use the License Key server to retrieve your current license key.



NOTE: If you enable IPSec during the trial period, but do not obtain a permanent license for it, be sure to disable IPSec before applying the permanent license. If you apply the license without disabling IPSec, you will not be able to access the device.

If you have any problems with the licensing process, open a case with the Juniper Case Manager at <http://www.juniper.net/cm>. To call from the United States, Canada, or Mexico, dial +1-888-314-JTAC. To call from other locations, check the list of local support centers at http://www.juniper.net/support/support_contacts.html or dial +1-408-745-9500.

To install a permanent license key:

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

Figure 26: Replacing the Temporary License with a Permanent License

The screenshot shows the Juniper Networks Device Setup web interface. The top navigation bar includes links for Device Setup, Compression, QoS, Acceleration, Monitor, Admin, and Help. The user is logged in as 'admin'. The left-hand navigation frame shows the 'Device Setup' menu with sub-items: Basic, AAA, Applications, Encryption, and Advanced. The 'License Key' page is displayed, showing the current license status. The current license is a 'Temporary license' with 'Unlimited' maximum throughput and 'Packet Flow, Acceleration, Encryption' additional modules. The expiration is '20 days: 8 hours: 31 minutes'. A text box indicates that a license key can be obtained by calling Juniper's Customer Care at 800-638-8296. Below this, there are input fields for 'Serial number' (containing '0500002013') and 'License key'. An 'Online Service...' button is also present. At the bottom, there 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** and log in.
 - b. Enter your contact information and the device serial number, and click **Submit**.
 - c. Enter your Authorization Code(s) for the desired device speed and other features, and click **Yes**. If you enter only the serial number, the device is licensed for the base speed.
 - 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

The following support is provided for SNMP:

- 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 427.

To enable SNMP:

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

Figure 27: Enabling SNMP

The screenshot shows the Juniper Device Setup interface for the '53/22-Carson' device. The left-hand navigation frame has 'Device Setup' selected, with a sub-menu showing 'Basic' (Addresses, Interfaces, Time, License Key, SNMP, Syslog Server, Local Routes, Registration Server), 'AAA', 'Applications', 'Encryption', and 'Advanced'. 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 field containing '.....'.
- Write Community String:** A text field containing '.....'.
- Trap Enabled:** A checkbox labeled 'Yes' is unchecked.
- Authentication Trap Enabled:** A checkbox labeled 'Yes' is unchecked.
- Trap Destinations:** A table with two columns: 'IP Address' and 'Community String'. The first row has '10.87.33.33' in the IP Address field and '....' in the Community String field. There are 'DELETE' and 'ADD' buttons next to each row.

At the bottom of the main content area are 'Submit' and 'Reset' buttons. A small text box on the right side of the Trap Destinations section provides instructions: 'To create a trap destination, enter the IP address and community string and click **ADD**. When you are finished adding trap destinations, click **Submit**.'

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). To add trap destinations (up to 10), enter the IP address and community string (up to 25 characters), and click Add. To delete a trap destination, click **Delete** next to the destination.
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

Syslog messages can be sent 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 the syslog messages, refer to “SNMP Traps and Syslog Messages” on page 427.

To send syslog messages to a specific facility (local1 through local7), use the CLI command described in “configure syslog” on page 394.

To enable syslog reporting:

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

Figure 28: Enabling Syslog Reporting

The screenshot shows the Juniper Device Setup interface. The left-hand navigation frame is expanded to show the 'Device Setup' section, with 'Syslog Server' highlighted. The main content area is titled 'Syslog Server' and contains the following configuration options:

- Syslog enabled:** A checkbox labeled 'Yes' is currently unchecked.
- Syslog servers:** A text area for entering IP addresses. A tooltip indicates: 'Enter IP addresses, one per line with a maximum of 5 servers.'
- Syslog message severity:** Four checkboxes are shown: 'Critical' (checked), 'Error' (checked), 'Informational' (unchecked), and 'Notice' (unchecked). A tooltip indicates: '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.
 - **Notice:** Informational messages about unusual events that are not errors.
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). 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 compression subnets. Compression subnets can then be advertised so that other devices in the community can compress and accelerate traffic sent to those subnets (refer to "Advertising Compression Subnets" on page 148). By default, only the subnets you select are advertised.

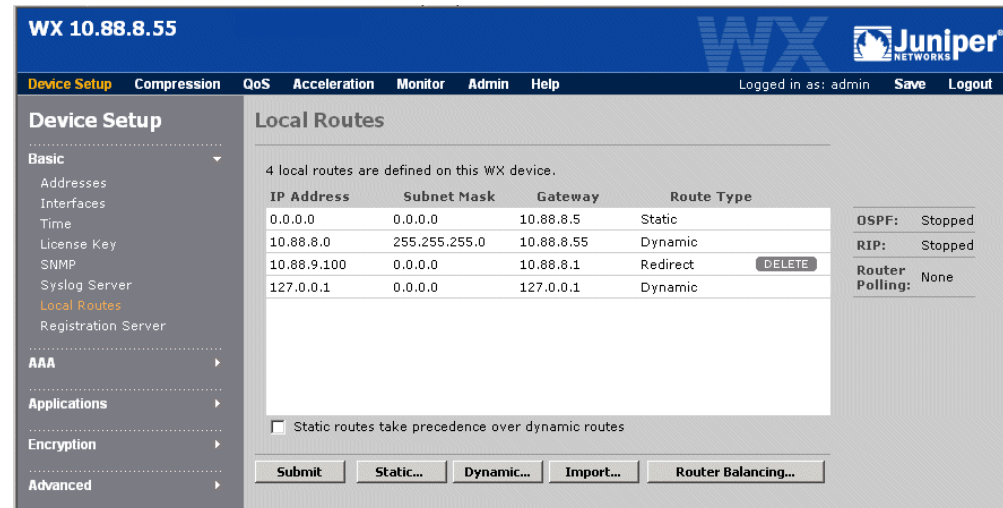


NOTE: In some environments, such as VLAN, some routes on the Local interface may be discovered only on the Remote interface. To advertise these subnets, you must enable the WAN compression subnet option through the CLI to display all routes on the list of compression subnets (refer to "configure reduction-subnet" on page 378).

To configure local routes:

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

Figure 29: Configuring Local Network Routes



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 compression 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 76
 - “Enabling Route Polling” on page 77
 - “Importing a Routing Table” on page 78
 - “Enabling Route-Based Router Balancing” on page 80
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:



NOTE: To include a relative cost for a static route, use the CLI command (refer to “configure route” on page 384).

1. On the Local Routes page, click **Static**.

Figure 30: 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 compression 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 compression subnets option is enabled (refer to “Advertising Compression Subnets” on page 148).

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 31: 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 compression subnets, but they are not advertised automatically to the other WX and WXC devices. If the WAN compression subnets option is enabled, all routes discovered on the Remote interface are also added to the list of compression subnets (refer to “Advertising Compression Subnets” on page 148).

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 387).

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 32: Dynamically Obtaining a Routing Table from a Cisco Router

3. Specify the following information:

Poll router	Enter the IP address of a Cisco router and the port number used for rsh (the standard port is 514). NOTE: 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 remote user name specified on the Cisco router.
Remote user name	Enter a remote user name that matches the local 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 compression subnets, but they are not advertised automatically to the other WX devices. If the WAN compression subnets option is enabled, routes discovered on the Remote interface are also added to the list of compression subnets (refer to “Advertising Compression Subnets” on page 148).

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 315).

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

To import routes from an FTP server:

1. On the Local Routes page, click **Import**.

Figure 33: 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 RIP, OSPF, or route polling is enabled.

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 compression subnets, but they are not advertised automatically to the other WX devices. If the WAN compression subnets option is enabled, routes discovered on the Remote interface are also added to the list of compression subnets (refer to “Advertising Compression Subnets” on page 148).

Enabling Route-Based Router Balancing

To balance the compressed 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 384.

Using route-based balancing, the device can distribute compressed traffic across up to four different gateways. In Figure 34, 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 34: 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 35).

Figure 35: Common Routes with Equal Cost Paths

53/22-Carson

WX

Juniper
NETWORKS

Device Setup Compression QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Device Setup

Basic

Addresses

Interfaces

Time

License Key

SNMP

Syslog Server

Local Routes

Registration Server

AAA

Applications

Encryption

Advanced

Local Routes

3 local routes are defined on this WX device.

IP Address	Subnet Mask	Gateway	Route Type
0.0.0.0	0.0.0.0	10.87.53.1	Static
10.87.53.0	255.255.255.0	10.87.53.22	Dynamic
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
	255.255.255.0	192.168.0.2	Dynamic

OSPF: Stopped
RIP: Stopped
Router Polling: None

☐ Static routes take precedence over dynamic routes

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

Equal cost paths to the same destination

To enable router balancing:

1. On the Local Routes page, click **Router Balancing**.

Figure 36: Configuring Router Balancing

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The user is logged in as 'admin'. The left sidebar shows the 'Device Setup' menu with options like Basic, Addresses, Interfaces, Time, License Key, SNMP, Syslog Server, Local Routes, Registration Server, AAA, Applications, Encryption, and Advanced. The main content area is titled 'Local routes > Router balancing'. It contains a description: 'The rule selected below determines how traffic is directed when more than one gateway exists for a given subnet.' There are four radio button options: 'Off' (selected), 'Per-destination', 'Per-packet', and 'Flow based'. Each option has a corresponding description of how traffic is directed. At the bottom, there are 'Submit', 'Reset', and 'Cancel' buttons.

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



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

- **Flow based**—Traffic is distributed over available routers based on source and destination IP addresses and ports.
3. Click **Submit** to 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 service tunnel to each of those devices (also called “endpoints”).

Since data compression occurs only between devices in the same community, in large deployments you can limit the number of devices in each community. To send compressed traffic between communities, you can create a hierarchical structure where selected devices reside in multiple communities (refer to “Configuring Tunnel Switching” on page 163).

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 compression between any two devices in the same community, as described in “Configuring Endpoints for Compression” on page 145.

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

- “Defining Registration Servers and Passwords” in the next section
- “Defining Communities” on page 84

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 **Device Setup** in the menu frame, and then click **Registration Server** in the left-hand navigation frame.

Figure 37: Configuring Registration Server Settings

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The 'Device Setup' sidebar on the left lists various configuration categories: Basic, AAA, Applications, Encryption, and Advanced. The 'Registration Server' section is selected, showing a form with the following options:

- ☐ Change registration server password: Includes fields for Old password, New password, and Verify new password. A warning message states: 'WARNING: Changing the password on the registration server will disrupt communication with reporting WX devices. Passwords on those devices must be updated to match the new password to restore communication.'
- ☐ Change SECONDARY registration server: Includes radio buttons for 'Use IP address' (selected, with IP 192.168.243.2) and 'No secondary registration server'. A note says: 'Also, if you want to preserve the changes, you must save the configuration to flash memory using the 'SAVE' tab after submit.'
- ☐ Transfer registration server designation to another device: Includes an IP address field.

Buttons at the bottom include 'Submit', 'Reset', and 'Communities...'.

- 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.

- 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.

- 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:

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

- 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.

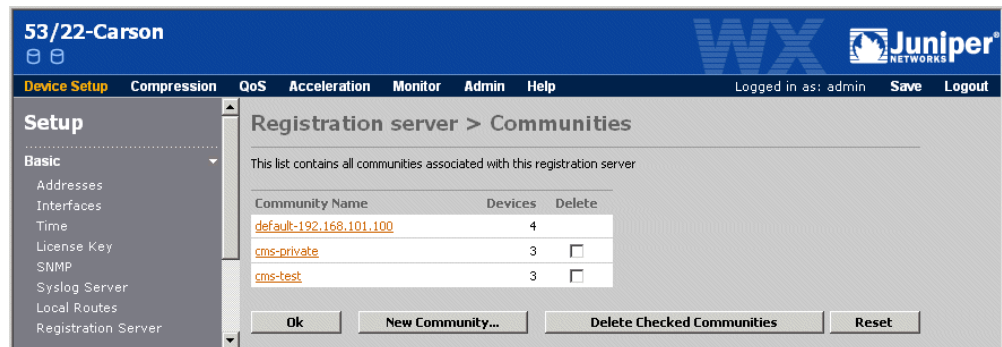
Defining Communities

Data compression occurs only between WX and WXC devices in the same community, so in large deployments you can limit the number of devices in each community. To send compressed traffic between communities, you can create hierarchical communities (refer to “Configuring Tunnel Switching” on page 163).

To configure the communities on a registration server:

- Log in to the device that acts as the registration server.
- Click **Device Setup** in the menu frame, click **Registration Server** in the left-hand navigation frame, and then click **Communities** at the bottom of the page.

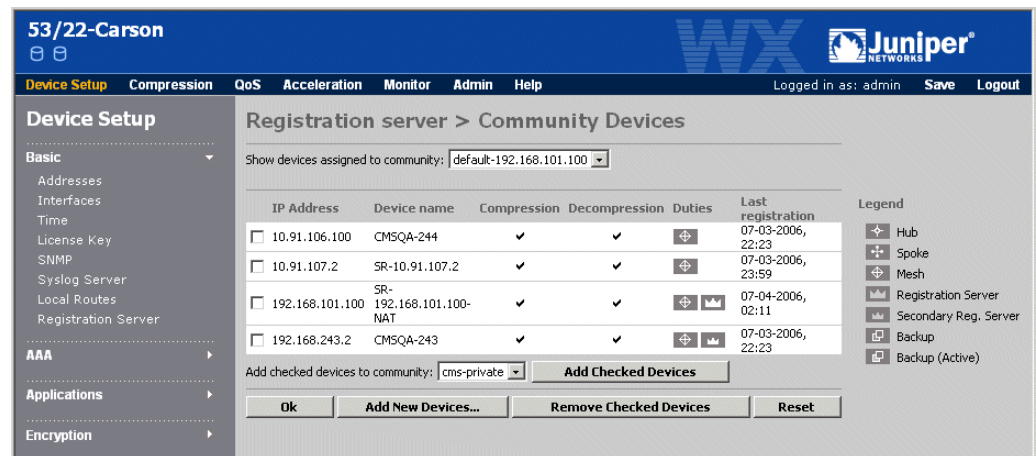
Figure 38: Viewing the Communities for a Registration Server



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.
- To define the devices in a community, click the community name.

Figure 39: 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 **Compression** and **Decompression** columns indicate whether the compression and decompression engines are activated. To activate or deactivate the compression and decompression engines, refer to “Configuring Endpoints for Compression” on page 145.
- The **Duties** column can contain the following icons:

Duty icon	Description
	Hub —The device is designated as a Hub. Each device attempts to form a service tunnel with a hub before creating tunnels to other WX devices (refer to “Configuring Topology Settings” on page 108).
	Spoke — The device is designated as a spoke in a Hub and Spoke topology. By default, a spoke compresses and decompresses data only for the hub device(s).
	Mesh — The device is designated as part of a mesh topology.
	Registration Server — Indicates that this device is the Registration Server for the community.
	Secondary Registration Server — Indicates that this device is the Secondary Registration Server for the community.
	Backup and Backup (Active) — 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 336.

- 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 by user name and password or by means of 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 the local database and remote RADIUS and/or TACACS+ servers. RADIUS and TACACS+ support 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 and TACACS+ servers provide the connection to the back-end authentication infrastructure, and existing user entries in the directory can be used for authentication and authorization.

Multiple RADIUS and/or TACACS+ servers can be configured for redundancy. You can use both the local database and remote servers, so that some users are authenticated locally and others are authenticated remotely.

The following topics describe how to define the authentication and authorization settings, remote servers, and other security features:

- “Selecting Authentication Methods” in the next section
- “Enabling Authorization Checking” on page 88
- “Defining RADIUS Servers and Server Groups” on page 89
- “Defining TACACS+ Servers” on page 91
- “Defining Local Users” on page 92
- “Securing Operator Access” on page 94
- “Securing Front Panel Access” on page 95

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, TACACS + servers, 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, TACACS + servers, and local user accounts, refer to “Defining RADIUS Servers and Server Groups” on page 89, “Defining TACACS + Servers” on page 91, and “Defining Local Users” on page 92.

To select the authentication methods for each user interface:

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

Figure 40: Selecting Authentication Methods

The screenshot shows the Juniper WX 10.88.8.55 Device Setup page. The left-hand navigation frame has 'AAA' expanded, with 'Authentication' selected. The main content area is titled 'Authentication' and contains three sections: Console, SSH, and Web. Each section has a table with 'Order' and 'Method' columns. The Console section has 4 rows, with the first row set to 'Local'. The SSH section has 4 rows, with the first row set to 'Local'. The Web section has 4 rows, with the first row set to 'Local'. Below the SSH section, there is a 'Disconnect user' section with a radio button for 'After' and a dropdown menu set to '3' failed attempts. There is also a 'Never' radio button. At the bottom of the page, there are 'Submit' and 'Reset' buttons. On the right side of the page, there is explanatory text about authentication methods and an exception rule.

WX 10.88.8.55

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

- Basic
 - Addresses
 - Interfaces
 - Time
 - License Key
 - SNMP
 - Syslog Server
 - Local Routes
 - Registration Server
- AAA**
 - Authentication**
 - Authorization
 - RADIUS
 - TACACS+
 - Local Users
 - Operator Access
 - Front Panel Access
- Applications
- Encryption
- 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 3 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' or 'TACACS+', then if the Local method does not find a username entry in the local database, instead of issuing a 'fail', the second method will be used.

If there is no response from any of the selected methods, then access is denied. 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> <ul style="list-style-type: none"> ■ RADIUS: <i>group_name</i>. Attempts to authenticate users by accessing the RADIUS servers in the specified group (refer to “Defining RADIUS Servers and Server Groups” on page 89). The servers are accessed in the order specified in the group. If all RADIUS servers are down or do not respond, the next method is tried. ■ TACACS+ . Attempts to authenticate users by accessing the TACACS+ servers in the order specified (refer to “Defining TACACS+ Servers” on page 91). If all TACACS+ servers are down or do not respond, the next method is tried. ■ Local. Attempts to authenticate users locally. ■ None. Login not required. Can be used alone or after the last RADIUS group. Cannot be used directly after Local. <p>Each method is tried in the order specified. Authentication stops with the first success or failure. However, if Local is the first method, the next method is tried if the user is not defined locally.</p>
SSH	<p>Select up to four authentication methods for users logging in using the SSH protocol. Same options as the console, except that None is not available (authentication is required).</p> <p>Select the number of unsuccessful SSH login attempts allowed before a user is disconnected (1 to 10) or select Never.</p>
Web	<p>Select up to four authentication methods for users logging in through the Web. Same options as the console, except that None is not available (authentication is required).</p>

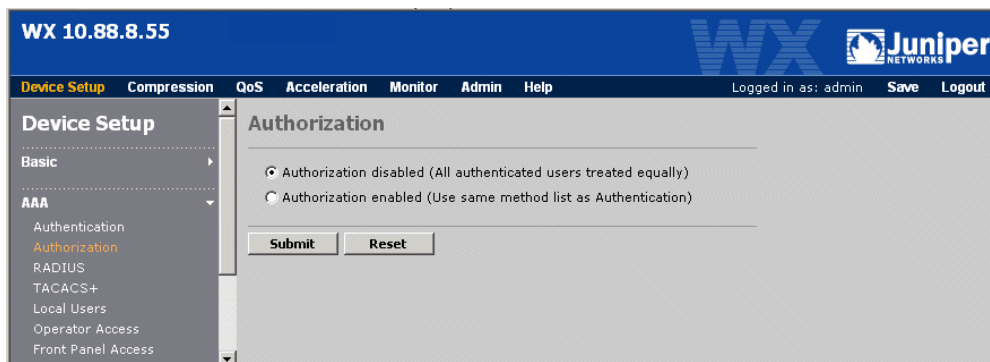
3. Click **Submit** to activate the changes.
4. To retain your changes when the device is restarted, click **Save** in the menu frame.

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

Figure 41: 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 a RADIUS or TACACS+ server 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

A WX or WXC device acts as 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.

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 87.

To define RADIUS servers and server groups:

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

Figure 42: Defining RADIUS Servers and Server Groups

53/22-Carson

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
 - Authentication
 - Authorization
 - RADIUS**
 - Local Users
 - Operator Access
 - Front Panel Access
- Applications
- Encryption
- Advanced

RADIUS

RADIUS Client

Source IP Address: 10.87.53.22

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 ☐

Delete

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:

RADIUS 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 (1 to 65535) used for authentication (default is 1812).
Timeout	Enter the number of seconds (1 to 60) that the device waits for the server to respond (default is 3).
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 (default is 3).
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 TACACS+ Servers

You can define up to five TACACS+ servers to authenticate WX users. The servers are accessed in the order specified. WX devices conform to the TACACS+ protocol specification 1.78 (draft-grant-tacacs-02.txt).

The following attributes can be returned from the TACACS+ server:

- **idletime = n**. Indicates the number of consecutive minutes a user session can be idle before the connection is closed (a zero indicates no idle timeout).
- **priv-lvl = n**. Indicates a user's access privileges (0 to 15).
- **packet-capture-allowed = 1/0**. Indicates whether packet captures are allowed.

To enable the use of TACACS+ servers for authentication, refer to “Selecting Authentication Methods” on page 87.

To define TACACS+ servers:

1. In the Device Setup page, click **AAA** in the left-hand navigation frame, and then click **TACACS+**.

Figure 43: Defining TACACS+ Servers

The screenshot shows the configuration interface for a WX device (10.88.8.55). The left-hand navigation pane is expanded to 'AAA', and 'TACACS+' is selected. The main content area is titled 'TACACS+' and contains the following sections:

TACACS+ Client

Source IP Address:

Order	TACACS+ Servers	IP Address	Port	Timeout (sec)	Re-transmit	Delete
1	TACACS1	10.10.20.30	49	10	3	<input type="checkbox"/>

New Server... Maximum 5 TACACS+ servers

Buttons: **Submit** **Reset**

From the TACACS+ page, you can:

- Add new servers, as described in Step 2.
- Change a server. Click the server name, make any needed changes, and click **Submit**.
- Change the IP address in the **Source IP Address** field (defaults to the WX device's IP address). Replies from the TACACS+ server are sent to the source address.
- Delete servers. Select the check box next to the servers you want to delete, and click **Submit**.

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

TACACS+ Server Name	Enter the TACACS+ server name (up to 31 characters).
IP Address	Enter the IP address of the server.
Authentication Port	Enter the UDP port number (1 to 65535) used for authentication (default is 49).
Timeout	Enter the number of seconds (1 to 60) that the WX waits for the server to respond (default is 10).
Retransmit	Enter the number of times (1 to 100) that requests are retransmitted to a server before trying the next server, if any (default is 3).
Shared Secret Key	Enter the secret key (up to 31 characters) used to access the server. The same key must be configured on the TACACS+ server.

3. Click **Submit** to activate the changes.
4. 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 for local authentication. Each user can have full (admin) or read-only access privileges. The predefined **admin** account has a default password of **juniper**. To ensure secure access to the device, you should change the password periodically. To specify how users are authenticated (locally and/or remotely), refer to “Selecting Authentication Methods” on page 87.

To define local user accounts:

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

Figure 44: Defining Local Users

WX 10.88.8.55

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
 - Authentication
 - Authorization
 - RADIUS
 - TACACS+
 - Local Users
 - Operator Access
 - Front Panel Access
- Applications
- Encryption
- Advanced

Local Users

User Name	Privilege Level	Packet Capture	Idle Timeout (seconds)	Delete
admin	Read Write	Allow	Never	<input type="checkbox"/>

New User...

Submit Reset

From the Local Users page, you can:

- Add a new user account, as described in Step 2.
- 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 (read-write) or read-only privileges.
Packet Capture	Select Allow or Disallow to indicate whether the user can run packet capture.
Idle Timeout	Enter the number of minutes before an idle user is logged out (the default is 30) or select Never .
Password	Enter the password twice (from 4 to 64 characters).



NOTE: Authorization checking is disabled by default, so that all authenticated users have read-write access and a 30-minute idle timeout. If you create read-only user accounts or change the default idle timeout, you must enable authorization checking (refer to “Enabling Authorization Checking” on page 88).

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

Figure 45: 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 that appear 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.

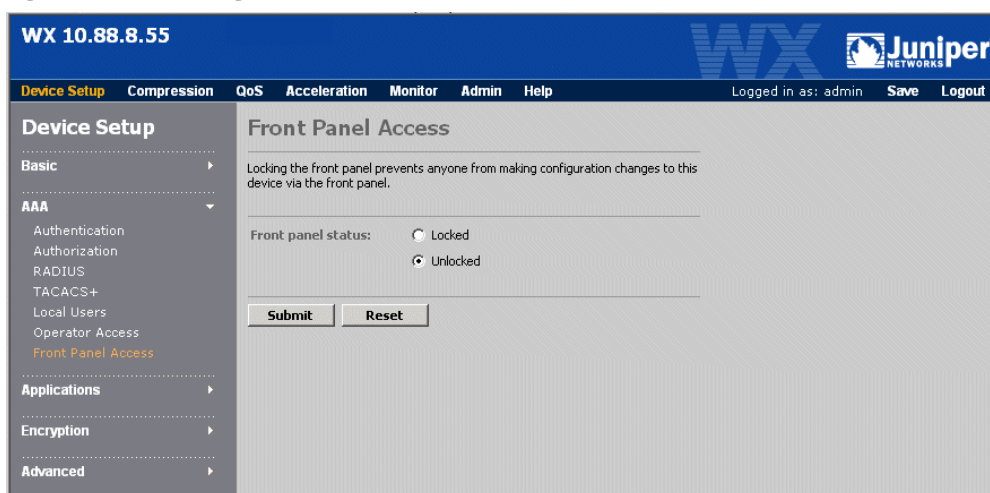


NOTE: The WX 15, WX 20, WXC 250, WXC 590, and the latest model of the WX 100 do not have a front panel keypad. Also, locking the front panel on the earlier WX 100s does not lock the front panels of the client devices.

To lock the front panel keypad:

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

Figure 46: 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.

Managing Applications

The following topics describe how to manage applications.

- “About Application Definitions” in the next section
- “Viewing the Application Overview” on page 98
- “Configuring Application Definitions” on page 99
- “Assigning Applications to Traffic Classes” on page 104
- “Monitoring Applications” on page 105

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 compression, as described in “Compressing Traffic by Application” on page 151.
- Accelerate the application’s traffic (if data compression is enabled), as described in “Accelerating WAN Traffic” on page 203.
- Assign the application to a traffic class, as described in “Assigning Applications to Traffic Classes” on page 104. 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 105.
- Trigger events based on the application’s performance, as described in “Configuring Events” on page 140.
- View compression and acceleration statistics for monitored applications, as described in “Monitoring and Reporting” on page 245.

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	47	5190-5193
CIFS	16	139, 445
Clearcase	34	371
CVS	44	2401
DNS	25	53

Application	Order	Port Numbers
Exchange	30	135 Note: Port 135 is the startup port; other ports are learned dynamically. This definition applies only to Exchange traffic for Windows clients, not Web clients.
Filenet	51	32768-32774
FTP	11	20-21 Note: Non-default FTP ports are learned dynamically.
Groupwise	40	1677
H.248	8	2945, 1039 TCP
H.323	9	1719-1720
Hostname Resolution	31	42
HTTP	14	80, 8080
HTTPS	22	443
ICA (Citrix)	19	1494
ICMP	32	Protocol 1 (no ports specified)
Kerberos	27	88
LDAP	26	389
Lotus Notes	17	1352
Mail	13	25,110,143
Microsoft SQL Monitor	1	1434
MS Streaming	41	1755
MS Terminal Services	28	3389
MySQL	6	3306
NetApp SnapMirror	50	10566
NetBios	15	137, 138
NFS	43	2049
Novell NCP	38	524
Oracle	21	No ports specified
Oracle SQL*Net	4	1529 TCP
Oracle SQL*Net v1	3	1525
Oracle SQL*Net v2	2	1521 TCP
PCAnywhere	48	5631-5632
Printer	37	515
RADIUS	42	1812, 1813
RTP	7	2048-3048 UDP
RTSP	39	554
SAP	46	3200, 3300-3388, 3390-3399, 3600-3699
Shell	35	514 TCP
SNMP	29	161-162

Application	Order	Port Numbers
SNTP	24	123
SQL Server	18	No ports specified
SSH	23	22
Sybase	20	No ports specified
Symantec Anti-Virus	45	2967
Syslog	36	514 UDP
TACACS	33	49
Telnet	12	23
Traceroute	52	33434-33534 UDP
UniSQL	5	1978
UniSQL Java	5	1979 TCP
XWindows	49	6000-6063

Viewing the Application Overview

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

To view the application overview:




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

Figure 47: Application Overview Page

The screenshot displays the 'Application Overview' page for a Juniper WXC-10.88.10.250 device. The page features a navigation sidebar on the left with options like Basic, AAA, Applications (selected), Encryption, and Advanced. Under 'Applications', 'Overview' is highlighted. The main content area shows a table of applications with their respective traffic classes and enabled features. The features include IPSec Encryption, SSL Optimization, Compression, NSC, TCP Acceleration, Fast Connection Setup, CIFS Acceleration, Exchange Acceleration, HTTP Acceleration, and Monitoring. Each feature is represented by a lock icon (enabled) or a circle with a slash (disabled).

Application Name	Traffic Class	IPsec Encryption	SSL Optimization	Compression	NSC	TCP Acceleration	Fast Connection Setup	CIFS Acceleration	Exchange Acceleration	HTTP Acceleration	Monitoring
AOL	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
CIFS	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Clearcase	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
CVS	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
DNS	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Exchange	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Filenet	Default	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled

- 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 104.
IPSec Encryption	Indicates whether the application’s traffic is encrypted with IPSec (refer to “Defining the IPSec Application Filter” on page 239.). A  indicates IPSec is required, a  indicates IPSec is used if available, and a  indicates IPSec is not configured or is not used for the application.
SSL Optimization	Indicates whether the application’s traffic is enabled for SSL optimization (refer to “Enabling Applications for SSL Optimization” on page 243.).
Compression	Indicates whether the application’s traffic is compressed (refer to “Compressing Traffic by Application” on page 151.).
NSC	Indicates whether Network Sequence Caching is used for data compression (refer to “Compressing Traffic by Application” on page 151). NSC requires a hard disk, and applies only to WXC devices.
TCP Acceleration	Indicates whether the application’s traffic is accelerated using TCP Acceleration (refer to “Enabling TCP Acceleration by Application” on page 212).
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 213).
CIFS Acceleration	Indicates whether CIFS traffic for the application is accelerated (refer to “Enabling Microsoft CIFS Acceleration” on page 218).
Exchange Acceleration	Indicates whether Exchange traffic for the application is accelerated (refer to “Enabling Microsoft Exchange Acceleration” on page 221).
HTTP Acceleration	Indicates whether HTTP traffic for the application is accelerated (refer to “Enabling HTTP Acceleration” on page 223).
Monitoring	Indicates whether you can view statistics for the application (refer to “Monitoring Applications” on page 105).

Configuring Application Definitions

For each defined application, the Applications Definitions page lists the application’s type, source and destination ports, and whether the application is encrypted by SSL.

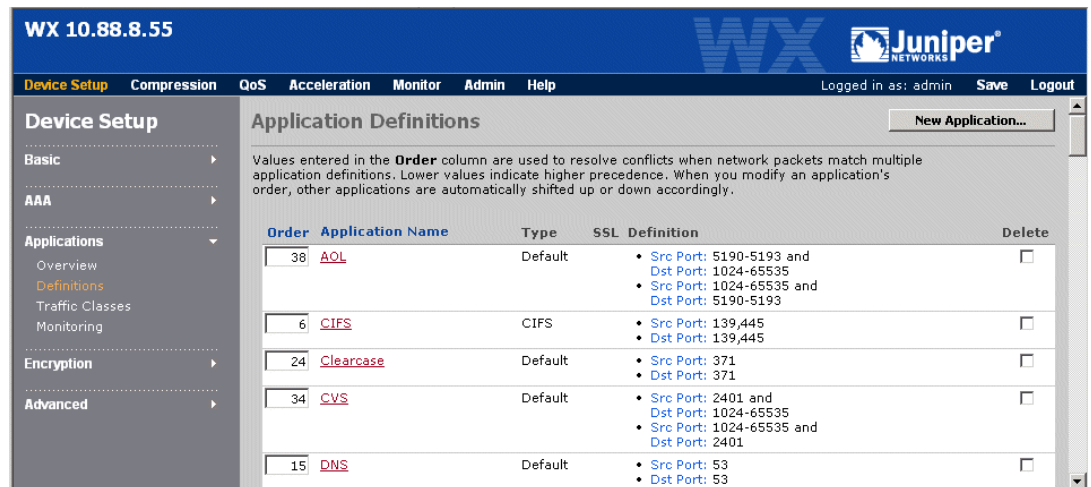


NOTE: To add an application definition by selecting an undefined application from the Traffic report, refer to “Traffic Statistics” on page 274.

To add or change application definitions:

- In the Device Setup page, click **Applications** in the left-hand navigation frame, and then click **Definitions**.

Figure 48: Application Definitions Page



From the Application Definitions page, you can:

- Add a new application definition, as described in Step 2 through Step 4.
- Change an application definition. Click the application name, make any needed changes, and click **Submit**.
- Change a definition's order number. Type a new value in the **Order** field, and click **Submit** to renumber the definitions. The new value cannot exceed the highest value in the current range. The definitions are compared against the traffic in ascending order.
- Delete application definitions. Select the check box next to the applications you want to delete, and click **Submit**.

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

Figure 49: Defining New Applications

WX 10.88.8.55

Device Setup Compression QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
- Applications
 - Overview
 - Definitions
 - Traffic Classes
 - Monitoring
- Encryption
- Advanced

Application Definitions > New

Application Name:

Application Type:

SSL Encrypted ☐ Yes

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"/>

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	<p>Select one of the following application types:</p> <ul style="list-style-type: none"> ■ Default. No special processing. ■ CIFS. Apply to CIFS application definitions whose traffic you want to accelerate (refer to “Enabling Microsoft CIFS Acceleration” on page 218). 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 221). 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 223). Also allows a URL to be specified for pattern matching.
SSL Encrypted	<p>Select the check box if the application is encrypted by SSL. Only applications that have this setting can be enabled for SSL optimization (refer to “Enabling Applications for SSL Optimization” on page 243). Note that the application type must be Default for applications that use SSL.</p>

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	<p>Enter a source IP address or subnet. The general format is:</p> <p>address/subnetmask</p> <p>A blank or an asterisk (*) with no subnet mask indicates any source IP address.</p>
Source Port	<p>Enter a source port number, a series of comma-separated port numbers, or a range of port numbers separated by a hyphen (-). A blank indicates any port. For a list of common application ports, refer to “Common Application Port Numbers” on page 451.</p>
Destination Address	<p>Enter a destination IP address or subnet (same format as the source address). A blank or asterisk (*) indicates any destination IP address. Typically, source and destination addresses are specified in separate rules so that a match occurs on either one. A rule that specifies source and destination addresses will match only the traffic between those addresses.</p>
Destination Port	<p>Enter one or more destination port numbers (same format as the source port). A blank indicates any port. Typically, source and destination ports are specified in separate rules so that a match occurs on either one. A rule that specifies source and destination ports will match only the traffic between those ports.</p>
Protocol	<p>Select an application protocol or select Any to indicate TCP or UDP. You can also type in a protocol number (0 to 134). By default, a match can occur on any TCP or UDP packet.</p> <p>NOTE: Any protocol defined by number is added to the Any list of defaults that applies to each rule that does not specify a protocol. To use application pattern matching (described below), select TCP.</p>

- 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>IP Precedence. Select an IP precedence value (0 through 7).</p> <p>DSCP. 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 196.</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>< host > / < uri ></p> <p>Where:</p> <p><host> 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><uri> 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**.
- 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 compression begins automatically. The new application can be viewed individually on monitoring reports, provided you have not exceeded the maximum number of monitored applications (40). To view or change the monitored applications, refer to “Monitoring Applications” on page 105.

If you do not see any traffic for the application (refer to “Monitoring and Reporting” on page 245), 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 compression, as described in “Compressing Traffic by Application” on page 151. 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 203.

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.

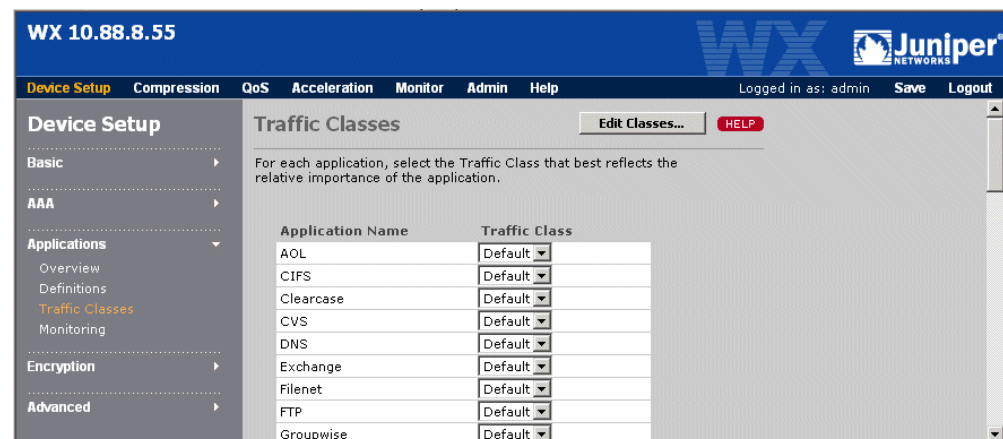
For more information about outbound QoS and Multi-Path, refer to:

- “Understanding Outbound Bandwidth Management” on page 168
- “Configuring Policy-Based Multi-Path” on page 129

To define traffic classes and assign applications to each class:

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

Figure 50: Assigning Applications to 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.
4. To retain your changes when the device is restarted, click **Save** in the menu frame.

Monitoring Applications

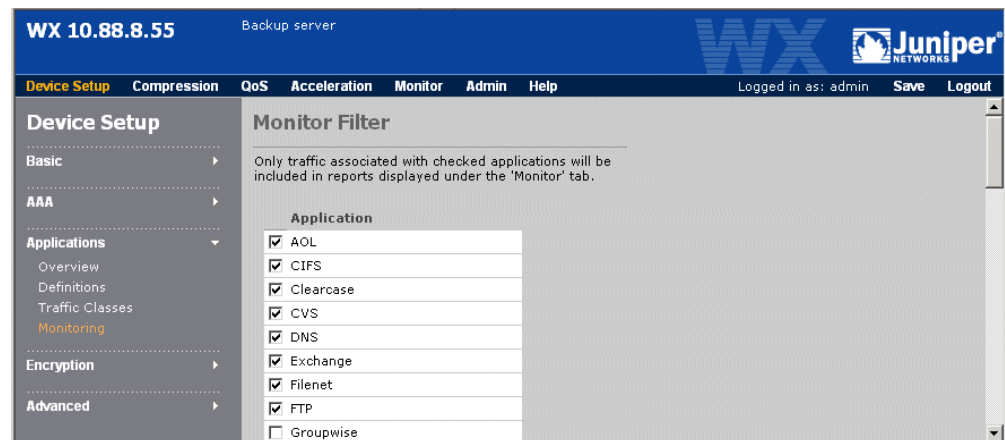
Monitoring an application lets you view compression and acceleration statistics for the application. You can monitor up to 40 applications. All unmonitored applications are placed in the “Others” category on reports. For more information about monitoring statistics, refer to “Monitoring and Reporting” on page 245.

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 99.

To select applications to be monitored:

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

Figure 51: Selecting Applications for Monitoring



2. Select the check box next to each application (up to 40) for which you want to view compression and acceleration statistics. All uncompressed or 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 compression 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:

- “Configuring Topology Settings” on page 108.
- “Using Source/Destination Filters” on page 112
- “Configuring the ARP Table” on page 114
- “Defining the Prime Time” on page 115
- “Configuring Packet Interception” on page 116
- “Configuring Policy-Based Multi-Path” on page 129
- “Configuring WAN Performance Monitoring” on page 138
- “Configuring Events” on page 140
- “Configuring Multiple Tunnels Between WX 100 Servers” on page 144

Configuring Topology Settings

The topology settings determine whether the device attempts to form a tunnel with all WX devices in the same community, and affects the maximum number of tunnels the device can support.



NOTE: The topology setting is not directly related to the topology of your network, but determines the automatic creation of tunnels between WX devices.

Selecting a Topology


Table 2 describes the topology settings and how they should be used.

Table 2: Topology Settings and Recommended Use


Setting	Description	Recommended Use
Hub	Hub devices attempt to form outbound tunnels to (and accept inbound tunnels from) all devices in the same community. You must also select a community size (small or large) to indicate the range of devices supported by the hub (refer to “Selecting the Community Size” on page 110).	If you have a mixture of WX and/or WXC models, or tunnels are not required between all devices, set one or more of the highest-capacity devices as hubs. Select the same community size on all hubs. NOTE: A hub and spoke topology assumes that traffic volume is greater from the hub to the spokes. If traffic is substantially greater in the reverse direction (from the spokes to the hub), use the mesh topology setting.
Spoke	Spoke devices attempt to form tunnels only with hub devices. Spokes also accept tunnels from all devices in the same community, but give preference to hubs when resources are limited. Each spoke uses the same community size as the hubs (see Table 3 on page 110).	Use for non-hub devices in a hub and spoke topology. On each spoke device, you can manually enable tunnels to other spokes as needed (refer to “Configuring Endpoints for Compression” on page 145). A WX 15 must be a spoke, but a WX 100 is never a spoke.
Mesh	Mesh devices attempt to form outbound tunnels to (and accept inbound tunnels from) all devices in the same community. You must also select the small or large community size.	Use if all devices are the same WX or WXC model, or tunnels are required between all (or most) endpoints. Select the same community size on all mesh devices.
Point-to-Point	Same as mesh, except that the community must be limited to two devices of the same type (formerly called the “max-mem” setting).	Use only if both devices are the same WX or WXC model, with the same version of WXOS. Typically used to maximize throughput between two data centers.

Partial Mesh Example

For the partial mesh shown in Figure 52, in most cases you would designate the core devices as hubs and the remote devices as spokes, and then manually configure additional tunnels between the spokes. However, if the traffic volume from the remote sites to the data center is substantially greater than in the reverse direction, then set each device to the mesh topology, and manually disable any unneeded tunnels.

Figure 52: Partial Mesh Example**Tiered Network Example**

In general, mixing hub, spoke, and mesh devices in the same community is not recommended. In Figure 53, if the core devices form tunnels primarily with the regional hubs, you can configure the core sites as mesh devices, and then manually disable tunnels to the remote spokes (the spokes form outbound tunnels only to the hubs). If tunnels are needed between the core devices and remote sites, designate the core devices as hubs.

Figure 53: Tiered Network Example

Selecting the Community Size

Table 3 shows the range of devices supported by each device type for the small and large community sizes, and for the Point-to-Point option. The Point-to-Point option allocates all available memory for a limited number of tunnels, but all devices must be the same model and have the Point-to-Point setting. All devices in the same community should have the same community size.

Note that the maximum number of devices (tunnels) supported is reduced by one for each tunnel between the following types of devices:

- WX and WXC devices
- WX or WXC devices that have different versions of WXOS
- WX or WXC devices that have different community sizes

In effect, tunnels between the above pairs of devices must be counted as two tunnels when calculating the maximum number of devices supported.

Table 3: Community Size by Device Type

Device	Small	Large	Point-to-Point
WX 15 (must be a spoke)	Up to 3	Up to 7	Up to 2
WX 20	Up to 4	Up to 10	Up to 2
WX 60	Up to 50	Up to 110	Up to 2
WXC 250	Up to 2	Up to 10	Up to 2
WXC 500	Up to 10	Up to 50	Up to 2
WXC 590	Up to 60	Up to 140	Up to 2
WX 100 (no clients)	Up to 48	Up to 105	Up to 2
WX 100 (with clients)	When a WX 100 has one or more client devices (stack configuration), the client model type is detected, and the community sizes 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 has one WXC 500 client, the large community size is displayed as 300 devices, which is the number of devices supported by six WXC 500 clients (6 x 50).		

To review or change the topology settings:

1. In the Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Topology**.

Figure 54: Reviewing and Changing the Topology Settings

The screenshot shows the Juniper WX 53/22-Carson web interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The left-hand navigation frame has 'Device Setup' expanded, with 'Advanced' selected and 'Topology' highlighted. The main content area is titled 'Topology' and contains the following text: 'Information from this step is used to automate the initial formation of service tunnels between this device and other WX devices in the community. It is still possible to create or delete service tunnels manually if necessary. Select the option below that most accurately describes the topology of your WX community. Then select a community size.' There are four radio button options: 'Hub' (selected), 'Spoke', 'Mesh', and 'Point-to-Point'. Each option has a brief description. Below the options is a 'Community Size' dropdown menu set to 'Small: Up to 32 devices'. At the bottom are 'Submit' and 'Reset' buttons.

2. Select one of the following topology settings:

Hub	A hub attempts to form tunnels with all devices in the community. Select the range of devices in the community (refer to “Selecting the Community Size” on page 110). If a community has multiple hubs, each hub should specify the same community size.
Spoke	By default, a spoke attempts to form tunnels only with devices that are designated as hubs. To enable tunnels between spoke devices, refer to “Configuring Endpoints for Compression” on page 145. Note that a WX 15 must be a spoke, but a WX 100 can never be a spoke (in standalone or stack mode).
Mesh	A mesh device attempts to form tunnels with all other devices in the community. Select the range of devices in the community (refer to “Selecting the Community Size” on page 110). Select the same community size on each mesh device.
Point-to-Point	Same as mesh, except that the community is limited to two devices of the same WX or WXC model, with the same version of WXOS, and with both devices set to the Point-to-Point topology. Typically used to maximize throughput between two data centers.

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.

Using Source/Destination Filters

You can enable or disable data compression 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 compression. A source/destination filter applies to all application traffic sent from the LAN to the WAN. To enable or disable data compression by application, refer to “Compressing Traffic by Application” on page 151. The source/destination filter is applied before the application filter, and is more efficient.

For example, to disable data compression for all traffic from a local subnet, create a “Do not optimize” entry and specify the subnet as the source and enter an asterisk (*) as the destination. To disable data compression for all traffic sent to the subnet by all WX and WXC devices, disable the advertisement of the subnet (refer to “Advertising Compression Subnets” on page 148).

Note the following:

- If you disable data compression 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 Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Source/Destination Filter**.

Figure 55: Using Source/Destination Filters

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The left-hand navigation frame is expanded to 'Advanced', showing options like 'Topology', 'Source/Destination Filter' (highlighted), 'ARP', 'Prime Time', 'Packet Interception', 'WAN Performance Monitor', 'Event Definitions', and 'Multi-Path'. The main content area is titled 'Source/Destination Filter'. It has three radio buttons: 'Off (default)' (selected), 'Compress data between the following source/destination pairs ONLY', and 'DO NOT compress data between the following source/destination pairs'. Below these are two identical sections for defining source and destination subnets. Each section has input fields for 'Source', 'Destination', and a 'Bidirectional' checkbox. The first section shows 'No Source and Destination subnets defined'. The second section has empty input fields. A note below the second section states: 'Click on "Submit" button to add a new source/destination pair. Enter IP address or address/subnet. Enter asterisk (*) to indicate that source or destination can be ANY address. Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0'. At the bottom are 'Submit' and 'Reset' buttons.

2. Select the type of source/destination filter you want to create.

- **Off (default).** Data is compressed for all eligible application traffic from all local routes to all remote routes advertised by other WX and WXC devices.
- **Compress data between the following source/destination pairs ONLY.** Data is compressed only for the specified source and destination pairs. Specify at least one address pair.
- **DO NOT compress data between the following source/destination pairs.** All data is compressed, except for traffic between the specified source and destination pairs.

Note that the excluded 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. For more information about outbound QoS, refer to “Understanding Outbound Bandwidth Management” on page 168.

3. Specify the following information:

Source	Enter a source IP address or subnet. The general format is: address/subnetmask The default subnet mask is “255.255.255.255”. An asterisk (*) with no subnet mask indicates any source IP address.
Destination	Enter a destination IP address or subnet (same format as the source address). An asterisk (*) indicates any destination IP address.
Bidirectional	Select the check box to include traffic sent from the destination to the source. This option is particularly useful for creating “do not optimize” lists in Demo Mode. In Demo Mode, you should exclude all traffic sent to the subnet where the WX device is installed. For more information about Demo Mode, refer to “Demo Mode” on page 453.

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:

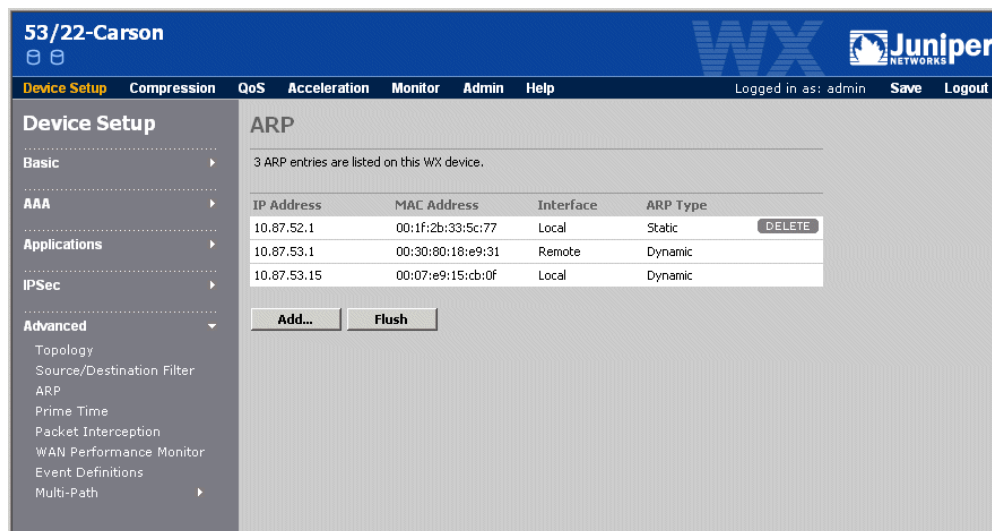
- Determine whether the gateway for a route is on the Local or Remote interface
- Discover the hardware (MAC) addresses of devices that are directly addressable on the Local and Remote interfaces

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 Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **ARP**.

Figure 56: 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 compression 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 Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Prime Time**.

Figure 57: Defining the Prime Time

The screenshot shows the Juniper Device Setup interface for the device 53/22-Carson. The left-hand navigation pane is expanded to the **Advanced** section, where **Prime Time** is selected. The main content area is titled **Prime Time** and contains the following elements:


- A descriptive text block: "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."
- A checkbox labeled **Use Prime Time**, which is currently unchecked.
- Time selection fields:
 - Hours:** Two dropdown menus labeled "From" and "To", both set to 12 AM.
- Day selection fields:
 - Days:** A row of seven checkboxes for Sun, Mon, Tue, Wed, Thu, Fri, and Sat. All checkboxes are checked.
- Two buttons at the bottom: **Submit** and **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 58). 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 58: Off-Path Deployment



NOTE: In off-path deployments, inbound QoS is not supported, and outbound QoS is limited to the WAN traffic that is routed through the WX. Also, an off-path device cannot have a default decompressor (compression will fail), but it can serve as a default decompressor for remote WX endpoints deployed in the data path.

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.

- “Methods of Packet Interception” in the next section
- “Configuring Packet Interception for Off-Path Devices” on page 118
- “RIP Router/Switch Configuration Commands” on page 120
- “WCCP Router Configuration Commands” on page 123
- “External Policy-Based Router Commands” on page 127
- “Alternatives to Packet Interception” on page 127

Methods of Packet Interception

In an off-path deployment, the traffic to be compressed 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 compression subnets advertised by the remote WX devices in the community. Note the following:

- If a remote WX device advertises its own subnet for compression, 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 compression 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 compressed and uncompressed traffic is sent through the tunnel.

To configure a router to use RIP routes, refer to the sample router commands in “RIP Router/Switch Configuration Commands” on page 120.

WCCP

The Cisco Web Cache Communication Protocol (WCCP), which was originally developed to redirect HTTP traffic to Web caches, can be used to redirect any traffic (by protocol) from the router to an off-path WX or a group of WX devices. The router must support WCCP version 2. Note the following:

- The WX accepts any combination of GRE and Layer 2 (L2) encapsulation for forwarded (traffic to be tunneled) and return traffic (passthrough traffic). L2 takes precedence if offered by the router, provided the WX is directly connected to the router at Layer 2. Passthrough traffic is returned to the router as GRE, all other traffic is encapsulated by the WX in a service tunnel.

Note that L2 redirection provides much higher performance than GRE, but is supported only on a selected set of Cisco equipment (such as Catalyst 65xx or Catalyst 45xx).

- WCCPv2 multicast groups are supported, so that in high-availability environments you can define service groups where one or more routers can load-balance traffic across up to four off-path WX devices.
- Redirecting traffic through WCCPv2 can be expensive in terms of router CPU time (particularly when GRE encapsulation is used). To avoid encapsulating passthrough traffic, creating ACLs to bypass WCCP redirection is recommended.

To configure a router to use WCCP, refer to the sample router commands in “WCCP Router Configuration Commands” on page 123.



IMPORTANT: When redirecting traffic to multiple WXs in a service group, you can apply TCP Acceleration, and all services that depend on TCP Acceleration, such as NSC and Application Flow Acceleration, only if you define a cluster for the group (refer to “TCP Acceleration Clusters” on page 330).

External

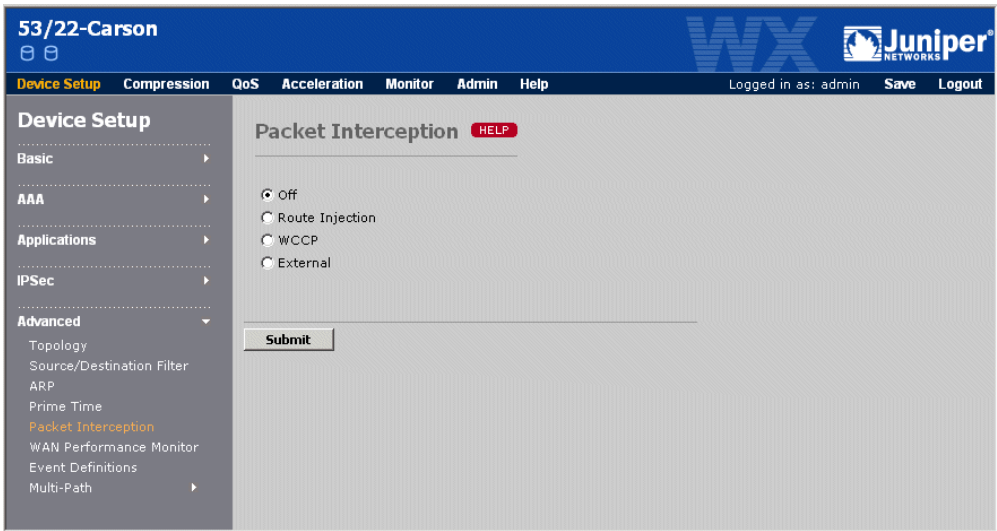
The WAN edge router is configured to route traffic to the off-path device. The off-path device should be connected directly to the router, and must be the only device on the port. You can also connect the off-path device to a dedicated VLAN on a Layer 3 switch. Refer to the sample router commands in “External Policy-Based Router Commands” on page 127.

Configuring Packet Interception for Off-Path Devices

To configure packet interception for an off-path device:

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

Figure 59: 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.

- **Route Injection.** To use RIPv2 for packet interception, click Route Injection, and specify the following:

Authentication Type	If the WAN edge router uses RIP authentication, click Password and enter the RIP password. This is the same password used to discover dynamic routes (refer to “Enabling RIP and OSPF Support” on page 76).
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 357.

- **WCCP.** To use WCCP for packet interception, click **WCCP**, and specify the following:

Address	<p>Enter the IP address of the WAN edge router (the router must support WCCP version 2), or the multicast address of a service group defined on the router. The multicast address must have the form “225.1.1.x” (“225.0.0.x” addresses are not supported).</p> <p>Note that multicast addresses 225.x.x.x through 238.x.x.x are recommended. IGMP snooping may have to be disabled on the Cisco device if addresses 224.x.x.x are used.</p> <p>All WX devices in the same service group (up to four) must specify the same multicast address, and must match the multicast address specified by the router. For load balancing to be effective, verify that tunnels exist between the members of the group and with the appropriate remote WXs.</p> <p>When redirecting traffic to multiple WXs in a service group, you can apply TCP Acceleration, and all services that depend on TCP Acceleration, such as NSC and Application Flow Acceleration, only if all the WXs in the service group belong to a cluster (refer to “TCP Acceleration Clusters” on page 330).</p>
WCCP Priority	<p>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.</p> <p>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 WX, specify a lower value to avoid diverting traffic from the Web cache.</p>
WCCP Auth. Password	<p>If the WAN edge router uses WCCP authentication, enter the WCCP password specified on the router.</p>

Specify the following for each service (up to five):

IP Protocol	<p>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:</p> <p>http://www.iana.org/assignments/protocol-numbers</p>
WCCP Service ID	<p>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.</p> <p>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.</p>

- **External.** 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 127.

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

- Review the compression subnets and be sure to advertise only the subnets on the LAN side of the off-path device (refer to “Advertising Compression Subnets” on page 148). Since only the Local interface is connected to the network, the WX cannot distinguish between LAN- and WAN-side subnets.



CAUTION: If you use RIP for packet interception, and multiple remote WX devices are installed on the same subnet, disable advertisement of the local subnet on all (or all but one) of the remote WXs. Otherwise, the off-path device cannot carve out the remote device addresses, and all traffic sent to them is returned by the router.

- 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 60). 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 60: Off-Path WX Device Connected to a Layer 3 Switch



- Enable RIP version 2:

```
router rip
version 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 on your switch, router, or security appliance (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 61, 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 decompressor.

Figure 61: 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 on your switch, router, or security appliance (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

Sample router commands are shown below for unicast and multicast configurations of WCCP. The actual commands will vary, depending on the network's topology and the type of traffic to be redirected. For more information about WCCP, go to <http://www.cisco.com/univercd/home/home.htm> and search for "wccp".


Note the following:

- **Router ID.** When using the WCCP multicast group function, make sure that the loopback address that the Cisco device chooses as the source address for GRE traffic is active and reachable. A known Cisco WCCP bug is that a configured, but down, IP address is sometimes used as the source address.
- **Enable Multicast Routing.** If using WCCPv2 multicast groups, you may need to enable multicast routing on the physical interface of the WCCP-enabled router (such as with "ip pim dense-mode"), and on any router between the WX and the WCCP-enabled router to ensure that multicast packets reach the WX. On a Catalyst 65XX, you must enable "ip pim dense-mode" on the connecting interface, even if the WX is directly connected.
- **L2 encapsulation.** For the highest level of performance, use Cisco models that can support L2 redirection with WCCPv2. The WX to Cisco connectivity must be Layer 2 for L2 redirection to be negotiated.
- **Miscellaneous.** Other Cisco caveats and testing notes:
 - On Cisco branch level routers, IOS versions 12.3(14) and higher are recommended.
 - On Catalyst 65xx/75xx, do not use CatOS.
 - Catalyst 65xx/75xx with the SUP 1, IOS 12.1(27) supports only two WX devices in a multicast group.
 - Catalyst 3550 has only limited support for WCCP and cannot be used to redirect TCP traffic to off-path WX devices.
- **Useful commands.** On the WX, use "show packet-interception". On the Cisco device, use "show ip wccp", "show ip wccp <service group> [view | detail]". Helpful debugging commands on the Cisco device include "debug ip wccp [events | packets]". Performing packet captures on UDP port 2048 (WCCP) is also beneficial.

Unicast Example

The following commands provide an example of how to configure WCCP on a Cisco router for a single off-path WX device, as shown in Figure 62.

Figure 62: Unicast Example for One Off-Path WX 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
```

Verify that Cisco Express Forwarding is enabled:

```
ip cef
```




NOTE: If you define a service ID on the router, but omit the redirect commands, no traffic is redirected to the WX, but entering a “show packet-interception” command on the WX will indicate the service is connected.

Multicast Example for a Cisco Branch Router

The following configuration is for a service group with two WX devices connected to a Cisco 3640 router running IOS 12.3(1). The key commands are highlighted. GRE encapsulation is automatically negotiated to forward traffic to the WX devices (L2 forwarding is not available on the 3640). The multicast address (225.1.1.1) and service IDs (80 and 90) must match those defined on WX1 and WX2. Because WX1 and WX2 are separated by an L3 boundary, multicast routing must be enabled.

Figure 63: Multicast Example for the Cisco 3640




```

version 12.3
ip wccp check services all
ip wccp 80 group-address 225.1.1.1
ip wccp 90 group-address 225.1.1.1
!
ip cef
!
ip multicast-routing
!
interface FastEthernet0/0
ip address 10.88.18.2 255.255.255.0
ip wccp redirect exclude in
ip wccp 80 group-listen
ip wccp 90 group-listen
no ip mroute-cache
!
interface FastEthernet2/0
ip address 10.88.22.1 255.255.255.0
ip wccp 80 redirect in
ip wccp 90 redirect in
!
interface FastEthernet1/0
ip address 10.88.23.1 255.255.255.0
ip wccp redirect exclude in
ip wccp 80 group-listen
ip wccp 90 group-listen
no ip mroute-cache
!
end
  
```

Multicast Example for the Catalyst 6509

The following configuration is for a service group with four WX devices connected to a Cisco Catalyst 6509 with a SUP 720 running IOS 12.2(18)SXF5, Release Software (fc3). The key commands are highlighted. Since L2 is always used to forward traffic to the WX devices, all WXs in the service group must be on a VLAN. The multicast address (225.1.1.100) and service IDs (51 and 55) must match those defined on WX1 through WX4. Note that even though the WXs and Catalyst 65xx are directly connected at Layer 2, multicast routing is still enabled.

Figure 64: Multicast Example for the Cisco Catalyst 6509



```

version 12.1
ip wccp 51 group-address 225.1.1.100
ip wccp 55 group-address 225.1.1.100
!
ip multicast-routing
!
interface Vlan50
ip address 10.87.105.254 255.255.255.0
no ip redirects
ip wccp 51 redirect out
ip wccp 55 redirect out
no mls ip
!
interface Vlan10
ip address 10.87.119.254 255.255.255.0
no ip redirects
ip wccp 51 group-listen
ip wccp 55 group-listen
ip pim dense-mode
no ip route-cache
no ip mroute-cache
!
interface Vlan20
ip address 10.87.120.254 255.255.255.0
no ip redirects
ip wccp 51 group-listen
ip wccp 55 group-listen
ip pim dense-mode
!
interface Vlan30
ip address 10.87.121.254 255.255.255.0
no ip redirects
ip wccp 51 group-listen
ip wccp 55 group-listen
!

```

```

interface Vlan40
ip address 10.87.122.254 255.255.255.0
ip wccp 51 group-listen
ip wccp 55 group-listen
ip pim dense-mode

```

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 65, 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 65: 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 66 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 66: 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 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 67, 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 67: 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 131
- “Defining Multi-Path Templates” on page 133

- “Defining Multi-Path Endpoints” on page 135
- “Configuring Routers to Support Multi-Path” on page 137

Procedure for Configuring Multi-Path

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

1. Verify that data compression is enabled in both directions between the two devices (refer to “Configuring Endpoints for Compression” on page 145).
2. Verify that the appropriate traffic classes are defined (refer to “Assigning Applications to Traffic Classes” on page 104).
3. Enable the multi-path feature and specify a secondary IP address (refer to “Enabling Multi-Path and Defining Marking Methods” on page 131).
4. 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 133).
5. 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 135).
6. If necessary, configure the WAN router to route traffic to the appropriate path (refer to “Configuring Routers to Support Multi-Path” on page 137).
7. Optionally, enable encryption for both paths or just the less-secure path (refer to “Configuring IP Security (IPSec) and SSL Optimization” on page 227).

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 Device Setup page, click **Advanced** in the left-hand navigation frame, click **Multi-Path**, and then click **Start/Stop**.

Figure 68: Multi-Path Start/Stop Page

The screenshot shows the Juniper Multi-Path Start/Stop configuration page. The left-hand navigation frame is expanded to 'Advanced' and 'Multi-Path'. The 'Multi-Path' section is active, showing the 'Multi-Path' toggle set to 'Enabled'. The 'Secondary IP Address' field is set to '0.0.0.0'. The 'Supplemental Marking Methods' section includes 'Gateway IP' with 'Primary' and 'Secondary' fields, and 'IP Precedence' and 'DSCP Primary' with 'Secondary' fields. There are 'Submit' and 'Reset' buttons at the bottom. The page also includes a 'HELP' button and a 'Logged in as: admin' status bar.

2. Specify the following information:

Multi-Path	Select Enabled to activate the multiple-path feature on this device.
Secondary IP Address	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.

Unless the WAN routers for the primary and secondary paths are also on this subnet (see **Gateway IP** below), the default gateway must be configured to route traffic with this source address to the appropriate WAN link (refer to “Configuring Routers to Support Multi-Path” on page 137).

NOTE: If you enter an address assigned to another device, the path will remain inactive (refer to “Defining Multi-Path Endpoints” on page 135). If you must change the address, do the following:


- Enter the new address and click **Submit**.
- Disable Multi-Path and click **Submit**, and then enable Multi-Path. The configuration for the old secondary path is disabled, including settings for QoS, acceleration, and IPsec.

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 69), 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 IP Precedence or DSCP and enter a ToS IP precedence value (0 to 7) or DSCP value (0 to 63) for packets sent on the primary and/or secondary paths.</p> <p>These values override the IP precedence or DSCP settings for:</p> <ul style="list-style-type: none"> ■ Outbound QoS (refer to “Changing Outbound ToS/DSCP Values” on page 196) ■ WX control packets (refer to “configure reduction” on page 371) <p>The multi-path DSCP values also override ToS marking for router-based balancing (refer to “configure route” on page 384).</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 69: Multi-Path with Primary and Secondary Gateways



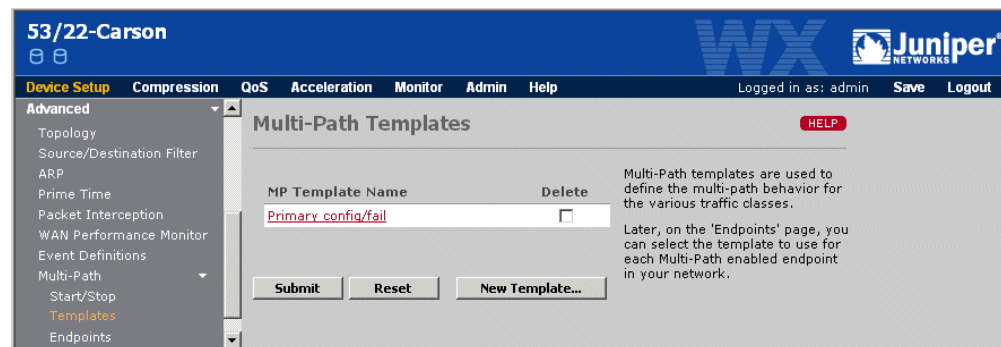
Defining Multi-Path Templates

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 135.

To define multi-path templates:

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

Figure 70: Defining Multi-Path Templates

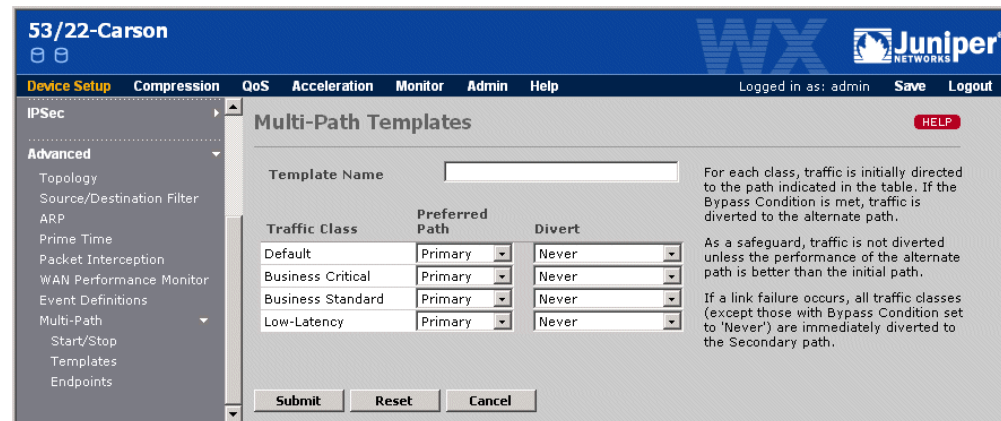


From the Multi-Path Templates page, you can:

- Add a new template, as described in Step 2.
- 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 71: Defining a New Multi-Path Template



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 104).	
Preferred Path	Select Primary or Secondary 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 service tunnel for the preferred path goes down and the tunnel for the alternate path is active.</p> <p>Congestion/Failure. The traffic class is diverted to the alternate path if the loss or latency threshold is exceeded on the preferred path or the service 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 Congestion/Failure is selected for any traffic class, probe packets are sent to the remote devices to measure the loss and latency of each path. To specify a latency threshold for each remote device, refer to “Defining Multi-Path Endpoints” on page 135. 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 352).</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 Device Setup page, click **Advanced** in the left-hand navigation frame, click **Multi-Path**, and then click **Endpoints**.

Figure 72: Defining Multi-Path Endpoints

53/22-Carson

Device Setup Compression QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
- Applications
- IPSec
- Advanced**
 - Topology
 - Source/Destination Filter
 - ARP
 - Prime Time
 - Packet Interception
 - WAN Performance Monitor
 - Event Definitions
 - 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 WX 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 Compression 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 Compression/Endpoints page and enable Compression 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 131).



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>NOTE: If you set the threshold too low, minor fluctuations in latency may cause constant switching between paths.</p> <p>By default, a probe tests the path 12 times per minute. 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 352.</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 133.</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 131). 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:

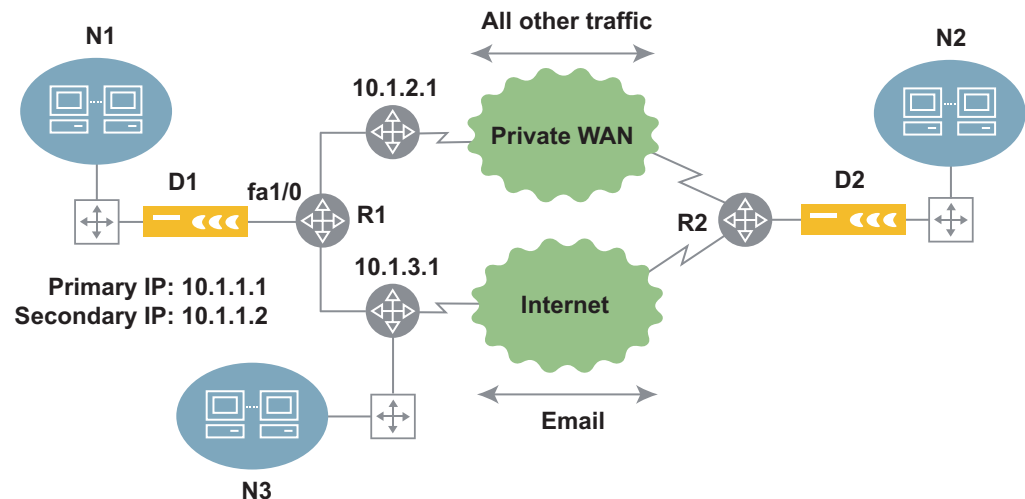
Icon	Description
	The service tunnel is up and the path's loss and latency are below the specified thresholds.
	<p>Connection or performance problem. Move the cursor over the icon to see which one of the following conditions applies.</p> <ul style="list-style-type: none"> ■ No secondary IP address for the remote endpoint (the address specified may belong to another device) ■ Outbound service tunnel is down ■ Loss threshold exceeded ■ Latency threshold exceeded <p>Note that when loss or latency thresholds are exceeded, traffic is switched to the alternate path only if the alternate service tunnel is up and the loss and latency are below the specified thresholds. If a service tunnel is down, traffic is switched to the alternate path regardless of the alternate's performance (provided the alternate service 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 73. A similar configuration is needed for R2.

Figure 73: 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 73, 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 249). 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 compression 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 129).

To enable WAN performance monitoring:

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

Figure 74: Configuring WAN Performance Monitoring

53/22-Carson

Device Setup Compression QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Device Setup

Basic

AAA

Applications

IPSec

Advanced

Topology

Source/Destination Filter

ARP

Prime Time

Packet Interception

WAN Performance Monitor

Event Definitions

Multi-Path

WAN Performance Monitoring

Find: GO HELP

☐ Enable WAN Performance Monitoring for checked endpoints

Device Name	Endpoint	Latency Threshold (msec)
<input type="checkbox"/> 52/22-CARSON	10.87.52.22	5000
<input type="checkbox"/> 54/22-SM250	10.87.54.22	5000
<input type="checkbox"/> 55/22-SR100	10.87.55.22	5000
<input type="checkbox"/> 5722/SR-50A	10.87.57.22	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	<p>Enter the round-trip time (RTT) threshold in milliseconds (20 to 5000). Traps, syslog entries, and report events are generated when the threshold is exceeded, and again when latency drops below the threshold.</p> <p>By default, a probe tests the path 12 times per minute. Traps are generated when the median latency exceeds the threshold for four consecutive minutes or if two or more probes are lost per minute for four consecutive minutes. To change these settings, refer to “configure wan-performance-monitor” on page 397.</p> <p>Note that availability on the WAN Performance report is measured as the percentage of minutes for which at least one probe was acknowledged.</p>
-------------------	--

5. Click **Submit** to 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 Events

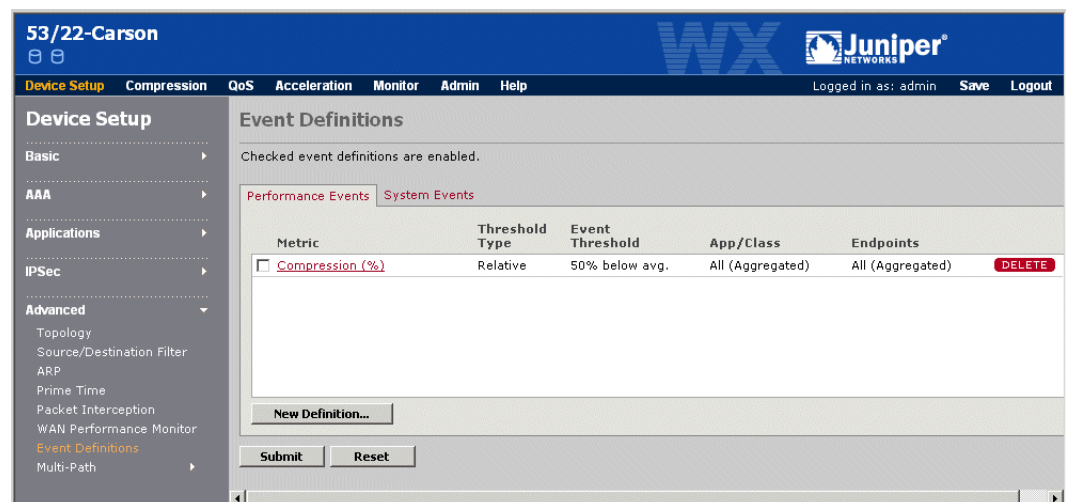
You can configure performance thresholds for compression, acceleration, throughput, and dropped traffic so that events are triggered when the average performance for the previous hour or day exceeds (or drops below) the specified threshold. You can also enable or disable the generation of SNMP traps and syslog messages for system events, such as login failures.

To view the performance and system events that have occurred, refer to “Events Summary” on page 280. Performance events are also sent to any SNMP trap destinations and syslog servers that you have defined (refer to “Enabling SNMP” on page 71 and “Enabling Syslog Reporting” on page 72).

To configure events:

1. In the Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Event Definitions**.

Figure 75: Configuring Performance Event Definitions



From the Event Definitions page, you can:

- Add a new performance event definition, as described in Step 2.
- Enable or disable the generation of specific system events, as described in Step 3.
- Change an event definition. Click the event name, change the generation criteria, and click **Submit**.
- Enable or disable an event definition so that it can generate events. To enable a definition, select the check box next to the metric name, and click **Submit**. To disable a definition, clear the check box and click **Submit**.
- Delete an event definition. Click **DELETE** next to the definition.

2. To add a performance event definition:
 - a. Click **New Definition** and specify the following information:

Metric	<p>Select one of the following metrics. You can create multiple event definitions for each metric. Table 4 describes how each metric is calculated.</p> <ul style="list-style-type: none"> ■ Application Acceleration (%) — for CIFS, Exchange, HTTP ■ Bytes Dropped Out (count) ■ Compression (%) ■ Compression Throughput Out (Kbps) ■ Packets Dropped Out (count) ■ QoS Throughput Out (Kbps) ■ TCP Acceleration (%) ■ TCP Acceleration Throughput In (Kbps) ■ WAN Throughput In (Kbps) ■ WAN Throughput Out (Kbps)
Threshold Type	<p>Select Absolute or Relative to indicate whether the event threshold (see below) is an absolute value or relative to the average performance for the past seven days.</p>
Event Threshold	<p>Select Above or Below and enter the threshold value (the selected metric indicates the appropriate units). For example, to generate an event if compression falls below 80 % of the average of the last seven days, select a Relative threshold type, select Below, and enter “80”.</p>
Application/Class	<p>Select a specific application or traffic class to be monitored, or one of the following (all metrics are for applications, except for QoS Throughput, Bytes Dropped, and Packets Dropped):</p> <ul style="list-style-type: none"> ■ All (Aggregated). An event occurs based on the overall performance of all applications or traffic classes. ■ Any. An event occurs if any application or traffic class violates the specified threshold. <p>If you select the Application Acceleration metric, and you select a specific application, be sure to select a CIFS, Exchange, or HTTP application (others applications have no effect).</p>
Destination	<p>Select a specific endpoint to be monitored, or one of the following:</p> <ul style="list-style-type: none"> ■ All (Aggregated). An event occurs based on the overall performance of all remote endpoints. ■ Any. An event occurs if any endpoint violates the specified threshold. ■ Other Traffic. An event occurs if the “other” traffic violates the specified threshold (applies only to the WAN and QoS Throughput, Bytes Dropped, and Packets Dropped). The “other” traffic excludes traffic sent to all WX and customized non-WX endpoints.
Period	<p>Select Hourly or Daily to indicate whether the average performance is evaluated at the end of each hour or once a day at midnight. A new event is triggered for each hour or day that the average performance violates the threshold.</p> <p>Select the Prime Time Only check box to evaluate performance only for prime time days and hours. No events are generated if you select Prime Time Only and prime time is not defined (refer to “Defining the Prime Time” on page 115).</p>

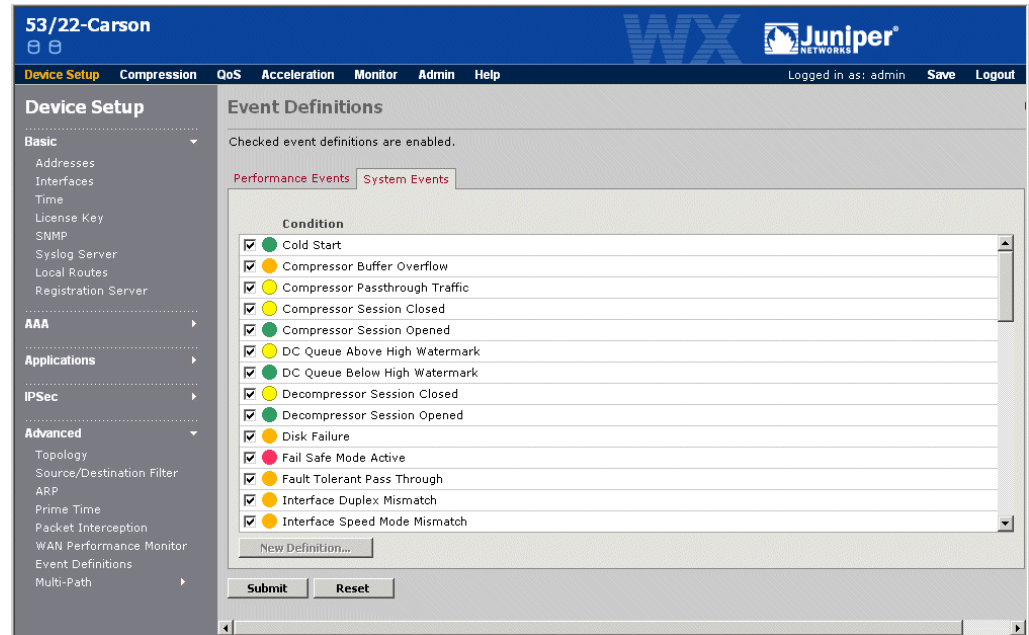
Severity	<p>Select a severity level for events triggered by this definition. The corresponding severity level used on syslog events is shown in parentheses.</p> <ul style="list-style-type: none"> ■ OK (Notice) ■ Warning (Info) ■ Major (Error) ■ Critical (Critical)
Enabled	Select the Yes check box to enable monitoring for this event definition.

Table 4: Performance Metric Calculations

Performance Metric	Calculation
Compression (%)	Average compression for the selected period: $(1.0 - \text{BytesOut} / \text{BytesIn}) * 100$
Compression Throughput Out (Kbps)	Average compression throughput for the selected period: $(\text{BytesOut} * 8) / (\text{Number of seconds in the period})$
TCP Acceleration (%)	Percentage increase in throughput due to TCP Acceleration: $(\text{AccelerationFactor} * 100) - 100$ For example, if the acceleration factor is 3, the acceleration percentage is 200 %.
TCP Acceleration Throughput In (Kbps)	Average throughput for all TCP accelerated sessions in the period. $(\text{BytesIn-over-all-sessions} * 8) / (\text{Active-transaction-time-over-all-sessions})$ Note that this value is averaged over the active TCP sessions, and may be higher than the broader-based compression, QoS, and WAN outbound throughput.
Application Acceleration (%)	Percentage of time saved over the period for CIFS, Exchange, and/or HTTP acceleration: $(\text{TimeWithoutAppAccel} - \text{TimeWithAppAccel}) * 100 / (\text{TimeWithoutAppAccel})$
Wan Throughput In (Kbps)	Average throughput from the WAN over the period: $(\text{BytesFromWAN} * 8) / (\text{Number of seconds in the period})$
Wan Throughput Out (Kbps)	Average throughput to the WAN over the period: $(\text{BytesToWAN} * 8) / (\text{Number of seconds in the period})$
QoS Throughput Out (Kbps)	Average Qos throughput out over the period: $(\text{BytesOut} * 8) / (\text{Number of seconds in the period})$
Bytes Dropped (count)	Absolute number of bytes dropped in the period.
Packets Dropped (count)	Absolute number of packets dropped in the period.

- b. Click **Submit** to activate the changes, or click Cancel to discard them.

3. To enable or disable the generation of specific system events:
 - a. Click **System Events**.

Figure 76: Configuring System Events

- b. To enable or disable a system event, select or clear the check box next to the event name. The colored icons (green, yellow, orange, and red) indicate the severity of each event (OK, Warning, Major, and Critical). For a description of each system event, refer to “SNMP Traps and Syslog Messages” on page 427.
 - c. 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 Multiple Tunnels Between WX 100 Servers

You can increase throughput between two WX 100 stack servers by configuring up to six tunnels between them (one tunnel for each client). Both WX 100 stack servers must have the same number of clients (at least two). The same number of tunnels should be configured on both servers.



NOTE: Disable bandwidth detection to the remote WX 100, as described in “Defining Outbound QoS Endpoints” on page 191 (bandwidth detection reduces throughput for multiple tunnels). Also, configuring Policy-Based Multi-Path for a remote server overrides multiple tunnels. After Multi-Path is configured, adding multiple tunnels has no effect.

To configure multiple tunnels to remote WX 100s:

1. In the Device Setup page, click Advanced in the left-hand navigation frame, and then click WX 100 Tunnels.

Figure 77: Configuring Multiple Tunnels Between WX 100s

SR-10.88.9.100 Number of Active Clients -- 2

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
- Applications
- Encryption
- Advanced
 - Topology
 - Source/Destination Filter
 - ARP
 - Prime Time
 - Packet Interception
 - WAN Performance Monitor
 - Event Definitions
 - Multi-Path
 - WX100 Tunnels**

WX100 Tunnels

This page allows you to specify the maximum number of tunnels that can be formed between this and other remote WX100 endpoints. To add an endpoint, enter the IP address of the remote device and the maximum number of tunnels (2-6), then click **ADD**. To delete an endpoint, click the **DELETE** button to the right of the endpoint. When you are finished editing endpoints, click **Submit**.

IP Address	Maximum Tunnels	
10.10.22.33	2	DELETE
		ADD

Submit Reset

2. To establish multiple outbound tunnels to one or more remote WX 100 stack servers:
 - a. Enter the remote IP address and the number of tunnels (2 to 6). The local and remote servers must have the same number of clients as the number of specified tunnels.
 - b. Click Add.
3. To delete multiple tunnels to a remote server, click **DELETE** next to the server entry. Note that one tunnel to the remote server will be retained.
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.

Chapter 5

Configuring Compression Policies

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

- “Configuring Basic Compression Policies” in the next section
- “Configuring Advanced Compression Policies” on page 153

Configuring Basic Compression Policies

The following topics describe how to configure basic compression policies:

- “Configuring Endpoints for Compression” in the next section
- “Advertising Compression Subnets” on page 148
- “Configuring Network Sequence Caching” on page 150
- “Compressing Traffic by Application” on page 151

Configuring Endpoints for Compression

When you install a new WX device and specify a registration server, the device attempts to form a service 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 compressed data to remote devices, and IN tunnels that convey the compressed data to be decompressed.

Data compression and decompression begins automatically for the compression subnets that are advertised (refer to “Advertising Compression Subnets” on page 148). At any time, you can disable decompression and/or compress data only for specific WX devices in the community.

To configure the endpoints for service tunnels:

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

Figure 78: Configuring Endpoints for Compression

55/22-SR100 Number of Active Clients -- 2

Device Setup **Compression** QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Compression

Basic

- Endpoints
- Compression Subnets
- Network Sequence Cache
- Application Filter

Advanced

Endpoints Find: GO

☒ Enable this device to DECOMPRESS traffic from all other WX devices

☒ Enable this device to COMPRESS traffic destined for:

- ☐ ALL discovered WX devices
- ☐ ONLY WX devices designated as hubs
- ☒ ONLY checked WX devices below

Device name	IP address	Duties	Tunnel Status		Description
			OUT	IN	
<input checked="" type="checkbox"/> 52/22-CARSON	10.87.52.22	Hub	1	1	
<input checked="" type="checkbox"/> 53/22-CARSON	10.87.53.22	Hub	2	✗	No request received
<input checked="" type="checkbox"/> 54/22-SM250	10.87.54.22	Hub	1	1	
<input checked="" type="checkbox"/> 56/22-SR100	10.87.56.22	Hub	2	2	

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 Compression tunnel from this device to remote device
- IN Compression tunnel from remote device to this device
- Tunnel established
- Tunnel established by server
- Tunnel established by client
- on port 1
- No tunnel established
- Broken tunnel
- Temporarily Unavailable

- To stop this device from decompressing compressed data, clear the **Enable this device to DECOMPRESS traffic from all other WX devices** check box. All devices in the community will stop compressing data for this device.
- To stop this device from compressing data for other devices, clear the **Enable this device to COMPRESS traffic destined for:** check box. Otherwise, select one of the following options:
 - All discovered WX devices.** Data is compressed for all other WX devices (default).
 - ONLY WX devices designated as hubs.** Data is compressed only for WX devices designated as a hub.
 - ONLY checked WX devices below.** Data is compressed 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
	Hub — The device is designated as a hub in the community. Each device attempts to form a service tunnel with a hub before creating tunnels to other WX devices (refer to “Configuring Topology Settings” on page 108).
	Spoke — The device is designated as a spoke in a Hub and Spoke topology. By default, a spoke compresses and decompresses data only for the hub device(s).
	Mesh — The device is designated as part of a mesh topology.
	Registration Server — The device is the primary registration server for the community.
	Secondary Registration Server — Indicates that this device is the secondary registration server for the community.
	Backup and Backup (Active) — 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 336.
	Tunnel established — A service tunnel exists between this device and the remote device at the specified IP address. On a 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.
	No tunnel established — No service tunnel exists between this device and the remote device due to a policy setting. For example, if you disable data compression to a remote device by clearing the check box next to its IP address, this icon is displayed in the OUT column, and the message “Disallowed by policy” is displayed in the Description column.
	Broken tunnel — No service tunnel exists between this device and the remote device because of a policy setting or an error. If you manually disable data compression to a remote device, the remote device displays this icon in the corresponding IN column and “No request received” in the Description column.
	Temporarily unavailable — The service tunnel is in a transitory state.

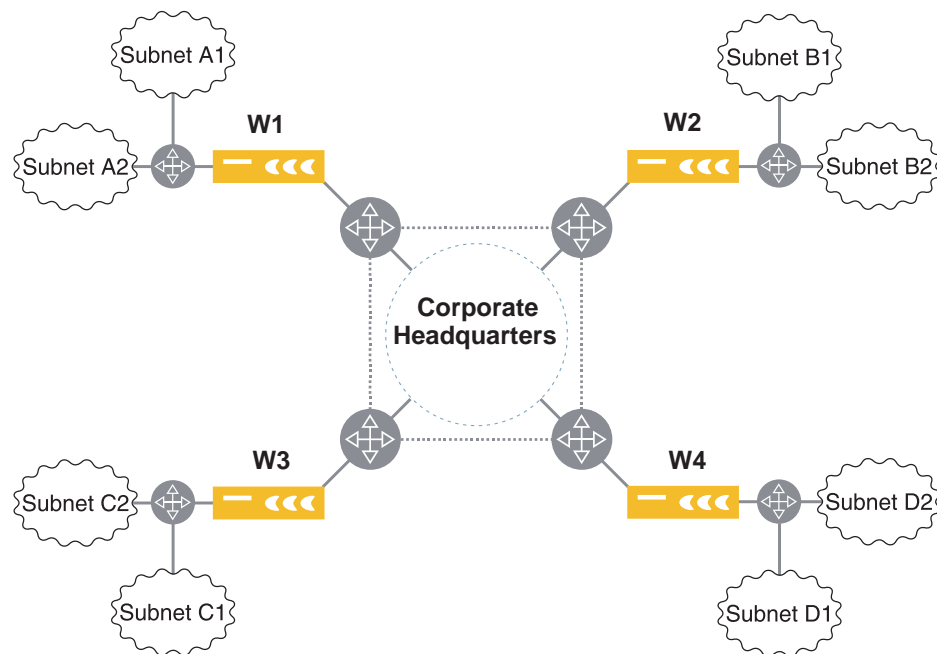
4. Click **Submit** to activate the changes, or click **Reset** to discard them.
5. To retain your changes when the device is restarted, click **Save** in the menu frame.

Advertising Compression Subnets

Compression 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 compress and accelerate traffic sent to the advertised subnets. Initially, the only compression subnet is the subnet where the WX device is installed. To identify more LAN-side subnets, refer to “Configuring Local Routes” on page 73.

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

Figure 79: Selecting Specific Subnets for Data Compression



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

For further control of the traffic being compressed, you can specify application filters, as described in “Compressing Traffic by Application” on page 151, and source/destination filters, as described in “Using Source/Destination Filters” on page 112.

If a WX device has 4000 or more compression 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 compression subnets. For more information, refer to “configure reduction-subnet” on page 378.



NOTE: If a host or gateway in an advertised subnet becomes unreachable, the WX device can dynamically adjust 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 154. To enable or disable the carve-out feature (refer to “configure reduction-subnet” on page 378). The carve-out feature is disabled by default.

To advertise compression subnets:

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

Figure 80: Configuring Compression Subnets

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



NOTE: Normally, compression 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 116) or if the WAN compression subnet option is enabled manually (refer to “configure reduction-subnet” on page 378). 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 compression 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 compression 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 compression is disabled, changes are propagated every hour unless the fetch interval is changed (refer to “Viewing and Fetching Remote Routes” on page 154).

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 compression 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 service 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 compression; 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, service tunnels must exist between them in both directions (refer to “Configuring Endpoints for Compression” on page 145), and TCP Acceleration and outbound QoS must be enabled on both devices (refer to “Enabling Packet Flow Acceleration by Endpoint” on page 208). Applications that are enabled for TCP Acceleration can then be enabled for NSC (refer to “Compressing Traffic by Application” on page 151).

When you install a new WXC, service tunnels, outbound QoS, TCP Acceleration, 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 **Compression** in the menu frame, and then click **Network Sequence Cache** in the left-hand navigation frame.

Figure 81: Configuring Compression Subnets

The screenshot shows the Juniper WX configuration interface for a device named 53/22-Carson. The left-hand navigation frame has 'Compression' selected, and 'Network Sequence Cache' is highlighted. The main content area is titled 'Network Sequence Cache' and contains the following elements:

- A checked checkbox labeled 'Enable Network Sequence Cache when sending traffic to:'.
- Two radio button options:
 - ☒ All NSC-capable WX devices
 - ☐ ONLY NSC-capable WX devices checked below
- A table with three columns: IP address, Device name, and Circuit Speed (Kbps).

IP address	Device name	Circuit Speed (Kbps)
<input type="checkbox"/> 10.87.52.22	52/22-CARSON	10000
<input type="checkbox"/> 10.87.54.22	54/22-SM250	2000
<input type="checkbox"/> 10.87.55.22	55/22-SR100	10000
- Buttons: 'Select All', 'Clear', 'Submit', and 'Reset'.

2. To disable NSC on this device so that standard data compression is used for all remote devices, clear the **Enable Network Sequence Cache...** check box. Otherwise, select one of the following options:

- **All NSC-capable WX devices.** NSC is used for all remote WXC devices in the community (default).
- **ONLY NSC-capable WX 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.

Compressing Traffic by Application

For each application, you can enable or disable data compression and Network Sequence Caching (NSC). To conserve system processing capacity, you should disable compression for applications whose traffic is encrypted or already compressed. However, you must compress 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 compressed, they are monitored automatically. To define additional applications, refer to “Managing Applications” on page 95.

To select applications to be compressed:

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

Figure 82: Selecting Applications for Compression

53/22-Carson

Device Setup **Compression** QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Compression

Basic

- Endpoints
- Compression Subnets
- Network Sequence Cache
- Application Filter

Advanced

Application Filter

Application name	Compress	Network Sequence Cache
AOL	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CIFS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Clearcase	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
CVS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
DNS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Exchange	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Filenet	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
FTP	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Groupwise	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Hostname Resolution	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
HTTP	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
HTTPS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
ICA	<input checked="" type="checkbox"/>	<input type="checkbox"/>
ICMP	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Kerberos	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
LDAP	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Lotus Notes	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Mail	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
MS Streaming	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
XWindows	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Undefined applications	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Only applications checked in the 'Compress' column will be compressed. Unchecked applications will be passed through without compression.

If Network Sequence Cache (NSC) is enabled for traffic destined for NSC-capable WX devices (See the "Network Sequence Cache" page.) then applications checked in the "Network Sequence Cache" column will be compressed using NSC.

Applications for which TCP Acceleration (AFP) has not been enabled cannot be enabled for NSC. This includes 'Undefined Applications'.

Compress All NSC 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:.

Compress	<p>Select the check box next to each application to be compressed. By default, all applications are compressed (except Groupwise, HTTPS, SMTP, SSH, and Traceroute). If an application is not compressed, its traffic passes through the device without compression. To compress all applications, click Compress All.</p> <p>To conserve processing capacity, disable compression for applications whose traffic is encrypted or already compressed. However, you must compress all TCP applications that you want to accelerate (refer to “Accelerating WAN Traffic” on page 203).</p>
NSC	<p>On a WXC device, you can enable Network Sequence Caching (NSC) for compressed 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 150), then NSC is used to compress 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 compression is used for traffic sent to WX devices or to WXC devices where NSC is disabled.</p> <p>To enable NSC for all compressed applications, click NSC All. To use NSC for an application, the application must be enabled for compression and for TCP Acceleration (refer to “Enabling TCP Acceleration by Application” on page 212).</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 Compression Policies

The following topics describe the advanced data compression policies:

- “Viewing and Fetching Remote Routes” in the next section
- “Configuring Tunnel Load Balancing Policies” on page 155
- “Defining Default Decompressors” on page 157
- “Defining Preferred Decompressors” on page 159
- “Configuring Tunnel Mode Settings” on page 160
- “Configuring Pre-Synchronization for Network Sequence Caching” on page 161
- “Configuring Tunnel Switching” on page 163

Viewing and Fetching Remote Routes

Remote routes are the compression subnets advertised by the other WX devices in the community. Each device can compress 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.

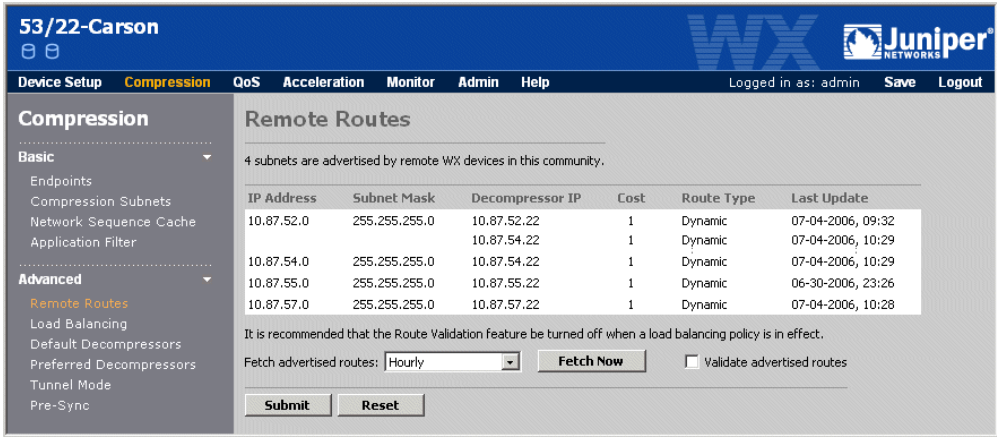


NOTE: The remote routes shown for a device may not match the list of advertised subnets shown on the device. Each device can dynamically adjust its advertised subnets to exclude (carve out) unreachable addresses. To exclude an address from an advertised subnet, multiple smaller subnets are generated, so that one advertised subnet may produce several remote routes. To enable or disable the carve-out feature, refer to “configure reduction-subnet” on page 378.

To view the remote routes:

- 1. Click **Compression** in the menu frame, click **Advanced** in the left-hand navigation frame, and then click **Remote Routes**.

Figure 83: Displaying and Updating Remote Routes



The **Decompressor IP** column shows the address of one or more remote devices that can decompress 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 compression 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 155).

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 decompressing the compressed data. Alternatively, you can specify preferred decompressors, as described in “Defining Preferred Decompressors” on page 159. If neither load balancing nor preferred decompressors are used, the path selection is arbitrary.

For example, in Figure 84, 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 85).

Figure 84: Configuring Tunnel Load Balancing Policies




Figure 85: Remote Routes with Equal Cost Paths

The screenshot shows the Juniper WX configuration interface for device 53/22-Carson. The 'Compression' tab is selected, and the 'Advanced' sub-tab is active, displaying the 'Remote Routes' section. A table lists four subnets advertised by remote WX devices, all with a cost of 1, indicating equal cost paths. An arrow points to the 'Last Update' column with the text 'Common destinations with equal cost paths'.

IP Address	Subnet Mask	Decompressor IP	Cost	Route Type	Last Update
10.87.52.0	255.255.255.0	10.87.52.22	1	Dynamic	07-04-2006, 09:32
10.87.54.0	255.255.255.0	10.87.54.22	1	Dynamic	07-04-2006, 10:29
10.87.55.0	255.255.255.0	10.87.55.22	1	Dynamic	07-04-2006, 10:29
10.87.57.0	255.255.255.0	10.87.57.22	1	Dynamic	06-30-2006, 23:26

Common destinations with equal cost paths

To configure tunnel load balancing policies:

1. Click **Compression** in the menu frame, click **Advanced** in the left-hand navigation frame, and then click **Load Balancing**.

Figure 86: Configuring Tunnel Load Balancing Policies

The screenshot shows the Juniper WX configuration interface for device 53/22-Carson. The 'Compression' tab is selected, and the 'Advanced' sub-tab is active, displaying the 'Load Balancing' section. The interface explains that the selected rule determines how traffic is routed when more than one compression tunnel exists for a given subnet. Four load balancing policies are listed: Off (selected), Per-destination, Per-packet, and Flow based.

The rule selected below determines how traffic is routed when more than one compression tunnel exists for a given subnet. The selected rule also applies to Default Decompressors 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.

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

3. Click **Submit** to 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 Default Decompressors

You can sometimes simplify route administration by designating a WX device as the default decompressor for one or more remote devices. The default decompressor need not discover and advertise all of its local routes because the remote devices automatically compress 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 decompressors, regardless of whether compression is enabled.



NOTE: Default decompressors cannot be used in Demo Mode or for an off-path device that uses RIP for packet interception (compression will fail). Be sure you understand which WX devices will use a default decompressor, and which local routes the default decompressor supports.


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

Note that traffic sent to the default decompressor is not compressed when:

- The sending device has a static or dynamic route to one of the default decompressor's local subnets that the default decompressor 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 "Using Source/Destination Filters" on page 112).

Figure 87 shows a simple example of a remote site with one outbound connection to the corporate network. If D1 is the default decompressor for D2, all traffic that uses the default route on D2 is compressed and sent to D1.

Figure 87: Setting a Default Decompressor



To disable data compression for traffic sent to subnet S4, you can add S4 to the exclusion list on D2. You can specify up to six default decompressors. If you specify more than one default decompressor, the current load balancing policies are applied (refer to “Configuring Tunnel Load Balancing Policies” on page 155).

To define default decompressors:

1. Log in to the device where you want to specify default decompressors.
2. Click **Compression** in the menu frame, click **Advanced** in the left-hand navigation frame, and then click **Default Decompressors**.

Figure 88: Creating a Default Decompressor List

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The 'Compression' section is expanded, showing 'Basic' and 'Advanced' sub-sections. Under 'Advanced', 'Default Decompressors' is selected. The main content area is titled 'Default Decompressors' and contains two text boxes: 'Default Decompressors' and 'Exclude List'. The 'Default Decompressors' box has a text area for entering IP addresses of WX devices, one per line, with a maximum of 6. The 'Exclude List' box has a text area for entering addresses/subnets, one per line, with a maximum of 128. Examples provided are '123.123.123.123' and '123.123.123.1/255.255.255.0'. At the bottom, there are 'Submit' and 'Reset' buttons.

3. In the Default Decompressors box, enter the IP address of up to six default decompressors (one per line). If load balancing is disabled, the precedence of the default decompressors 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 compressed before being sent to the default decompressor. 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 decompressor 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 73.
7. To retain your changes when the device is restarted, click **Save** in the menu frame.

Defining Preferred Decompressors


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 decompressor. Alternatively, you can use load balancing to vary the selected path, as described in “Configuring Tunnel Load Balancing Policies” on page 155. If neither load balancing nor preferred decompressors are used, the path selection is arbitrary.



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

For example, in Figure 89, data from Subnet 1 has two network paths to Subnet 2. If the WX device D1 designates D2 as a preferred decompressor, all compressed data destined to Subnet 2 is sent to D2. If D2 is unavailable, D3 is used.

Figure 89: Designating a Preferred Decompressor

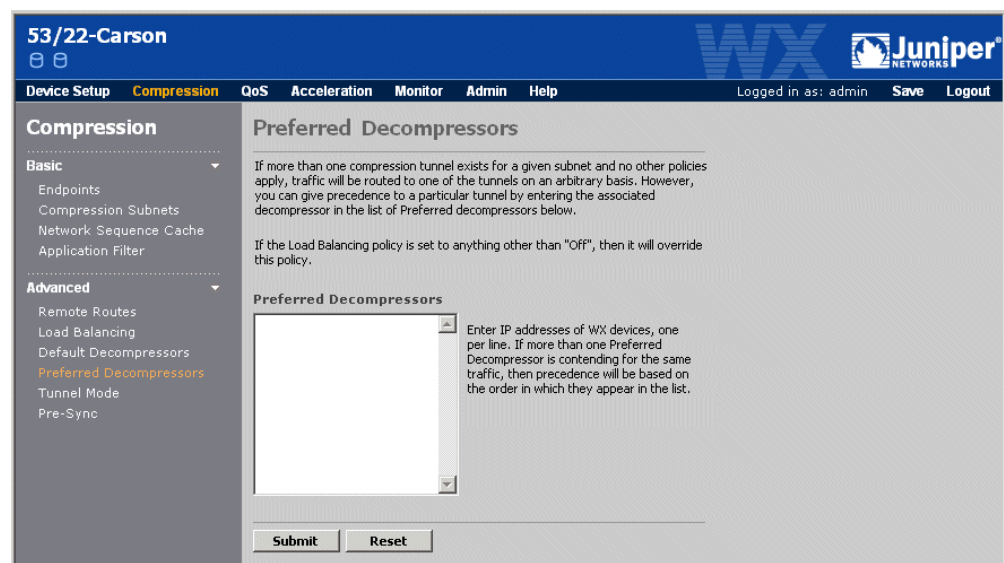


Note that a preferred decompressor is used even for routes that have a lower cost on an alternate WX device.

To create a list of preferred decompressors:

1. Click **Compression** in the menu frame, click **Advanced** in the left-hand navigation frame, and then click **Preferred Decompressors**.

Figure 90: Defining Preferred Decompressors



- 2. Enter the IP address of a remote preferred decompressor. You can specify up to 80 preferred decompressors (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.

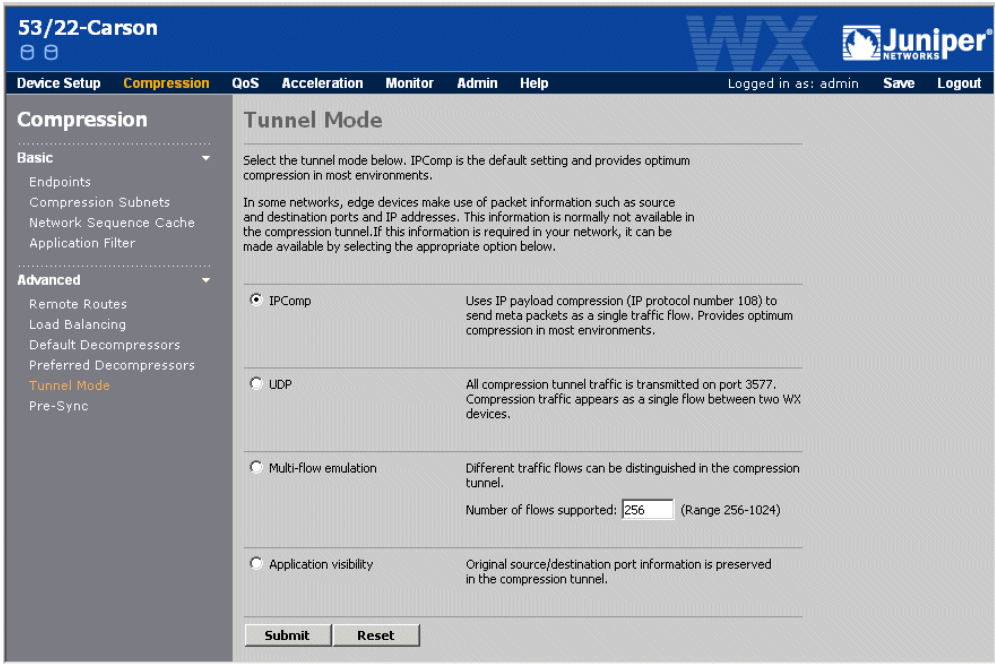
Configuring Tunnel Mode Settings

The tunnel mode determines how a WX device sends compressed traffic to the remote WX devices in the same community. By default, compressed packets are enclosed in meta packets and sent over a service 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 **Compression** in the menu frame, click **Advanced** in the left-hand navigation frame, and then click **Tunnel Mode**.

Figure 91: Configuring Tunnel Mode Settings



- 2. Select one of the following tunnel modes:

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

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 service 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 compression rates for large files, such as database files and software updates.

During pre-synchronization, the repeated patterns in the files are added to the compression dictionaries of the selected devices, so that compression occurs when the first user requests the files. Compression also occurs if the files are sent in the reverse direction.

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 374).



NOTE: If you have WX CMS 5.5 or later, you can use the content distribution feature to manage and schedule all pre-synchronizations from the CMS web interface (refer to the *WX Central Management System (CMS) Administrator's Guide*).

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 compress 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 150.
- NSC must be enabled for the application that users will access to retrieve the files (refer to “Compressing Traffic by Application” on page 151).

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 WXC's without being compressed.

2. Click **Compression** in the menu frame, click **Advanced** in the left-hand navigation frame, and then click **Pre-Sync**.

Figure 92: Configuring Pre-Synchronization for Network Sequence Caching

53/22-Carson

Device Setup **Compression** QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Compression

Basic

- Endpoints
- Compression Subnets
- Network Sequence Cache
- Application Filter

Advanced

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

Pre-Sync

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

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

FTP file pathnames should be of the form *ftp://host-ip-address:user:password/path*.

File Pathname

(Max length 2000)

IP address	Device name
<input type="checkbox"/> 10.87.52.22	52/22-CARSON
<input type="checkbox"/> 10.87.54.22	54/22-SM250
<input type="checkbox"/> 10.87.55.22	55/22-SR100

Select All Clear

Go

3. Specify the FTP server location of each file (one file per line). Each location is limited to 128 characters (2000 characters total). 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.

4. 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.

5. 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 compression (form service tunnels) for a varying number of remote WX devices, depending on the device type. Tunnel switching allows each device to compress data for every other WX device in the network, without having to form a tunnel with each remote device.



NOTE: Tunnel switching supports only data compression and QoS, and cannot be enabled on the same device with CIFS, Exchange, or HTTP acceleration. Also, tunnel switching is not supported on a WX 100 with one or more client devices.

Devices can be deployed hierarchically in either of the following ways:

- Assign devices to separate communities (only devices in the same community can compress and decompress data for each other)
- Apply hub and spoke topology designations to selected devices in the same community (by default, spoke devices compress and decompress 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 that 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 decompressors for traffic leaving the community, you must enable tunnel switching to compress the traffic for the next device in the path.

In Figure 93, tunnel switching is enabled on D2 and D3 in Community 3. These are the common devices that convey compressed traffic between Community 1 and Community 2.

Figure 93: Example of Tunnel Switching Between Communities

When D1 in Figure 93 encounters traffic destined for a subnet advertised by D4, the following processing occurs:

1. D1 cannot match the destination to an advertised compression subnet (D4 is in a separate community), so D1 compresses the traffic and sends it to the default decompressor (D2).
2. D2 decompresses the traffic from D1, matches the destination to a compression 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 decompresses the compressed traffic from D2, matches the destination to a compression 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 decompressor for D4 (and the other devices in Community 2).

Note that Figure 93 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 decompression and recompression on D3.

Procedure for Tunnel Switching Between Communities

To configure tunnel switching between communities:

1. Identify the common devices in each community that convey compressed traffic between communities.
2. On each of the other devices in a community, designate the common device as the default decompressor (refer to “Defining Default Decompressors” on page 157). 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 371). For off-path devices, LAN/WAN checking is disabled by default.
 - Add static routes for the compression subnets to be advertised to the other common devices (refer to “Adding Static Routes” on page 75). For example, on D2 in Figure 93, define static routes for the compression 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 compression subnets on the common devices (refer to “configure reduction-subnet” on page 378). For example, in Figure 93, if the other devices in Community 1 are on the Remote (WAN) side of D2, WAN compression 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 compression 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 compressed traffic between any two spoke devices. By default, spoke devices compress and decompress 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 94 is very similar to Figure 93, 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.

Figure 94: Example of Tunnel Switching Between Spoke Devices

To configure tunnel switching between spoke devices:

1. Identify the hub devices that convey compressed traffic between spokes.
2. On each of the spoke devices, designate the hub device as the default decompressor (refer to “Defining Default Decompressors” on page 157). 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 371). In some routing environments, on each hub you may need to add static routes for the compression 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 compression subnets on the hub devices (refer to “configure reduction-subnet” on page 378). For example, in Figure 93, if the spoke devices for Hub 1 are on the Remote (WAN) side of D2, WAN compression 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 75).

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 168
- “Configuring Outbound QoS Policies” on page 177
- “Configuring Inbound QoS Policies” on page 199
- “Summary of Key Terms” on page 202

Using Outbound QoS to Enhance Performance

Outbound QoS provides two key benefits:

- **Basic bandwidth allocation.** Data compression 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 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 191. Note the following:
 - If you know the local and remote WAN speeds, be sure to adjust them to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 170).
 - If you do not know the WAN speeds, you can enable bandwidth detection for each remote endpoint (refer to “Bandwidth Detection” on page 172).
2. Start outbound QoS using Weighted Fair Queuing (WFQ) or Weighted Strict Priority (WSP), as described in “Starting and Stopping Outbound QoS” on page 198. 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 compressed. 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 169
- “QoS Templates and Endpoints” on page 170
- “WAN Circuit Speeds and Router Overhead” on page 170
- “Dedicated and Oversubscribed WANs” on page 171
- “Bandwidth Detection” on page 172

- “Direct Setup Versus Wizard Configuration Results” on page 174
- “Class Priorities and Excess Bandwidth Allocation” on page 176
- “ToS/DSCP Values” on page 177
- “Unadvertised Subnets” on page 177
- “Summary of Key Terms” on page 202

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 95 shows the default settings for the standard traffic classes created by the Setup Wizard. You can have up to 16 traffic classes.

Figure 95: 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 176).

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 95).

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 115.
- **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 262).

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 and Oversubscribed WANs” on page 171).
- **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 bandwidth detection for that endpoint.



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

If bandwidth detection is NOT enabled, the endpoint WAN circuit speeds must be set 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.



NOTE: The overhead on Asynchronous Transfer Mode (ATM) is typically 15 %.

For an oversubscribed WAN, always set the outbound speed as accurately as possible, even if bandwidth detection 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 and Oversubscribed 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 bandwidth detection is enabled:

- **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 96, 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 96: 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 97, 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 97: Oversubscribed WAN

Bandwidth Detection

The bandwidth detection feature dynamically alters the bandwidth allocation per-endpoint based on the measured real-time available WAN bandwidth. Bandwidth detection is recommended for networks that have variable WAN bandwidths, such as:

- Frame Relay networks that support a sustained CIR and bursts to a peak rate
- MPLS networks, which are inherently "connectionless"

- Shared satellite uplink environments where several routers may share a single satellite connection

Enabling bandwidth detection simplifies the QoS configuration. However, if your environment has large amounts of loss or large amounts of out-of-order traffic, then bandwidth detection alone may not be the best choice. Bandwidth detection responds to loss and out-of-order traffic by reducing the amount of traffic it sends. If you have applications that require a specific amount of bandwidth, then bandwidth detection may not be suitable for your environment.

Bandwidth detection also lets you set minimum speeds on a per-endpoint basis, guaranteeing that a certain amount of bandwidth is always available for those endpoints. For example, if you have a 2 MB circuit with a 1 MB CIR that occasionally bursts to 2 MB, you can set the maximum bandwidth level to 2 MB in QoS, and a minimum speed of 1 MB in bandwidth detection.

In a network with mixed technologies, you might do the following.

- Configure dedicated WANs directly without bandwidth detection
- Configure bandwidth detection with a minimum speed for oversubscribed MPLS networks where you need to guarantee minimum bandwidth levels for each site, but still need to send at the maximum data rate.

Note that WX/WXC devices will not reduce the bandwidth below the configured minimum, even if there is loss or congestion in the network.

- Configure bandwidth detection with no minimum for Internet VPNs with no available traffic rate guarantees.

In the case of loss, you can configure bandwidth detection to be more aggressive (back off less) by changing the bandwidth detection mode, as follows:

1. Go to Admin > Tools > Command line interface.
2. Enter the following commands, and click **Submit**:

```
configure qos outbound set congestion-control-action-on-loss scps-slow-backoff
commit
```

Remember to save your configuration. Since bandwidth detection dynamically adjusts to the available bandwidth, the WAN speed specified for each remote endpoint is not critical. For example, you can enable bandwidth detection 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 170).



NOTE: Bandwidth detection 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 191).

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 5 shows the Wizard and direct setup results when D1 in Figure 97 is configured with two traffic classes and the same guaranteed bandwidths for each endpoint.

Table 5: 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 5 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 98 shows the “Oversubscribed” template specifying the 15 % and 40 % guarantees; Figure 99 shows the guaranteed bandwidths in Kbps displayed on the Outbound QoS Overview page when the template is applied to each endpoint.

Figure 98: 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 99: 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	EDIT Oversubscribed	60	9	24	33
192.168.53.5	EDIT Oversubscribed	500	75	200	275
192.168.52.22	EDIT Oversubscribed	1500	225	600	825
Total			309	824	1133

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 5. 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 100 shows the guaranteed bandwidths in Kbps generated by the Setup Wizard and displayed on the Outbound QoS Overview page.

Figure 100: 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	EDIT Wizard-PrimeTime	60	6	17	23
192.168.53.5	EDIT Wizard-PrimeTime	500	54	145	199
192.168.52.22	EDIT Wizard-PrimeTime	1500	163	436	599
Total			223	598	821

The Wizard adjusts the bandwidths for oversubscribed WANs only when there are multiple remote endpoints. For example, in Figure 97 on page 172, 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 6. The allocation depends on whether the WAN is dedicated or oversubscribed.

Table 6: 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 5 on page 174, 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 196.

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 199).

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 compression for specific destinations, as needed (refer to “Using Source/Destination Filters” on page 112).

Also, if advertised subnets automatically exclude hosts or gateways that become unreachable, traffic to those “carved out” addresses is also attributed to the “Other traffic” endpoint. To enable or disable the carve-out feature, refer to “configure reduction-subnet” on page 378.

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 178
- “Using the Outbound QoS Setup Wizard” on page 179
- “Defining Outbound QoS Settings by Endpoint” on page 186
- “Defining Traffic Classes” on page 188
- “Defining Outbound QoS Templates” on page 189
- “Defining Outbound QoS Endpoints” on page 191
- “Changing Outbound ToS/DSCP Values” on page 196
- “Starting and Stopping Outbound QoS” on page 198

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 compression for specific destinations (refer to “Using Source/Destination Filters” on page 112).
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 188).
 - b. Define one or more templates to specify the priorities and bandwidths for each traffic class (refer to “Defining Outbound QoS Templates” on page 189).
 - 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 170 and “Defining Outbound QoS Endpoints” on page 191).
 - d. Assign a prime-time and nonprime-time template to each endpoint (refer to “Defining Outbound QoS Settings by Endpoint” on page 186).
 - e. Enable QoS and select a queuing model (refer to “Starting and Stopping Outbound QoS” on page 198).
- 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 186).
 - Change traffic class names (refer to “Defining Traffic Classes” on page 188).
 - Add new templates, change a template name, or change just one of the Wizard templates (refer to “Defining Outbound QoS Templates” on page 189)

- Define virtual endpoints or exclude address or subnet pairs from bandwidth management (refer to “Defining Outbound QoS Endpoints” on page 191).
- Change the ToS/DSCP values for one or more traffic classes (refer to “Changing Outbound ToS/DSCP Values” on page 196).

Using the Outbound QoS Setup Wizard

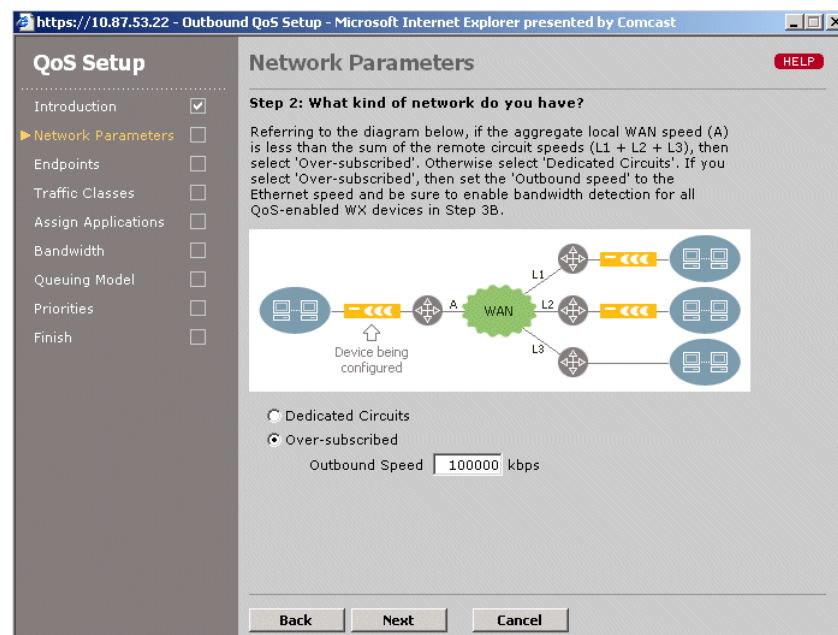
Use the Setup Wizard the first time you define outbound QoS policies. The Setup Wizard creates two identical templates and assigns them to the selected endpoints:

- **Wizard-PrimeTime.** Applies to the prime time hours (critical business hours). To specify the prime time, refer to “Defining the Prime Time” on page 115.
- **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 262.

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 189.

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.

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 **Outbound Speed** field. Be sure to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 170).

https://10.87.53.22 - Outbound QoS Setup - Microsoft Internet Explorer presented by Comcast

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

Endpoint	Name	Circuit Speed
<input type="checkbox"/> Other traffic		1000000
<input checked="" type="checkbox"/> 10.87.52.22	10.87.52.22	10000
<input type="checkbox"/> 10.87.54.22	54/22-SM250	2000
<input type="checkbox"/> 10.87.55.22	55/22-SR100	10000
<input type="checkbox"/> 10.87.57.22	5722/SR-50A	100000
<input type="checkbox"/> 10.87.58.22		1000

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 acceleration is enabled for an endpoint, outbound QoS cannot be disabled.



CAUTION: If you do not enable Bandwidth Detection (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 170). 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 129.

- 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 WX” 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 191), 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 186).

5. Click **Next**.

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 Bandwidth Detection and indicate the minimum speed (in kbps) for the relevant endpoints. If you don't know the minimum speed, enter '0'.

☒ Enable Bandwidth Detection when sending compressed traffic to:

☒ All QoS-enabled WX devices that also have compression enabled

☐ ONLY checked WX devices below

IP Address	Device Name	Minimum Speed
<input type="checkbox"/> 10.87.54.22	54/22-SM250	<input type="text" value="0"/>
<input type="checkbox"/> 10.87.57.22	5722/SR-50A	<input type="text" value="0"/>
<input type="checkbox"/> 10.87.58.22	10.87.58.22	<input type="text" value="0"/>

6. If the WAN bandwidth to a remote WX device is variable, such as for MPLS, Frame Relay, or shared satellite links, enable Bandwidth Detection for traffic sent to that device.

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

To enable bandwidth detection:

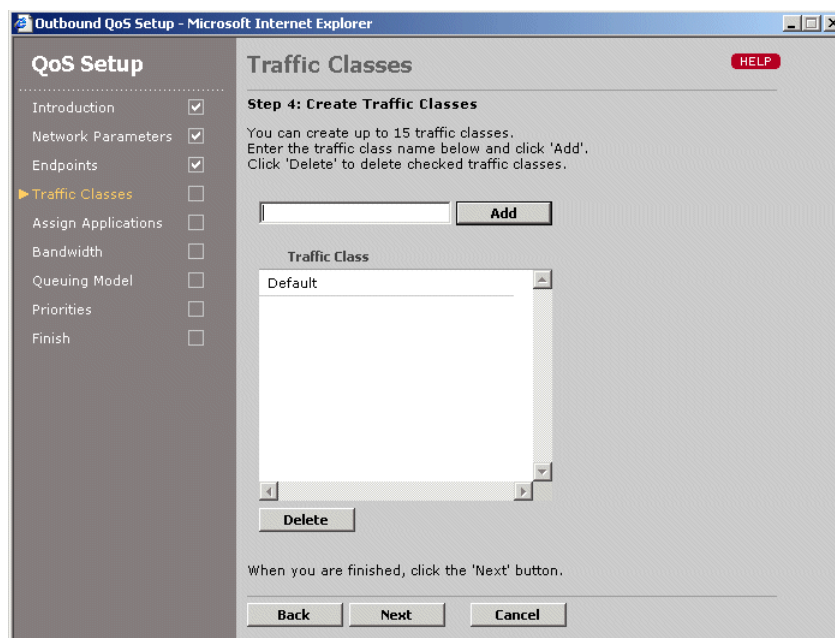
- a. Select Enable Bandwidth Detection and select one of the following options:
 - **All QoS-enabled WX devices.** Applies bandwidth detection to all remote WX devices for which QoS is enabled (default).
 - **ONLY checked WX devices below.** Select the check box for one or more QoS-enabled endpoints.

- 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: Bandwidth Detection manages only traffic sent to other WX endpoints for which service 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 191.

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



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**.

QoS Setup

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

Assign Applications to Classes HELP

Step 5: How important are your applications?

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

Application	Traffic Class
AOL	Default
CIFS	Default
Clearcase	Default
CVS	Default
DNS	Default
Exchange	Default
Filenet	Default
FTP	Default
Groupwise	Default

When you are finished, click the 'Next' button.

Back Next Cancel

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 95. Click **Next**.

QoS Setup

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

Traffic Class Bandwidth HELP

Step 6: How much bandwidth should each traffic class have?

Enter Guaranteed and Maximum Bandwidth limits for the following Traffic Classes.

Traffic Class	Guaranteed Bandwidth	Maximum Bandwidth
Default	0.00 %	100.00 %
Business Critical	40.00 %	100.00 %
Business Standard	20.00 %	100.00 %
Low-Latency	20.00 %	100.00 %
Prohibited	0.00 %	0.00 %

The Guaranteed Bandwidth is the amount of bandwidth that is guaranteed to be available for a traffic class, regardless of the volume of competing traffic from other classes.

The Maximum Bandwidth is the amount of bandwidth that this traffic class will be limited to even if additional bandwidth is available.

Guaranteed and Maximum bandwidths are expressed as a percent of the bottleneck link speed. The total Guaranteed bandwidth across all traffic classes should not exceed 80%.

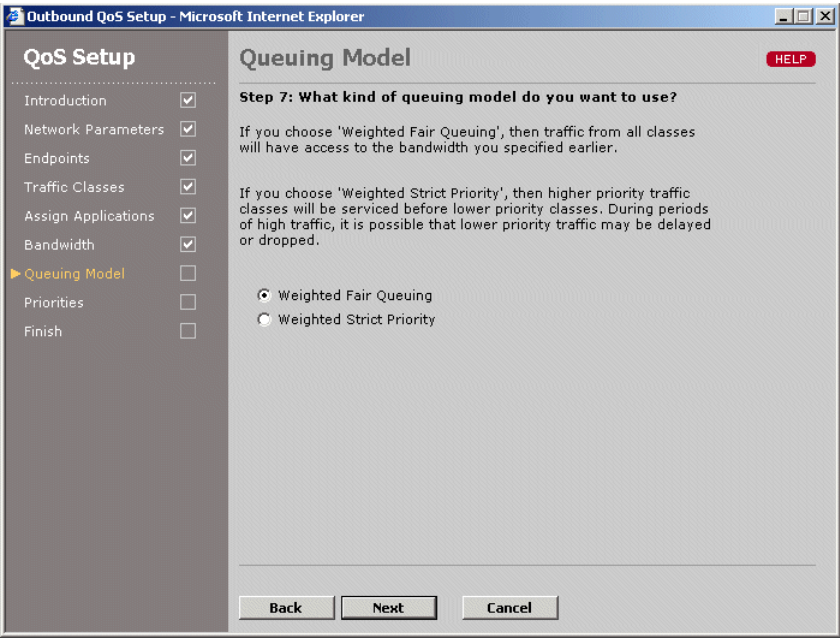
Back Next Cancel

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	<p>Queues are created for each traffic class, and the excess bandwidth is allocated by processing the queues based on their priority and guaranteed bandwidth.</p>
Weighted Strict Priority	<p>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.</p>

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.

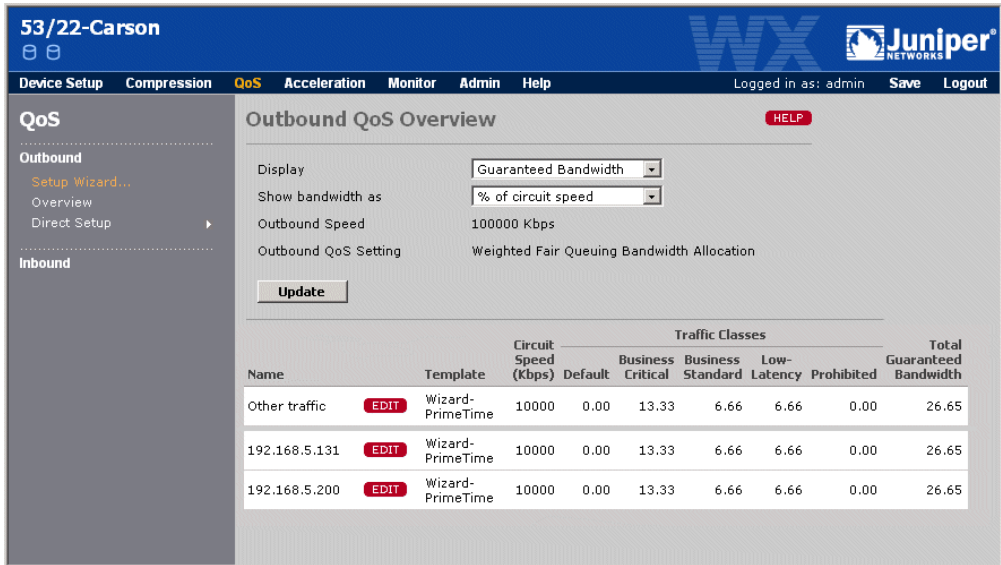
Defining Outbound QoS Settings by Endpoint

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 191.

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 101: Outbound QoS Overview



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 bandwidth detection is enabled.



NOTE: If bandwidth detection 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 115).
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 102: Changing Endpoint Templates or Template Settings

53/22-Carson

Device Setup Compression **QoS** Acceleration Monitor Admin Help

Logged in as: admin Save Logout

QoS

Outbound
Setup Wizard...
Overview
Direct Setup

Inbound

Outbound QoS Overview > Other traffic **HELP**

This page determines bandwidth limits for Outbound QoS to the selected endpoint.

Endpoint: Other traffic
Circuit Speed: 1000000 Kbps

☐ Use QoS template Wizard-PrimeTime
☒ Use custom setting

Show bandwidth as: % of circuit speed

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 %
Total		26.67 %	

Bandwidth limits are stored as a percent of circuit speed. If circuit speed is modified, the bandwidth Kbps values will also change.

Submit Cancel

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 189.
 - 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 179).

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

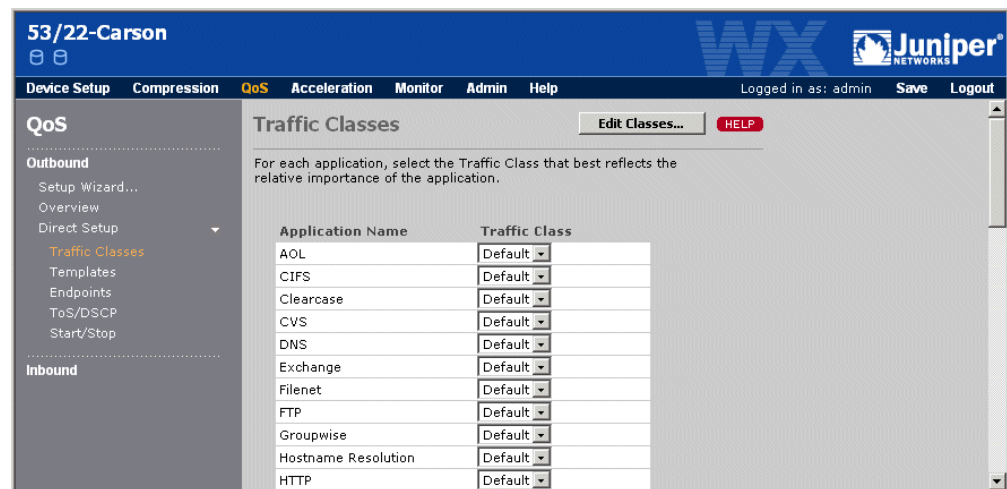
Defining Traffic Classes

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 103: 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.
4. To retain your changes when the device is restarted, click **Save** in the menu frame.

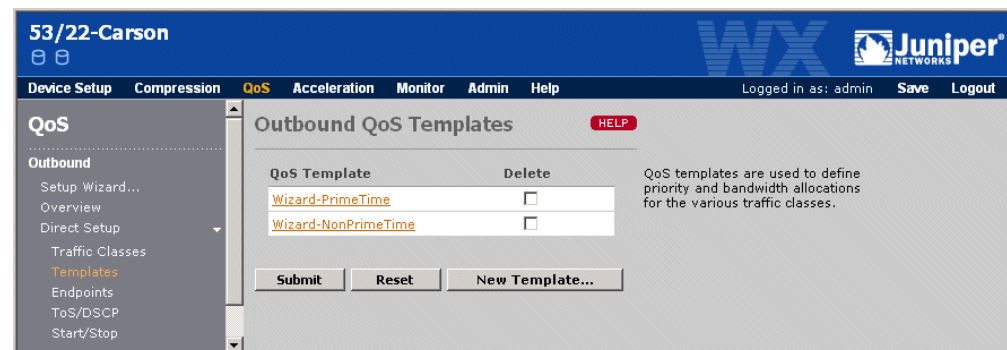
Defining Outbound QoS Templates

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 186.

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 104: Defining Outbound QoS Templates



From the Outbound QoS Templates page, you can:

- Add a new template, as described in Step 2.
- 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:
 - a. Click **New Template**.

Figure 105: Defining a New QoS Template

53/22-Carson

WXJuniper

Device SetupCompressionQoSAccelerationMonitorAdminHelp

Logged in as: adminSaveLogout

QoS

Outbound

Setup Wizard...
Overview
Direct Setup
Traffic Classes
Templates
Endpoints
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

SubmitResetCancel

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.

- b. Enter the following information:

Template Name	Enter the name of the template (up to 20 characters).
Priority	Select a priority value (0 to 7), where 7 is the highest priority. These values are used by the Weighted Fair Queuing and Strict Priority queuing models to allocate excess bandwidth to the competing classes of applications.
Guaranteed Bandwidth	<p>Enter a percentage of the bandwidth that is guaranteed to be allocated to the applications in the traffic class. Lower values indicate that the traffic in the class is more likely to be delayed. Traffic may be dropped when the guaranteed bandwidth is exceeded, such as during a burst of higher-priority traffic.</p> <p>The total guaranteed bandwidth across all traffic classes cannot exceed 80 %. Also, the total guaranteed bandwidth across all endpoints cannot exceed 80 % of the outbound speed.</p>
Maximum Bandwidth	Enter the maximum percentage of the bandwidth that can be allocated to the applications in the traffic class. Traffic is dropped when the maximum bandwidth is exceeded. A zero indicates that all traffic in the class is dropped.



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

- 3. Click **Submit** to 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 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 bandwidth detection 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 179).

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 106: Enabling Bandwidth Management by Endpoint

WXC-10.90.64.131

Device SetupCompressionQoSAccelerationMonitorAdminHelp

QoS

Outbound

Setup Wizard...
Overview
Direct Setup
Traffic Classes
Templates
Endpoints
ToS/DSCP
Start/Stop

Inbound

Outbound QoS Endpoints

Find: GO HELP

Select the WAN mode that best describes your network.

☒ Dedicated Circuits ☐ Over-subscribed -- Outbound Speed Kbps

Outbound QoS is currently enabled. You can enable/disable Outbound QoS by using the Setup Wizard or from the 'Start/Stop' page. If enabled, QoS policies will be applied to:

☒ All discovered WX devices and non-WX endpoints ☐ Only checked endpoints below

For some endpoints, actual circuit speeds may vary. In order to optimize Bandwidth Management for these endpoints, you should enable Bandwidth Detection and indicate the minimum speed (in Kbps) for the relevant endpoints. If you don't know the minimum speed, enter '0'. NOTE: Endpoints enabled for Bandwidth Detection must also be enabled for Compression. (See the [Endpoints](#) page under **Compression**.)

☒ Enable Bandwidth Detection when sending traffic to:

WX devices are automatically included in the list below. If you want to enable QoS to endpoints that are NOT reachable through a WX device, you can manually add the endpoint to the list by clicking **ADD**. To [view a list of remote networks NOT accessed through a remote WX device](#), click this link.

Endpoint	IP Address	Circuit Speed (Kbps)	Bandwidth Detection	Min. Speed (Kbps)
Other traffic		<input type="text" value="default"/>		
<input checked="" type="checkbox"/> Branch1	No remote WX	<input type="text" value="512"/>		
<input checked="" type="checkbox"/> WXC-10.90.64.67	10.90.64.67	<input type="text" value="1000000"/>	<input checked="" type="checkbox"/>	<input type="text" value="0"/>
<input checked="" type="checkbox"/> 5722/SR-50A	10.87.57.22	<input type="text" value="100000"/>	<input type="checkbox"/>	<input type="text"/>
<input checked="" type="checkbox"/> 10.87.58.22	10.87.58.22	<input type="text" value="1000"/>	<input type="checkbox"/>	<input type="text"/>

Endpoints enabled for acceleration cannot be disabled for QoS.

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 **Over-subscribed** 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 170).

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 186).

- 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 bandwidth detection for one or more endpoints, as described in Step 4.
 - Exclude specific address or subnet pairs from bandwidth management, as described in Step 5.
2. To enable bandwidth management for a remote endpoint:
- a. Select one of the following options:
 - **All discovered WX devices and non-WX endpoints.** Applies bandwidth management to all current and future WX and non-WX endpoints. The maximum remote WAN circuit speed defaults to 1 Mbps for the WX 15, and 1 Gbps for all other WX endpoints.
 - **Only checked endpoints below.** Applies bandwidth management only to the selected endpoints. Select or clear the check box next to the appropriate endpoints in the Endpoints column. This option disables QoS for remote endpoints discovered in the future.

Virtual endpoints have a **DELETE** button next to them, and are listed first by default. 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 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 bandwidth detection 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 170). 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 129), 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 186.
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 107: Adding Virtual Endpoints

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The 'QoS' section is expanded, showing 'Outbound' and 'Inbound' options. The 'Outbound' section is further expanded, showing 'Setup Wizard...', 'Overview', 'Direct Setup', 'Traffic Classes', 'Templates', 'Endpoints', 'ToS/DSCP', and 'Start/Stop'. The 'Endpoints' option is selected, leading to the 'Outbound QoS Endpoints > Add Endpoint' dialog. The dialog contains the following fields: 'Name' (a text input field), 'Circuit Speed' (a text input field with 'kbps' as a unit), and 'Subnets' (a text area for entering subnets or IP addresses). A 'Submit' button and a 'Cancel' button are at the bottom. A note on the right side of the dialog states: 'It is possible to enable Outbound QoS to remote networks that are not accessed through a remote WX device. To do this, define an endpoint by filling in the fields below, and click 'Submit'.' An example of subnet format is provided: 'Example: 45.1.20.0/255.255.255.0 45.1.35.123'.

Specify the following information, and click **Submit**. The maximum number of virtual endpoints depends on the device type (refer to “WX Device Specifications” on page 421).

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 186).

4. If the WAN bandwidth to a remote WX device is variable, such as for MPLS, Frame Relay, or shared satellite links, enable bandwidth detection for traffic sent to that device.

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

To enable bandwidth detection:

- a. Select Enable Bandwidth Detection and select one of the following options:
 - **All QoS-enabled WX devices.** Applies bandwidth detection to all remote WX devices for which QoS is enabled (default).
 - **ONLY WX devices checked under “Bandwidth Detection”.** Select the **Bandwidth Detection** 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: Bandwidth detection manages only traffic sent to other WX endpoints (service 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:
 - a. Click **Exclusions**.

Figure 108: Excluding Subnets or Hosts from Bandwidth Management

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The 'QoS' section is expanded, showing 'Outbound' and 'Inbound' options. The 'Outbound' section is further expanded, showing 'Setup Wizard...', 'Overview', 'Direct Setup', 'Traffic Classes', 'Templates', 'Endpoints', 'ToS/DSCP', and 'Start/Stop'. The 'Exclusions' page is displayed, showing a table with two columns: 'Between LAN side network' and 'And WAN side network'. The first row has an asterisk (*) in the LAN column and '10.87.53.0/255.255.255.0' in the WAN column, with a 'DELETE' button. Below the table, there are input fields for 'Between LAN side network' and 'And WAN side network', with an 'ADD' button. A note below the input fields states: 'Click on "Add" button to add a new pair. Enter IP address or address/subnet. Enter asterisk (*) to indicate that source or destination can be ANY address. Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0'. At the bottom, there are 'Submit', 'Reset', and 'Cancel' buttons.

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.

- b. 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.

6. Click **Submit** to activate the changes, or click **Reset** to discard them.
7. 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 compressed 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 7 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 7: 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	–

Name	DSCP	IP Precedence
AF12	12	–
AF13	14	–
AF21	18	–
AF22	20	–
AF23	22	–
AF31	26	–
AF32	28	–
AF33	30	–
AF41	34	–
AF42	36	–
AF43	38	–
EF	46	–

To set ToS/DSCP values by traffic class:

1. Click **QoS** in the menu frame, click **Direct Setup** in the left-hand navigation frame, and then click **ToS/DSCP**.

Figure 109: Setting ToS/DSCP Values

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The 'QoS' section is expanded in the left-hand navigation frame, showing 'Outbound' and 'Inbound' sections. Under 'Outbound', 'Direct Setup' is selected, and 'ToS/DSCP' is highlighted. The main content area is titled 'ToS/DSCP' and contains the following options:

- ☒ 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 decompression

Traffic class: ToS/DSCP value:

☐ Default

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 131) or if ToS marking for router-based balancing is in use (refer to “configure route” on page 384).

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 188.



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 129).

5. After compressed traffic from remote WX devices is decompressed, the **Restore original ToS/DSCP bits after decompression** 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.

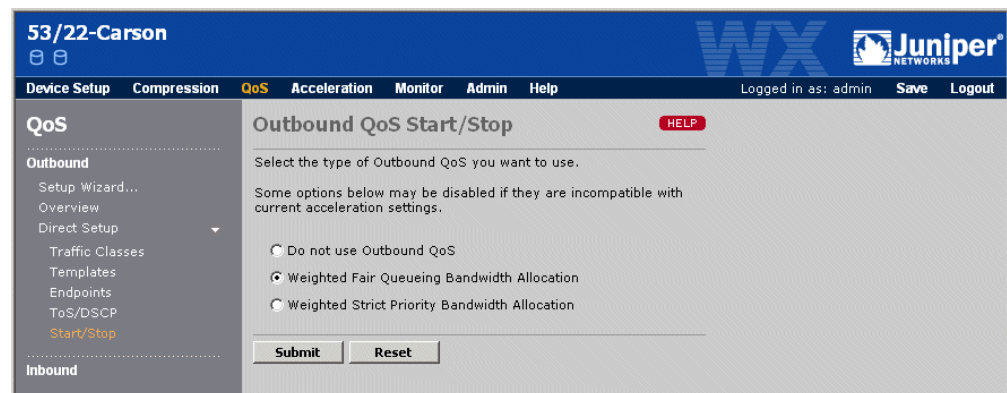
Starting and Stopping Outbound QoS

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 110: Starting and Stopping Outbound QoS



2. To stop the outbound QoS service, click Do not use Outbound QoS. If acceleration 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 Queuing 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 traffic classes for inbound bandwidth management.

Table 8: Inbound Bandwidth Management Classes

Class	Description
Compressed	Compressed traffic from other WX devices.
Intranet	Uncompressed 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 274).
TCP	TCP traffic that is not in the Compressed or Intranet class.
Default	All traffic that is not in the Compressed, Intranet, or TCP class.

For example, to enable inbound bandwidth management on D1 in Figure 111, set the inbound speed to 1500 Kbps (1.5 Mbps). You then set maximum bandwidth percentages for one or more of the 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 111: 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 112: Configuring Maximum Inbound QoS Bandwidths

53/22-Carson

Device Setup Compression **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

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
Compressed	<input type="text" value="100"/> %	Any traffic that has been compressed by a WX device.
Intranet	<input type="text" value="100"/> %	TCP traffic originating from the corporate network that has NOT been compressed.
TCP	<input type="text" value="100"/> %	TCP traffic NOT originating from the corporate network.
Default	<input type="text" value="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 113: Configuring Subnets for the Inbound QoS Intranet Class

53/22-Carson

Device Setup Compression **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

Inbound QoS > Intranet

Traffic that originates from within the corporate network (and has NOT been compressed) 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 214

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 207
- “Enabling Packet Flow Acceleration by Endpoint” on page 208
- “Enabling TCP Acceleration by Application” on page 212
- “Enabling Fast Connection Setup by Application” on page 213

Overview of Packet Flow Acceleration

While data compression effectively increases the available WAN bandwidth, application performance may be further constrained by network latency. Packet Flow Acceleration provides three methods to improve the throughput of compressed TCP application flows across high-speed, high-latency WAN links. For WX devices that support Multi-Path, you can enable Packet Flow Acceleration for the primary and/or secondary paths.

The following topics describe each acceleration method:

- “TCP Acceleration” in the next section
- “Forward Error Correction” on page 206
- “Fast Connection Setup” on page 206

TCP Acceleration

TCP Acceleration is intended primarily for high-latency environments, such as satellite connections, and long-haul high-bandwidth links, such as E3 and T3. TCP Acceleration is also beneficial when the compression percentage is very high.




NOTE: TCP Acceleration 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. TCP Acceleration 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 114: TCP Acceleration



TCP Acceleration is intended for applications that do large data transfers. In general, TCP Acceleration 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 compression 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, TCP Acceleration will improve performance if the host's TCP window size is 64 KB or less.




NOTE: Like high bandwidth and latency, high compression rates also increase the maximum window size, which increases the benefit of TCP Acceleration.

Asymmetric Routing for TCP Acceleration

For TCP Acceleration 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 115: TCP Acceleration Clusters for Asymmetric Routing Support



Note the following about asymmetric routing support:

- All devices in the same cluster must be the same model, such as all WXC 500s, and they must all have the same version of WXOS. Do not mix different device types or WXOS versions in the same cluster.
- 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, service 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 decompressors and tunnel load balancing settings (if any) defined on the WX device.

TCP Acceleration Statistics

Figure 116. shows an example of the statistics provided for TCP Acceleration. 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 116: Sample TCP Acceleration 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, TCP Acceleration may also benefit from Forward Error Correction, but TCP Acceleration 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.

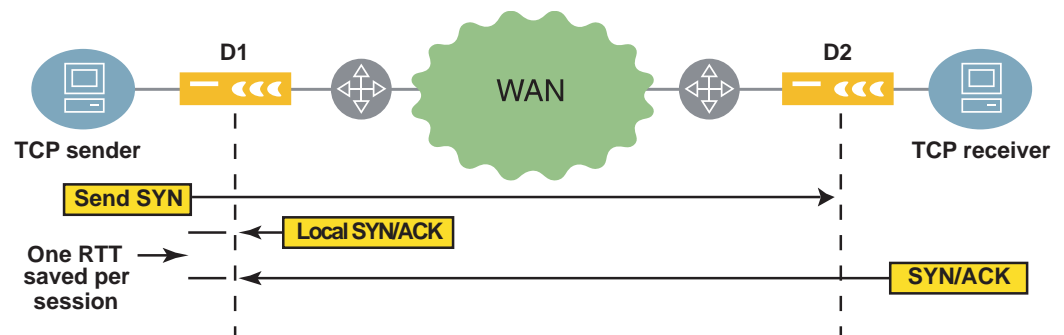
FEC is intended for use in high-loss, high-latency environments, such as satellite connections.

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 117: 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 118. 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 118, the FTP gains apply to a small number of sessions that probably affect only the traffic on the control port.

Figure 118: 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%	<div style="width: 24.3%;"></div>			
NetBios	714	68	9.5%	873.91	1088.90	19.7%	<div style="width: 19.7%;"></div>			

Requirements for Using Packet Flow Acceleration

To use Packet Flow Acceleration to accelerate application traffic between two WX devices, the following conditions must be met:

- The applications must be compressed (refer to “Compressing Traffic by Application” on page 151).
- A service tunnel must exist in both directions between the WX devices (refer to “Configuring Endpoints for Compression” on page 145).

- 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 167).

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 TCP Acceleration, you must 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 TCP Acceleration” on page 205 and “configure acceleration” on page 326). For TCP Acceleration 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 274).



NOTE: Packet Flow Acceleration is most effective in networks with high-speed connections and high latency, and/or very high compression rates. However, it 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. TCP Acceleration 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 service tunnels in both directions between the WX devices (refer to “Configuring Endpoints for Compression” on page 145).
- Enable compression for the applications you want to accelerate (refer to “Compressing Traffic by Application” on page 151).
- 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 167).

If you enable TCP Acceleration 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 213 and “Enabling TCP Acceleration by Application” on page 212).

To enable Packet Flow Acceleration by endpoint:

1. Verify that service tunnels exist in both directions between the WX devices that you want to support acceleration (refer to “Configuring Endpoints for Compression” on page 145).
2. Click **Acceleration** in the menu frame.

Figure 119: Enabling Packet Flow Acceleration

The screenshot shows the Juniper WX Acceleration Overview page. The left sidebar has a menu with 'Acceleration' selected, containing 'Packet Flow Acceleration' (with sub-items 'Overview', 'TCP Acceleration (AFP)', and 'Fast Connection Setup') and 'Application Flow Acceleration' (with sub-items 'CIFS', 'Exchange', and 'HTTP'). The main content area is titled 'Acceleration Overview' and includes a 'Find' search box. It is divided into two steps: 'Step 1: Enable desired Acceleration capabilities' and 'Step 2: Specify how enabled Acceleration capabilities are applied to endpoints'. Step 1 has checkboxes for 'TCP Acceleration (AFP)', 'Fast Connection Setup*', and 'Forward Error Correction†'. Step 2 has radio buttons for 'Accelerate all QoS-enabled endpoints using default settings' and 'Accelerate checked endpoints using custom settings'. Below this is a table of endpoints with columns for Name, IP Address, and checkboxes for TCP Acceleration (AFP), Fast Connection Setup, Forward Error Correction, Recovery Packets, and Data Packets. The table lists five endpoints: 52/22-CARSON, 54/22-SM250, 55/22-SR100, 5722/SR-50A, and 10.87.58.22. At the bottom are buttons for 'Select All', 'Clear', 'Show Advanced Settings', 'Submit', and 'Reset'. A note on the right states: 'Note: QoS must be enabled on endpoint before it can be accelerated. * Should only be used for connections with application generate many very short-TCP connections (e.g. HTTP) across high latency links. † Should only be used for connections that are subject to high loss and do not have FEC enabled on the satellite modem CSU/DSU.'

Name	IP Address	TCP Acceleration (AFP)	Fast Connection Setup	Forward Error Correction	Recovery Packets	Data Packets
<input checked="" type="checkbox"/> 52/22-CARSON	10.87.52.22	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input checked="" type="checkbox"/> 54/22-SM250	10.87.54.22	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input checked="" type="checkbox"/> 55/22-SR100	10.87.55.22	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input checked="" type="checkbox"/> 5722/SR-50A	10.87.57.22	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9
<input checked="" type="checkbox"/> 10.87.58.22	10.87.58.22	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1	9

3. At the top of the page, select the check box next to each of the acceleration methods that you want to use for one or more of the remote endpoints.
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 service tunnel exists and outbound QoS is configured correctly. The 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 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 service 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 acceleration for one or both paths. To configure Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 129.

5. Select the methods to be used for each endpoint or for all endpoints:.

TCP Acceleration	<p>Intended for high-latency environments, such as satellite connections, long-haul high-bandwidth links, such as E3 and T3, and networks where compression rates are very high.</p> <p>TCP Acceleration must be enabled on both the sending and receiving device, and cannot be used simultaneously on the same path with Fast Connection Setup. TCP Acceleration is required for Network Sequence Caching and Application Flow Acceleration.</p> <p>NOTE: In some cases, you may need to do one or more of the following (refer to “configure acceleration” on page 326):</p> <ul style="list-style-type: none"> ■ Adjust the buffer size for optimum performance. ■ Increase the number of lost heartbeat packets allowed on high-loss links (compression may stop when consecutive heartbeat packets are lost). ■ Enable clustering if the outbound and return traffic does not always traverse the same two WX devices. ■ If tunnel load balancing is enabled, verify that it is “Flow based” or “Per-destination” (refer to “Configuring Tunnel Load Balancing Policies” on page 155) ■ For device speeds of 20 Mbps or more, enable fast service tunnels for greater throughput if acceleration is more important than compression (refer to the “config reduction set fast-reduction-tunnel” command “Fast compression tunnels” on page 376).
Fast Connection Setup	<p>Intended for applications that have many short sessions, such as HTTP 1.0 and NetBios. The sending device locally acknowledges session requests for destinations known to be active. Short sessions are those that last less than ten times the round-trip time (RTT).</p>
Forward Error Correction	<p>Intended for high-loss environments. The sending device sends recovery packets with the data to reduce the number of retransmissions required when data packets are lost. By default, one recovery packet is sent for every nine data packets. To change the number of data and recovery packets, click Show Advanced Settings at the bottom of the page.</p> <p>After you enable FEC, check the monitoring report periodically. If losses are not persistent, disable FEC to avoid the overhead required to process recovery packets.</p>
Recovery Packets and Data Packets	<p>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 9.</p> <p>Note the following:</p> <ul style="list-style-type: none"> ■ 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 9: 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 %

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

You can now enable acceleration for specific applications, as described in the following sections.

Enabling TCP Acceleration by Application

After TCP Acceleration is enabled (refer to “Enabling Packet Flow Acceleration by Endpoint” on page 208), you can select the applications whose outgoing traffic you want to accelerate. TCP Acceleration 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 TCP Acceleration for one or more applications:

1. Click **Acceleration** in the menu frame, and then click **TCP Acceleration** in the left-hand navigation frame.

Figure 120: Enabling TCP Acceleration by Application

The screenshot shows the Juniper WX Acceleration (AFP) configuration page. The left-hand navigation frame shows the 'Acceleration' menu item selected. The main content area is titled 'TCP Acceleration (AFP)' and contains a list of applications with checkboxes for selection. The 'Application' list includes AOL, CIFS, Clearcase, CVS, DNS, Exchange, Filenet, FTP, Groupwise, Hostname Resolution, HTTP, HTTPS, ICA, ICMP, XWindows, and Undefined Applications. The 'Select All' button is highlighted. The page also includes a 'Submit' button and a 'Reset' button.

2. Select the check box next to each application that you want to accelerate using TCP Acceleration, or click **Select All**. The selected applications are accelerated only if they are also being compressed (refer to “Compressing Traffic by Application” on page 151).



NOTE: TCP Acceleration 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 208, 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 121: Enabling Fast Connection Setup by Application

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration' (highlighted), 'Monitor', 'Admin', and 'Help'. The left-hand navigation frame shows 'Acceleration' selected, with sub-options: 'Packet Flow Acceleration' (Overview, TCP Acceleration (AFP), Fast Connection Setup) and 'Application Flow Acceleration' (CIFS, Exchange, HTTP). The main content area is titled 'Fast Connection Setup' and includes a 'HELP' button. It contains the following text: 'When Fast Connection Setup is enabled, checked applications will be accelerated, provided they have also been checked for Compression. (See the Application Filter page under Compression.)' and 'Checked application traffic will be included in Acceleration report statistics only if the applications are also checked for Monitoring. (See the Monitoring page under Device Setup > Applications.)'. Below this is a table of applications with checkboxes:

Application	Check Box
AOL	<input type="checkbox"/>
CIFS	<input type="checkbox"/>
Clearcase	<input type="checkbox"/>
CVS	<input type="checkbox"/>
DNS	<input type="checkbox"/>
Exchange	<input type="checkbox"/>
Filenet	<input type="checkbox"/>
FTP	<input type="checkbox"/>
Groupwise	<input type="checkbox"/>
Hostname Resolution	<input type="checkbox"/>
HTTP	<input checked="" type="checkbox"/>
HTTPS	<input type="checkbox"/>
ICA	<input type="checkbox"/>
ICMP	<input type="checkbox"/>
XWindows	<input type="checkbox"/>

At the bottom of the application list is a 'Clear' button. Below the list are 'Submit' and 'Reset' buttons.

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 compressed (refer to “Compressing Traffic by Application” on page 151).



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 214
- Enabling Microsoft CIFS Acceleration on page 218
- Enabling Microsoft Exchange Acceleration on page 221
- Enabling HTTP Acceleration on page 223

Overview of Application Flow Acceleration

Though technologies such as compression (MSR and NSC) and TCP Acceleration (formerly called Active Flow Pipelining) 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 TCP Acceleration 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 acceleration is supported for currently supported Microsoft operating systems (Desktop 2000, Server 2000, Server 2003, Windows XP Desktop) and for Samba servers version 3.0 and higher. By default, traffic flows between Vista and non-Vista devices are downgraded from SMB2 to SMB and accelerated (refer to “Enabling Microsoft CIFS Acceleration” on page 218).


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 122), 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 122: Microsoft CIFS/Exchange Read Operations



In write operations (Figure 123), 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 123: Microsoft CIFS/Exchange Write Operations

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 124), 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 124: HTTP Caching—Header-Only Mode




In HTTP cache header-and-body mode (Figure 125), 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 125: HTTP Caching—Header-and-Body Mode (WXC Devices)



If pre-fetch is enabled (Figure 126), 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 126: HTTP Pre-Fetch

To view the current cache usage, refer to "configure acceleration" on page 326.



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. 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 99).
- TCP Acceleration must be enabled on both the client- and server-side WX devices (refer to "Enabling TCP Acceleration by Application" on page 212).

To enable CIFS acceleration for one or more applications:

1. To add new CIFS application definitions to accelerate specific CIFS traffic:
 - a. Click **Device Setup > Applications > Definitions**, 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 127: Adding New CIFS Application Definitions

SR-10.88.9.100

Device Setup Compression QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Device Setup

- Basic
 - Addresses
 - Interfaces
 - Time
 - License Key
 - SNMP
 - Syslog Server
 - Local Routes
 - Registration Server
- AAA
- Applications
 - Overview
 - Definitions
 - Traffic Classes
 - Monitoring
- Encryption
- Advanced

Application Definitions > CIFS

Application Name: CIFS_Server_1

Application Type: CIFS

SSL Encrypted: ☐ Yes

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

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 128: Enabling CIFS Acceleration

WX590-10.87.245.2

Device Setup Compression QoS **Acceleration** Monitor Admin Help

Logged in as: admin Save Logout

Acceleration

Packet Flow Acceleration

Overview
TCP Acceleration (AFP)
Fast Connection Setup

Application Flow Acceleration

CIFS
Exchange
HTTP

CIFS Acceleration HELP

TCP Acceleration (AFP) must be enabled before CIFS acceleration can be used. It is only necessary to enable CIFS acceleration on the WX device closest to the client.

☐ Disable CIFS Acceleration

☒ Enable CIFS Acceleration for checked applications

☒ Disable SMB signing when **not** required by the server*

☒ Apply SMB signing across the WAN when required by the server

Username

Enter Password

Verify Password

Domain (optional)

Application

☒ CIFS

* This option will not disable SMB signing for an end-station that has SMB signing required by default (e.g. Microsoft Domain Controllers). See the user guide for further information.

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 TCP Acceleration.
4. Select the following options to accelerate CIFS transactions when Server Message Block (SMB) signing is enabled or required by the server:
 - **Disable SMB signing when not required by the server.** Allows CIFS transactions to be accelerated for servers that have SMB signing enabled, but not required (enabled by default).
 - **Apply SMB signing across the WAN when required by the server.** Allows CIFS transactions to be accelerated for servers that require SMB signing. The SMB signature is based on a key derived from the login password. To allow the WX to log in to a server and create a signature, specify a user name, password, and Windows domain (optional) that matches an account on the appropriate Windows servers. Note the following:
 - SMB signing occurs between the client-side WX and the server, not between the WX and the client.
 - Traffic flows between Vista and non-Vista devices are downgraded from SMB2 to SMB to allow acceleration (to disable this feature, refer to page 329). CIFS traffic between two Vista devices using SMB2 is not accelerated.

Alternatively, you can disable SMB signing on Windows 2000 and Windows 2003 servers (see Step 7). When SMB signing is required, CIFS transactions are not accelerated unless **Apply SMB signing** is enabled on the client-side WX.

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.
7. If you want to disable SMB signing on Windows 2000 or Windows 2003 domain controllers, refer to the Microsoft Web site:

<http://support.microsoft.com/kb/887429>

Enabling Microsoft Exchange Acceleration

You can accelerate 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:

- Outlook 2000, 2002 or 2003 clients running on Windows 2000 or XP, 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 compression 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 99)
- TCP Acceleration must be enabled on both the client- and server-side WX devices (refer to “Enabling TCP Acceleration by Application” on page 212).

To enable Exchange acceleration for one or more applications:

1. To add new Exchange application definitions to accelerate specific Exchange traffic:
 - a. Click **Device Setup > Applications > Definitions**, 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 129: Adding New Exchange Application Definitions

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The left sidebar shows 'Device Setup' expanded with options like 'Basic', 'Addresses', 'Interfaces', 'Time', 'License Key', 'SNMP', 'Syslog Server', 'Local Routes', and 'Registration Server'. The main content area is titled 'Application Definitions > New'. It contains a form for creating a new application definition. The 'Application Name' is 'Exchange_Server_1' and the 'Application Type' is 'Exchange'. The 'SSL Encrypted' checkbox is unchecked. Below the form is a table for defining application rules. The table has columns for 'Source Address', 'Source Port', 'Destination Address', 'Destination Port', 'Protocol', and 'Advanced'. The first row shows '10.10.20.45' for Source Address, '135' for Source Port, '10.10.20.45' for Destination Address, '135' for Destination Port, and 'Any' for Protocol. The 'Advanced' column has a dropdown menu and a 'CLEAR' button. Below the table, there are instructions for entering IP addresses and ports. At the bottom, there are 'Submit' and 'Cancel' buttons.

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
10.10.20.45	135	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
				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 130: Enabling Exchange Acceleration

The screenshot shows the Juniper WX configuration interface. The top navigation bar includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The left sidebar shows 'Acceleration' expanded with options like 'Packet Flow Acceleration' and 'Application Flow Acceleration'. The main content area is titled 'Exchange Acceleration'. It contains a section for enabling Exchange acceleration. The 'Enable Exchange Acceleration for checked applications' radio button is selected. Below this, there is a table for selecting applications. The table has columns for 'Application' and 'Select All'. The first row shows 'Exchange' and 'Select All'. The second row shows 'Exchange_Server_1' and 'Select All'. At the bottom, there are 'Submit' and 'Reset' buttons.

Application	Select All
Exchange	Select All
Exchange_Server_1	Select All

Submit Reset

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 TCP Acceleration.

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 99)
- TCP Acceleration must be enabled on both the client- and server-side WX devices (refer to “Enabling TCP Acceleration by Application” on page 212).



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 **Device Setup > Applications > Definitions**, 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 131: Adding New HTTP Application Definitions

SR-10.88.9.100 Number of Active Clients -- 2

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

Basic

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

AAA

Applications

- Overview
- Definitions
- Traffic Classes
- Monitoring

Encryption

Advanced

Application Definitions > New

Application Name: HTTP_Server_1

Application Type: HTTP

SSL Encrypted ☐ Yes

Application traffic will be identified using the following rules

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
10.10.20.55	80			Any	Advanced CLEAR
		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

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 132: Enabling HTTP Acceleration

53/22-Carson

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Acceleration

Packet Flow Acceleration

- Overview
- TCP Acceleration (AFP)
- Fast Connection Setup

Application Flow Acceleration

- CIFS
- Exchange
- HTTP

HTTP Acceleration HELP

TCP Acceleration (AFP) must be enabled before HTTP acceleration can be used. It is only necessary to enable HTTP acceleration on the WX 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 TCP Acceleration.

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 326.

Chapter 8

Configuring IP Security (IPSec) and SSL Optimization

The following sections describe how to configure IP security (IPSec) between WX devices, and how to optimize SSL traffic:

- “Configuring IP Security (IPSec)” on page 227
- “Optimizing SSL Traffic” on page 240

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 229
- “Using the IPSec Setup Wizard” on page 229
- “Defining IPSec Settings by Endpoint” on page 234
- “Defining IPSec Templates” on page 236
- “Defining the Default IPSec Policy” on page 238
- “Defining the IPSec Application Filter” on page 239

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).



NOTE: IPSec is NOT supported on a WX 100 stack server with one or more clients.

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 compression and decompression are enabled (refer to “Configuring Endpoints for Compression” on page 145). Specific endpoints need not be enabled for compression.
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 229.



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 234).
 - Add new templates, or change a template name or key lifetimes (refer to “Defining IPSec Templates” on page 236).

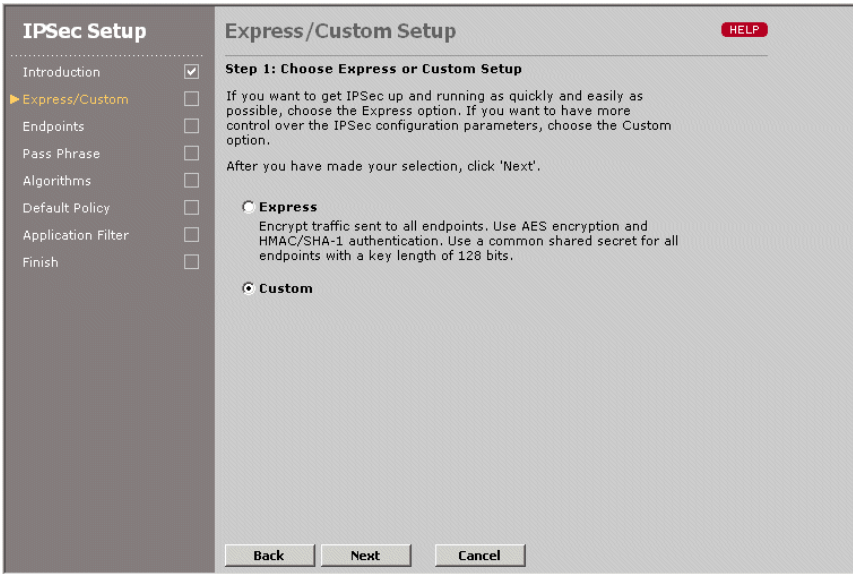
Using the IPSec Setup Wizard

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 236.

To run the IPSec Setup Wizard:

1. In the Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **IPSec Setup Wizard**.

2. Click **Enable IPSec** and click **Next**.



3. Select one of the following, and click **Next**.

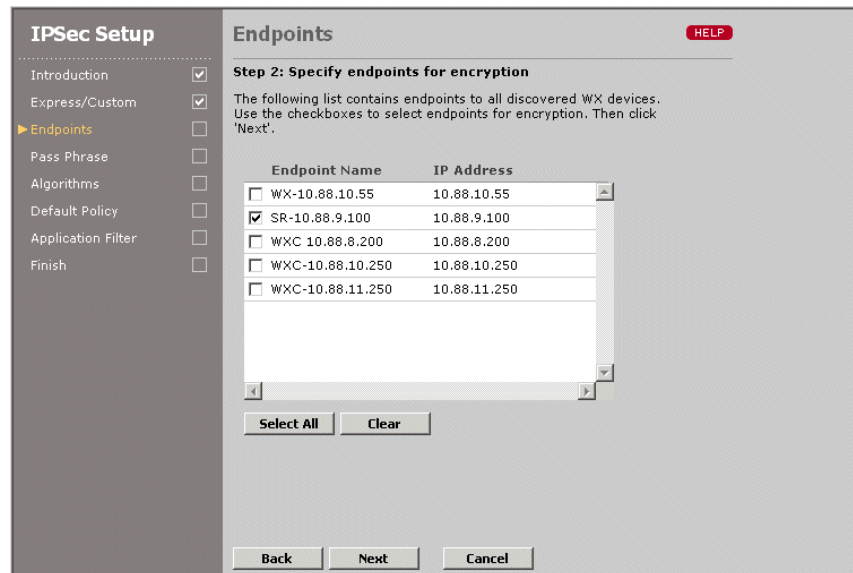
Express	Applies the default Wizard template to all endpoints, and prompts you to enter a single pass phrase that applies to all endpoints (up to 64 characters). Note that assigning the same pass phrase to all endpoints is a poor security practice, and is not recommended. The default template uses the following algorithms: <ul style="list-style-type: none">■ Encryption. Advanced Encryption Standard with a 128-bit key (AES-128)■ Authentication. Secure Hash Algorithm (HMAC/SHA-1) If you select this option, enter the pass phrase for all endpoints on the next page, and go to Step 9.
Custom	Allows you to select specific endpoints that support IPSec, specify a separate pass phrase for each endpoint, and select the template algorithms.

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 129.



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 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"> ■ Any. The algorithm selected for the peer endpoint is used. If both endpoints specify Any, AES-128 is used. ■ AES-128. Advanced Encryption Standard with a 128-bit key. ■ AES-192. AES with a 192-bit key. ■ AES-256. AES with a 256-bit key. ■ 3DES. Triple Digital Encryption Standard with a 168-bit key.
Authentication Algorithm	<p>Select the algorithm used to authenticate outbound traffic:</p> <ul style="list-style-type: none"> ■ Any. The algorithm selected for the peer endpoint is used. If both endpoints specify Any, HMAC/SHA-1 is used. ■ HMAC/SHA-1. Secure Hash Algorithm. ■ HMAC/MD5. Message Digest 5.

The screenshot shows the 'IPsec Setup' window with the 'Algorithms' tab selected. The left sidebar lists the setup steps: Introduction, Express/Custom, Endpoints, Pass Phrase, Algorithms (highlighted), Default Policy, Application Filter, and Finish. The main area is titled 'Algorithms' and contains 'Step 4: Select Encryption and Authentication Algorithms'. It explains that users can choose algorithms to encrypt traffic and authenticate endpoints. Below this, there are two dropdown menus: 'Encryption Algorithm' and 'Authentication Algorithm', both currently set to 'Any'. At the bottom are 'Back', 'Next', and 'Cancel' buttons.

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.

The screenshot shows the 'IPsec Setup' window with the 'Default Policy' tab selected. The left sidebar is the same as the previous screen, with 'Default Policy' highlighted. The main area is titled 'Default Policy' and contains 'Step 5: Define Default Encryption Policy'. It explains that for subnets not announced by a remote WX device, the following lists determine how packets will be handled. It also states that packets for subnets listed under 'Encryption Required' will be dropped and logged, while packets for subnets listed under 'Encryption Optional' will be passed-through unencrypted. Below this, there are two empty text boxes for 'Encryption Required' and 'Encryption Optional'. At the bottom are 'Back', 'Next', and 'Cancel' buttons.

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: <IP address>/<subnet mask>
Encryption Optional	Enter destination addresses or subnets (one per line) for which traffic can be sent unencrypted. For example, if subnet 10.10.0.0/255.255.0.0 is specified as encryption required, you can specify one or more smaller subnets in that range where encryption is optional, such as 10.10.20.0/255.255.255.0. If an address or subnet is in both lists, the traffic is sent unencrypted.

10. Select the encryption policy for each application:

- **If Configured.** Traffic is encrypted if IPSec is configured for the remote WX (the default).
- **No.** Traffic is never encrypted.
- **Required.** Traffic is dropped and logged if IPSec is not configured (or unavailable) for the remote WX.

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

12. Under Encryption in the left-hand navigation frame, click **IPSec Overview** to refresh the IPSec Overview page, which now shows the template name (Wizard) assigned to each of the selected endpoints.

13. 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 236.



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 Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **IPSec Overview**.

Figure 133: IPSec Overview

- 2. 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 129.

- 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 236.
- 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
	Normal operation — A secure connection is established between this device and the remote device.
	Configuration change — 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.
	No security association — 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 238).

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

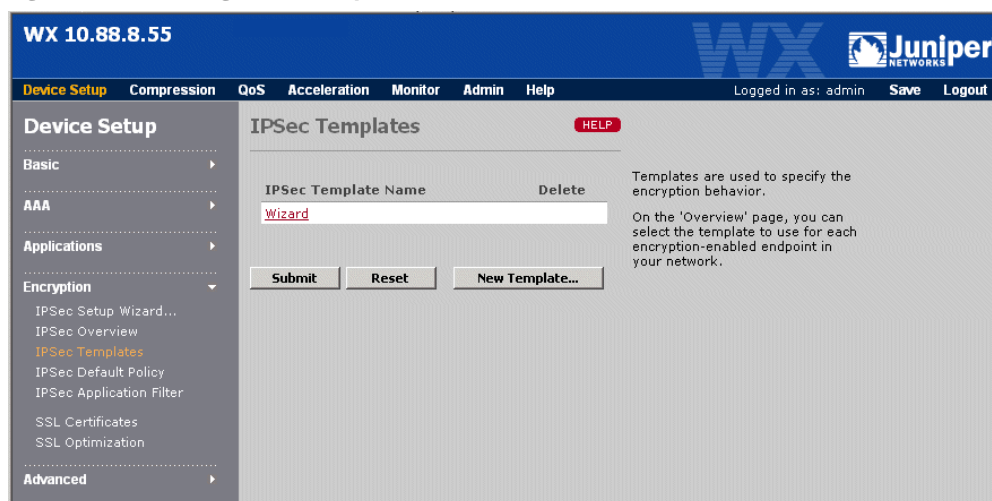
Defining IPSec Templates

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 234.

To define IPSec templates:

1. In the Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **IPSec Templates**.

Figure 134: Defining IPSec Templates



From the IPSec Templates page, you can:

- Add a new template, as described in Step 2.
 - 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:
 - a. Click **New Template**.

Figure 135: Defining a New IPSec Template

WX 10.88.8.55

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
- Applications
- Encryption**
 - IPSec Setup Wizard...
 - IPSec Overview
 - IPSec Templates
 - IPSec Default Policy
 - IPSec Application Filter
 - SSL Certificates
 - SSL Optimization
- Advanced

IPSec Templates HELP

Template Name

Encryption Algorithm

Authentication Algorithm

Key Lifetime

Time ☐ Never expires ☒ Expires in hours

Data ☐ Never expires ☒ Expires in MB

b. Enter the following information:

Template Name	Enter the name of the template (up to 20 characters).
Encryption Algorithm	Select the algorithm used to encrypt outbound traffic: <ul style="list-style-type: none"> ■ Any. The algorithm selected for the peer endpoint is used. If both endpoints specify Any, AES-128 is used. ■ AES-128. Advanced Encryption Standard with a 128-bit key. ■ AES-192. AES with a 192-bit key. ■ AES-256. AES with a 256-bit key. ■ 3DES. Triple Digital Encryption Standard with a 168-bit key.
Authentication Algorithm	Select the algorithm used to authenticate outbound traffic: <ul style="list-style-type: none"> ■ Any. The algorithm selected for the peer endpoint is used. If both endpoints specify Any, HMAC/SHA-1 is used. ■ HMAC/SHA-1. Secure Hash Algorithm. ■ HMAC/MD5. Message Digest 5.
Key Lifetime	Specify the time and data limits for generated keys: <ul style="list-style-type: none"> ■ Time. Enter the number of hours before a generated key expires (up to 2160), or select Never expires. ■ Data. Enter the number of megabytes of traffic allowed before a generated key expires (up to 4000), or select Never expires. <p>Key negotiation begins when the key lifetime reaches 80 % of the time limit or 50 % of the data limit. Keys should be negotiated periodically for security purposes.</p>

- 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 296.

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 compression is disabled.

To change the default IPSec policy:

1. In the Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **IPSec Default Policy**.

Figure 136: Defining the IPSec Default Policy

WX 10.88.8.55

Device Setup Compression QoS Acceleration Monitor Admin Help

Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
- Applications
- Encryption**
 - IPSec Setup Wizard...
 - IPSec Overview
 - IPSec Templates
 - IPSec Default Policy**
 - IPSec Application Filter
 - SSL Certificates
 - SSL Optimization
- Advanced

IPSec Default Policy HELP

For subnets not announced by a remote WX device, the following lists determine how packets destined for that subnet will be handled.

The lists also apply to subnets advertised by a WX 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

Encryption Required

Encryption Optional

2. In the two text boxes, specify the destination addresses and subnets where encryption is required or optional, as follows:

Encryption Required	Enter destination addresses or subnets (one per line) for which traffic must be dropped and logged. The subnet format is: <IP address>/<subnet mask>
Encryption Optional	Enter destination addresses or subnets (one per line) for which traffic can be sent unencrypted. For example, if subnet 10.10.0.0/255.255.0.0 is specified as encryption required, you can specify one or more smaller subnets in that range where encryption is optional, such as 10.10.20.0/255.255.255.0. If an address or subnet is in both lists, or in neither list, the traffic is not encrypted.

3. Click **Submit** to 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 the IPSec Application Filter

When IPSec is configured between two WX devices, all application traffic is encrypted by default. For each application, you can disable IPSec entirely or require that the application's traffic be dropped and logged, rather than sent unencrypted. To view the system log, refer to "Viewing and Saving System Logs" on page 296.

To change the default IPSec application filter:

1. In the Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **IPSec Application Filter**.

Figure 137: Defining the IPSec Application Filter

WX 10.88.8.55 Juniper

Device Setup Compression QoS Acceleration Monitor Admin Help Logged in as: admin Save Logout

Device Setup

- Basic
- AAA
- Applications
- Encryption
 - IPSec Setup Wizard...
 - IPSec Overview
 - IPSec Templates
 - IPSec Default Policy
 - IPSec Application Filter**
 - SSL Certificates
 - SSL Optimization
- Advanced

IPSec Application Filter

Application Name	Tunnel in IPSec		
	Required	If Configured	No
AOL	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
CIFS	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Clearcase	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
CVS	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
DNS	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Exchange	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Filenet	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
FTP	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Groupwise	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
H.248	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Select All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Submit Reset

The IPSec Application Filter determines whether or not application traffic will be encrypted in a WX-to-WX IPSec tunnel.

Required: Application traffic will be dropped if an IPSec tunnel is unavailable.

If Configured: If an IPSec tunnel is not available, the traffic will be sent in the clear.

No: Application traffic will be sent in the clear.

The **Required** radio button will be available only when IPSec Encryption is enabled.

2. Select the encryption policy for each application:
 - **If Configured.** Traffic is encrypted if IPsec is configured for the remote WX. (the default).
 - **No.** Traffic is never encrypted.
 - **Required.** Traffic is dropped and logged if IPsec is not configured (or unavailable) for the remote WX. Traffic is never sent unencrypted.
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.

Optimizing SSL Traffic

The following sections describe how to optimize application traffic that uses the Secure Socket Layer (SSL) for encryption:

- “Overview of SSL Optimization” in the next section
- “Importing SSL Certificates” on page 241
- “Enabling Applications for SSL Optimization” on page 243

Overview of SSL Optimization

Application traffic that uses SSL encryption can be decrypted and optimized for transmission between WX devices in the same community. Optimization includes compression, acceleration, and QoS management. On all WX platforms except the WX 100, the optimized traffic can be re-encrypted using IPsec (recommended) or sent as clear text. For each application to be optimized, the SSL server certificates and private keys must be imported on the WX closest to the server.

Traffic flows that meet the following conditions are eligible for SSL optimization:

- SSL version 3 (or later)
- Key exchange algorithm is RSA
- Encryption cipher is AES128, AES256, DES, 3DES, IDEA, RC2, or RC4
- Message digest algorithm is HMAC/SHA-1 or HMAC/MD5
- SSL compression is NOT used

For a pair of WX devices to use SSL optimization, SSL optimization must be enabled on both devices, but certificates are imported only on the server-side WX.



NOTE: SSL certificates and private keys are NOT copied to a backup WX device.

Importing SSL Certificates

For each application that you enable for SSL optimization, the SSL certificates and private keys for the application server must be imported on the WX device closest to the server. You can import up to 100 certificates. To enable applications for SSL optimization, refer to “Enabling Applications for SSL Optimization” on page 243.

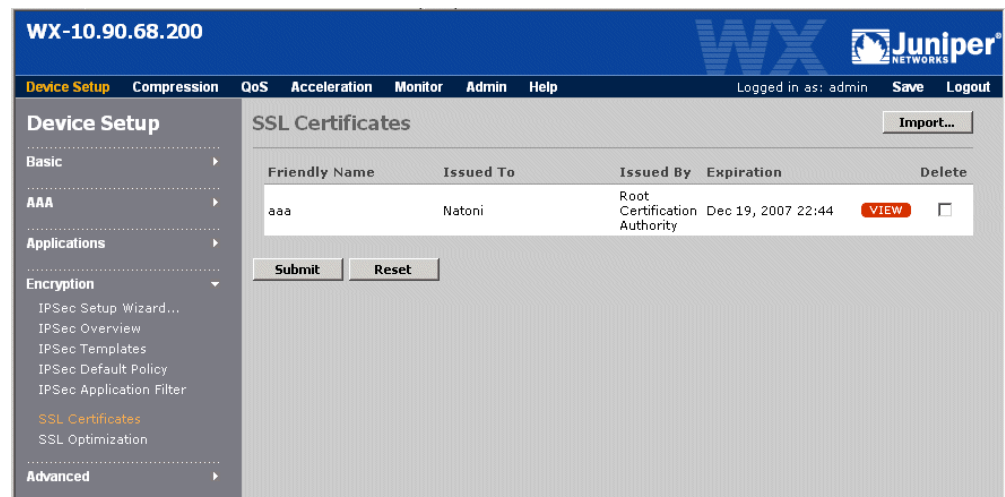


NOTE: Imported SSL certificates and private keys are NOT copied to a backup WX device.

To load SSL certificates and private keys:

1. In the Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **SSL Certificates**.

Figure 138: Viewing Imported SSL Certificates



2. To view the details of a certificate, click View. To delete outdated certificates, select the check box next to the appropriate certificates, and click **Delete**.
3. To import new certificates or update certificates that have a new private key:
 - a. Click **Import** at the top of the page.

Figure 139: Importing SSL Certificates

b. Enter the following information:

Friendly Name	Enter a unique name for the certificate to be imported (up to 15 characters). Use only letters, numbers, and underscores. To update the private key of an existing certificate, this name must match the existing name.
Certificate	Click Browse and select the certificate file. The supported formats are: <ul style="list-style-type: none"> ■ PKCS12 (extension “.p12” or “.pfx”) ■ PEM (extension “.pem”) ■ DER (extension “.der”) All file extensions are accepted, provided the certificate is in a supported format.
Private Key	If the server's private key is not in the certificate file, select Certificate and private key are provided as separate files , and click Browse to select the key file.
Pass Phrase	If the private key is encrypted, enter the password needed to access the key. Encrypted certificates are not supported. However, if you import a PKCS12 file, you must enter the password used to create the PKCS12 file.

c. 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 Applications for SSL Optimization

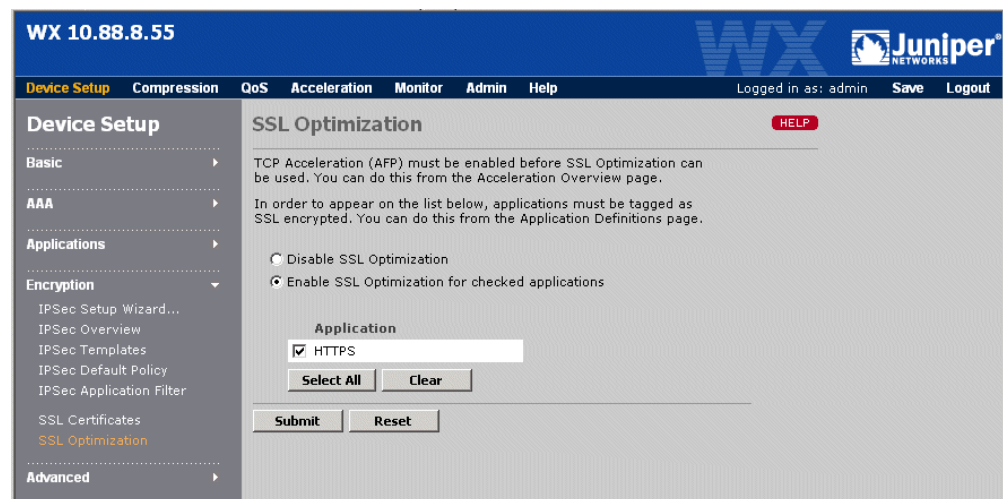
Applications that use SSL version 3 (or later) for encryption can be decrypted and optimized (compressed, accelerated, and managed by QoS). Note the following:

- TCP Acceleration must be enabled on both the client- and server-side WX devices (refer to “Enabling TCP Acceleration by Application” on page 212).
- IPSec encryption is recommended for use with SSL optimization, otherwise the decrypted and optimized SSL traffic is sent as clear text between the WX devices (refer to “Configuring IP Security (IPSec)” on page 227). Verify that IPSec is established between the WX devices before enabling SSL optimization.
- SSL certificates for each application server must be imported on the WX closest to the server (refer to “Importing SSL Certificates” on page 241).

To enable applications for SSL optimization:

1. Specify which applications use SSL encryption:
 - a. Click **Device Setup > Applications > Definitions**.
 - b. For each application that uses SSL encryption, click the application name, select the **SSL Encrypted** check box, and click **Submit**.
2. In the Device Setup page, click **Encryption** in the left-hand navigation frame, and then click **SSL Optimization**.

Figure 140: Enabling SSL Optimization



3. Select **Enable SSL Optimization for checked applications** and click the check box next to the appropriate applications, or click **Select All**. An application is listed here only if its application definition has the SSL Encrypted option enabled. Applications are grayed out unless they are also enabled for TCP Acceleration.
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.

Chapter 9

Monitoring and Reporting


This chapter describes how to view statistics for data compression, 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 246
- “Compression Statistics” on page 253
- “Outbound Bandwidth Statistics” on page 262
- “Inbound Bandwidth Statistics” on page 264
- “Acceleration Statistics” on page 266
- “Traffic Statistics” on page 274
- “Endpoints Summary” on page 277
- “Executive Summary” on page 278
- “Events Summary” on page 280

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 141). 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 141).

Figure 141: 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 115.
- 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 248
- “WAN Performance Statistics” on page 249

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 191. 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 “Compressing Traffic by Application” on page 151.

- 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 142: WAN Throughput Report

- 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 Demo Mode.

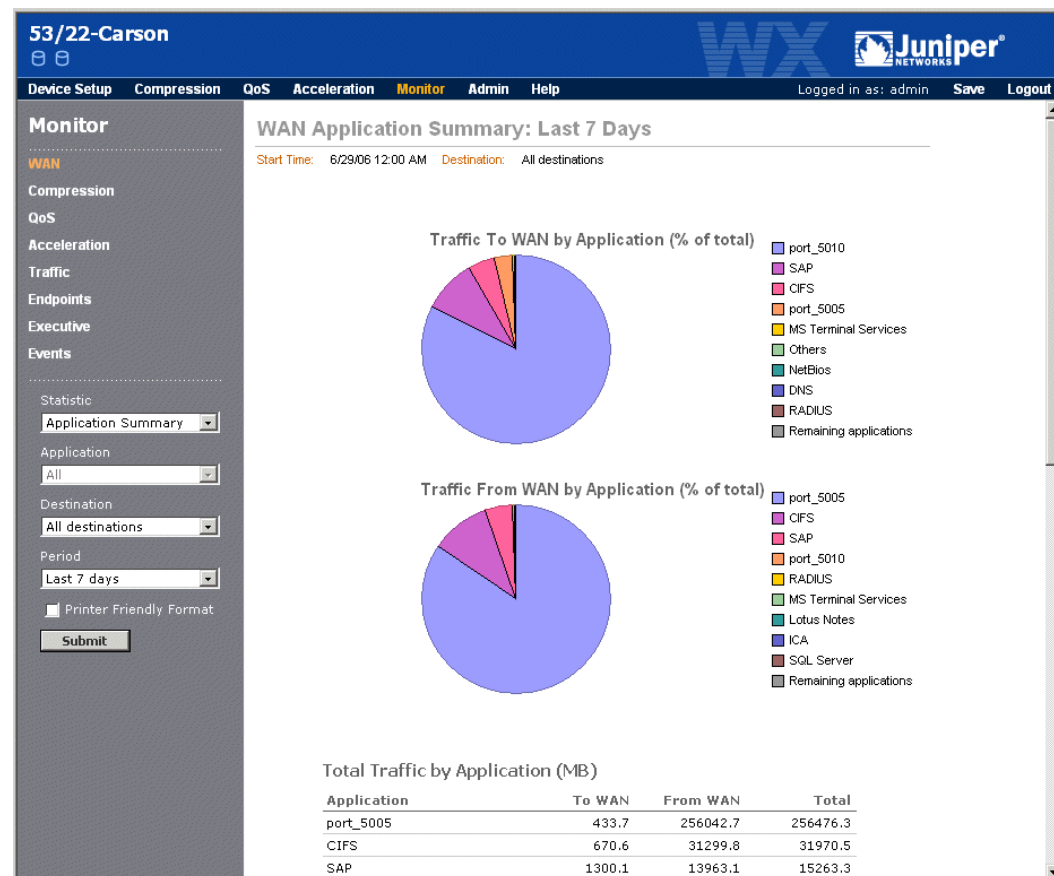
WAN Application Summary

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 191. The traffic to and from the WAN is shown for up to 40 monitored applications. To specify the monitored applications, refer to “Compressing Traffic by Application” on page 151.

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 143: 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 274).

WAN Performance Statistics

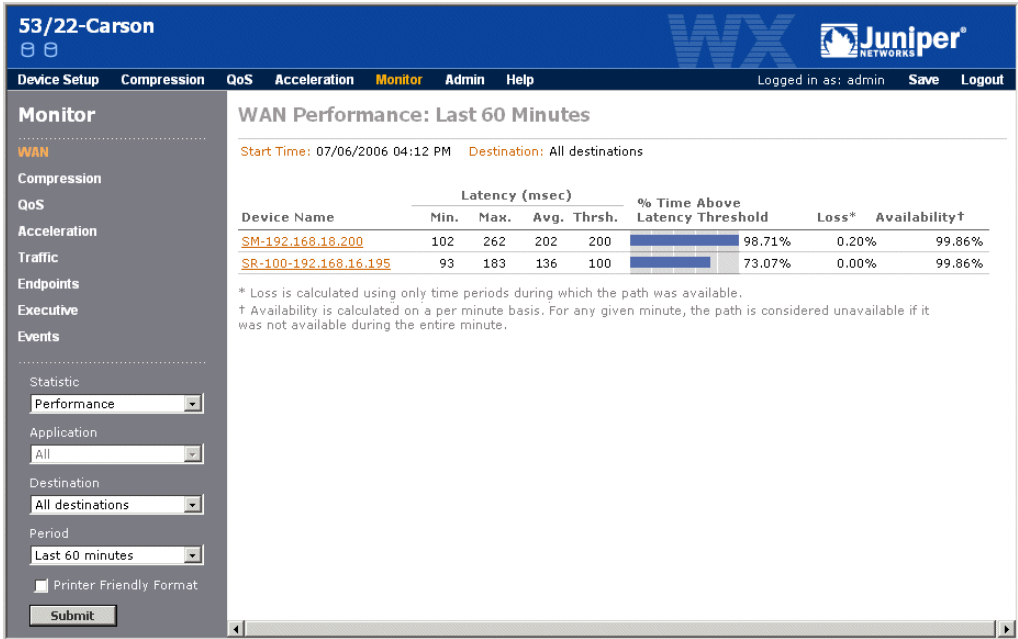
The WAN Performance report provides WAN loss and latency statistics, and performance events, between the current device and the remote WX devices that are enabled for either of the following:

- WAN performance monitoring (refer to “Configuring WAN Performance Monitoring” on page 138)
- 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 129)

To view WAN performance:

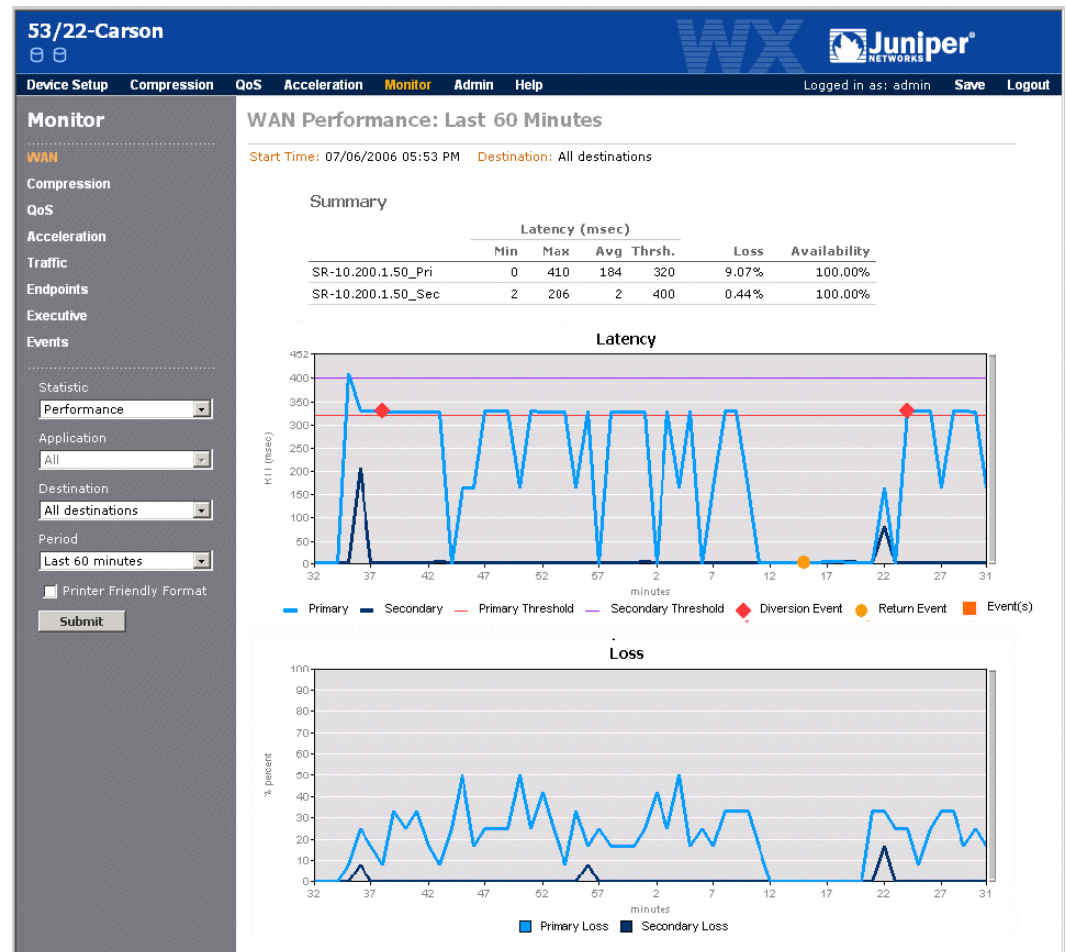
1. Click **Monitor** in the menu frame, and **WAN** in the navigation frame.
2. Select **Performance** from the **Statistic** menu, change one or more of the following report parameters, and click **Submit**.
 - Select a specific device from the **Destination** menu to view the performance graphs and 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 144: 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 145) or WAN performance monitoring (Figure 146).



Figure 145: Multi-Path WAN Performance Charts



For a Multi-Path device, the following information is shown on the Loss and Latency charts (Figure 145):

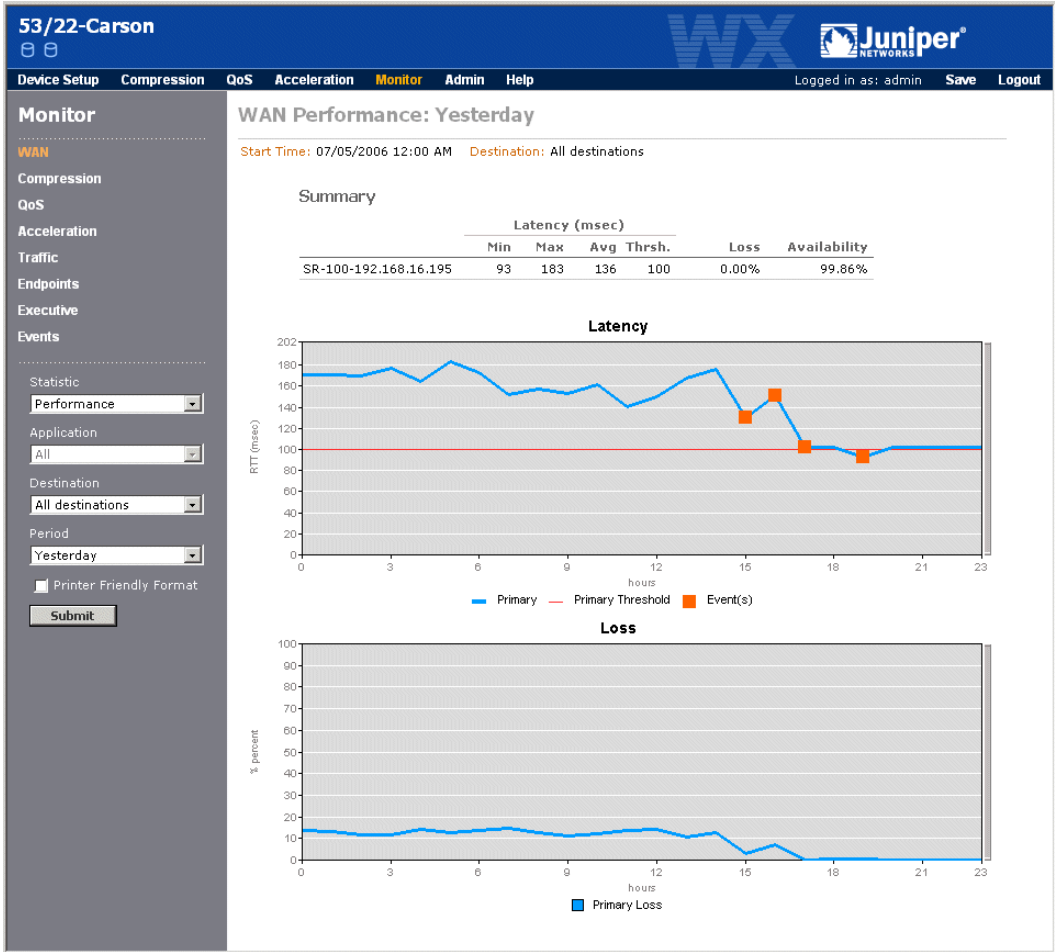
- The Latency chart shows the average round-trip time for the primary path (blue) and secondary path (black), and indicates the configured latency threshold for each path. The following icons are used to indicate performance events. An informational SNMP trap and a syslog entry are generated for each 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"> ■ Loss or latency threshold exceeded. Eligible traffic is diverted only if the alternate path's service tunnel is up and the loss and latency are below the specified thresholds. ■ Tunnel is down. Eligible traffic is diverted regardless of the alternate path's performance (if the alternate tunnel is up). Traffic that cannot be switched to the alternate path is passed through without compression (if the link is up and only the tunnel is down).

Icon	Description
	To view the status of the service tunnels, check the Multi-Path Endpoints page (refer to “Defining Multi-Path Endpoints” on page 135) or the Endpoints Summary report (refer to “Endpoints Summary” on page 277).
	Indicates that performance has returned to normal, and traffic was switched back to the preferred path (the service 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 146: Single-Path WAN Performance Charts



For WAN performance monitoring endpoints (Figure 146), 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.

Compression Statistics

This section describes the compression statistics displayed in the WXOS Web console:

- “Device Throughput Statistics” in the next section
- “Data Compression Statistics” on page 255
- “Application Summary Statistics” on page 258
- “Passthrough Statistics” on page 260
- “Packet Size Distribution Statistics” on page 261

Device Throughput Statistics

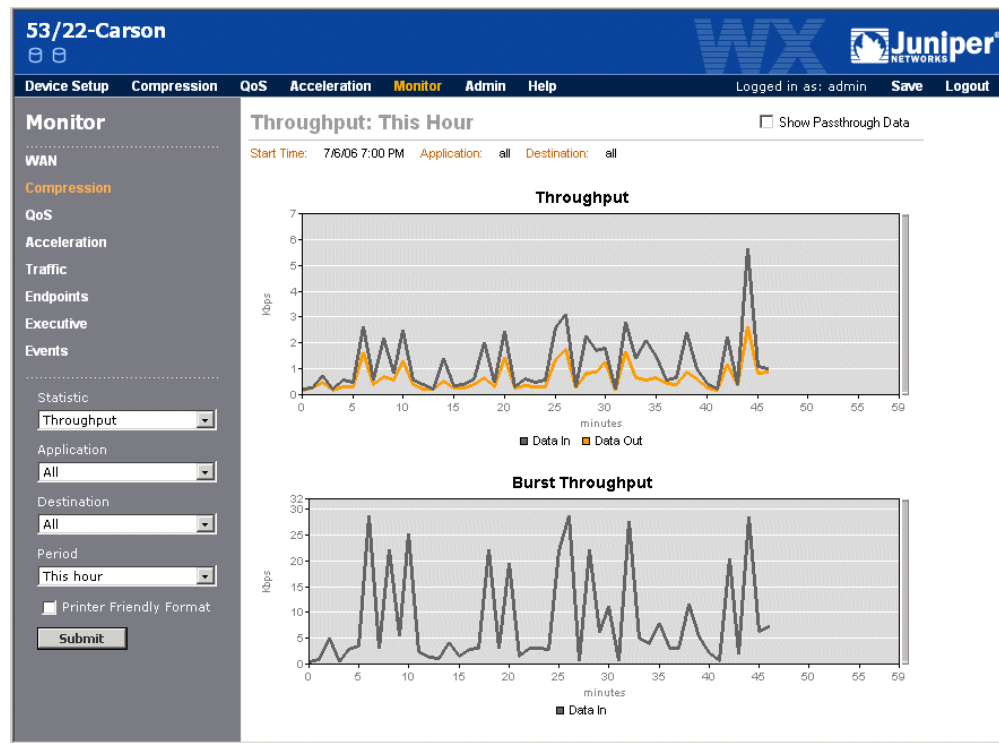
The device throughput statistics include a Throughput line graph and a Burst Throughput line graph. The Burst Throughput graph is shown only when you view the data for all applications and destinations. These statistics help you gauge the speed of traffic in and out of the device.

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 “Compressing Traffic by Application” on page 151.
 - 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 147).

Figure 147: 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 compression engine.
 - **Data Out** (orange line). Average data throughput out of the compression 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 compressed.
 - 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 compression 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 compressed.



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 260.

Data Compression Statistics

The data compression statistics include a Summary table, a Percent Compression 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 compression achieved for the traffic sent to each of the other WX devices.

Note that the percentage of data compression is not an average, but is based on the total number of bytes in and out of each device, as follows:


$$\% \text{ of Compression} = [(\text{Bytes In} - \text{Bytes Out}) / \text{Bytes In}] \times 100$$

To view data compression statistics:

1. Click **Monitor** in the menu frame.
2. Select **Compression** 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 “Compressing Traffic by Application” on page 151.
 - 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 Compression page opens (Figure 148).

Figure 148: Data Compression Statistics



3. Review the following information on the data compression graphs. Keep in mind that all values are for the selected application, destination, and time period.
 - The Summary table shows the following if All is selected from the **Destination** menu.
 - **Peak Data Compression.** Highest percentage of data compression 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 Compression.** Percentage of compressed data for the selected time period.
 - **Effective WAN Capacity.** Factor increase in WAN capacity resulting from the total data compression. For example, this value is 2.00 if total data compression is 50 %.
 - **Total Bytes Into Compression.** Number of bytes into the data compression engine.
 - **Total Bytes Out of Compression.** Number of bytes after data compression.
 - **Total Bytes Passed Through.** Number of bytes passed through without compression. To view the different types of passthrough traffic, refer to "Passthrough Statistics" on page 260.




NOTE: If a specific device is selected from the **Destination** menu, the Summary table shows the total data compression 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 **Compression By Endpoint** at the top of the page to view the data compression 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 compression by application for the device (to view the application statistics for all endpoints, refer to “Application Summary Statistics” on page 258).

- The Percent Compression graph shows how the average and peak percentage of data compression varied over the selected time period. Peak compression is shown only for all applications.
- The Bytes graph shows the number of megabytes in and out of the device (Figure 149).

Figure 149: Compression Bytes Graph



The Bytes graph includes the following:


- **Data In** (grey line). Number of bytes into the compression engine.
- **Data Out** (orange line). Number of bytes out of the compression 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 compressed. 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 260.

- 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 150).

Figure 150: Compression Packets Graph

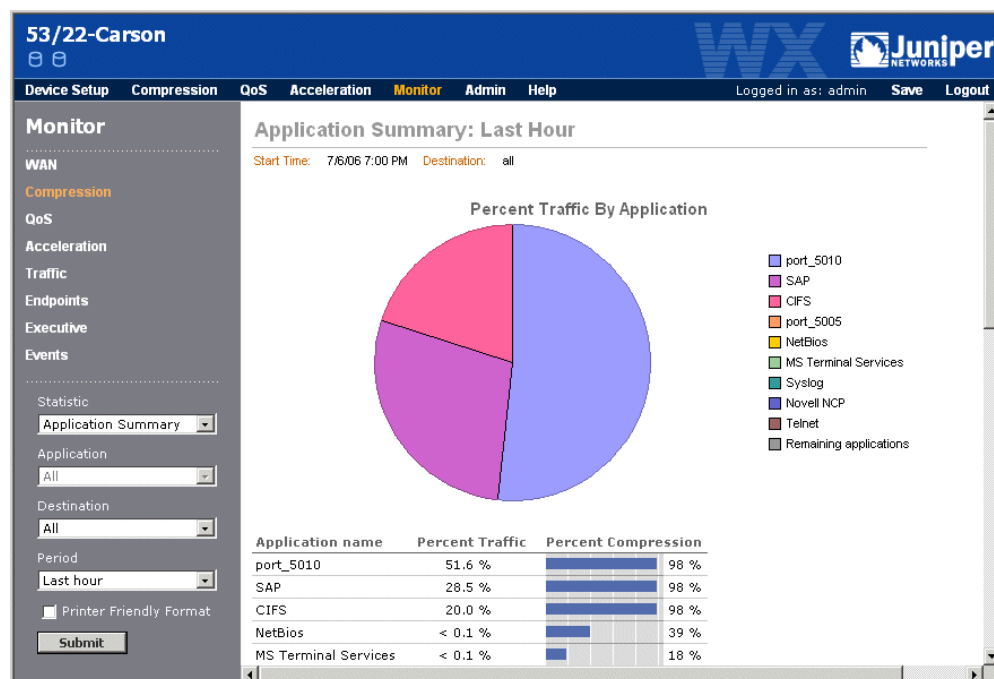


Application Summary Statistics

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 compression for each monitored application (up to 40). To specify the monitored applications, refer to “Compressing Traffic by Application” on page 151.

To view application summary statistics:

1. Click **Monitor** in the menu 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 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 151: 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 compression percentage. The **Others** category indicates the traffic for compressed 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 274).
 - **Percent Traffic.** Percentage of the total traffic into the device’s compression engine for each application.
 - **Percent Compression.** Percentage of data compression achieved for each application. A dash is shown for applications that have no traffic or cannot be compressed (such as encrypted applications). Data compression should be disabled for applications that consistently show little or no compression (refer to “Compressing Traffic by Application” on page 151).

Passthrough Statistics

Traffic that falls into one of several categories is passed through the WX device with no attempt at data compression. 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.

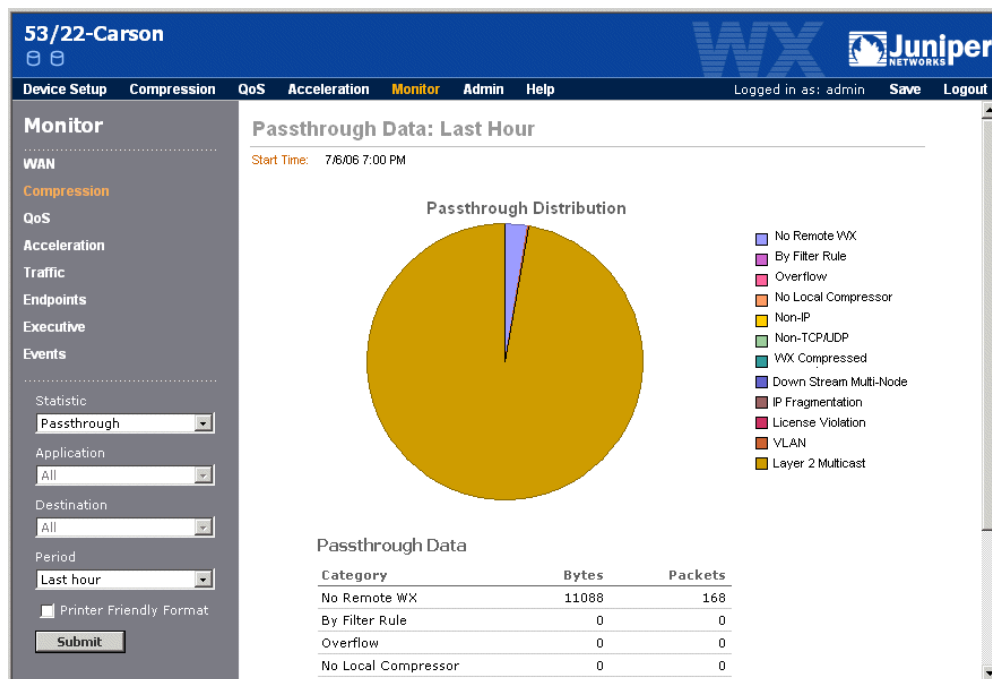


NOTE: There are no passthrough statistics for an off-path device where RIP is used to route traffic to the device. All traffic is sent through the service tunnel.

To view passthrough statistics:

1. Click **Monitor** in the menu frame.
2. Select **Passthrough** from the **Statistic** menu.
3. Select a time period from the **Period** menu, and click **Submit**.

Figure 152: Passthrough Statistics



The following table describes the passthrough categories.

Category	Description
No Remote WX	No WX device available to decompress the data, or compression is disabled for one or more devices.
By Filter Rule	Compression is disabled for specific applications or source/destination addresses (refer to “Compressing Traffic by Application” on page 151 and “Using Source/Destination Filters” on page 112).
Overflow	Traffic volume exceeded the device capacity.

Category	Description
No Local Compressor	Compression is disabled on this device (refer to “Configuring Endpoints for Compression” on page 145).
Non-IP	Non-IP traffic is not compressed.
Non-TCP/UDP	By default, only TCP/UDP application traffic is compressed. This category is invalid if you define non-TCP/UDP applications.
WX Compressed	Traffic was compressed by another WX device.
Down Stream Multi-Node	Traffic will be reduced by the next WX device.
IP Fragmentation	Always zero unless compression of IP fragments is disabled (refer to “configure filter” on page 342).
License Violation	The licensed throughput speed was exceeded.
VLAN	Total VLAN traffic that was not compressed 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 compressed because the intended destination is unknown.



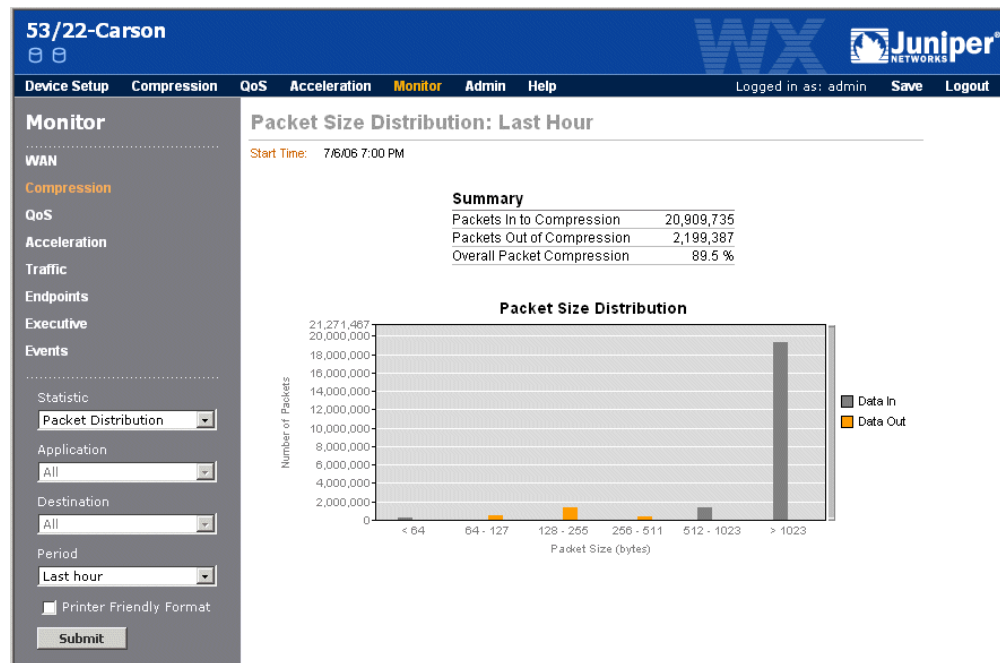
NOTE: Jumbo Gigabit Ethernet frames are also passed through without compression, but they are not counted in any of the above categories.

Packet Size Distribution Statistics

The Packet Size Distribution report shows the number of packets in and out of the compression engine, the percentage compression 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 153: 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 177.



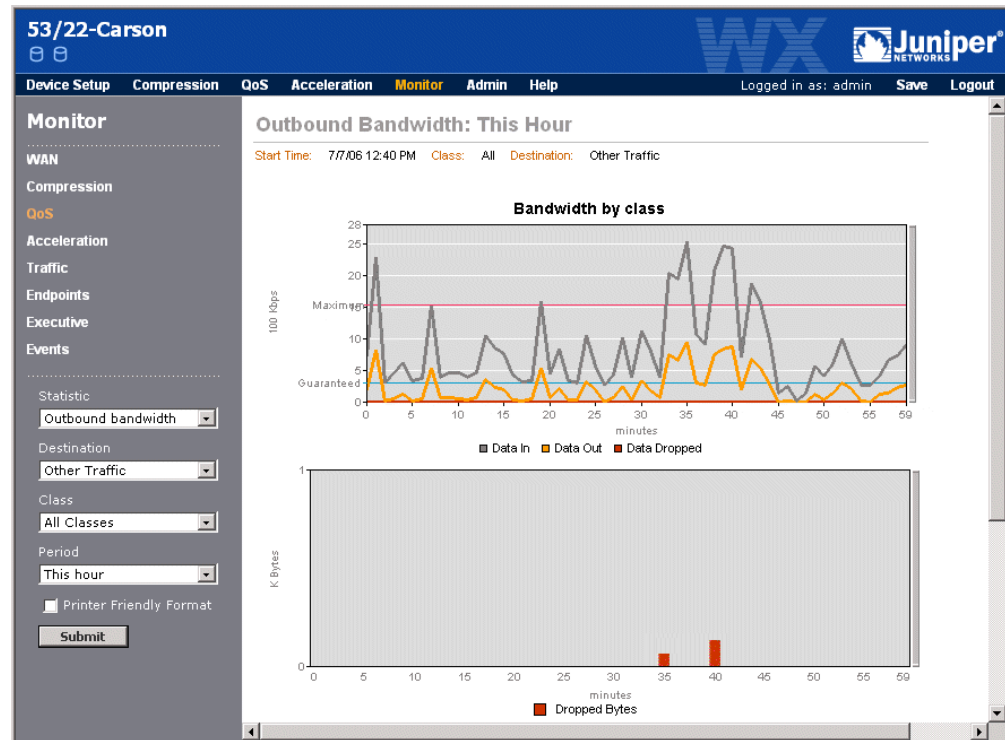
NOTE: Outbound bandwidth management is not effective for an off-path WX device unless all outbound WAN traffic is routed through the device.

To view outbound bandwidth statistics:

1. Click **Monitor** in the menu frame, and **QoS** in the navigation frame.
2. Select **Outbound bandwidth** 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 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** 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 154: 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 compression 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 (Compressed, Intranet, TCP, and Default) exceeds its maximum allocated bandwidth. To configure inbound QoS settings, refer to “Configuring Inbound QoS Policies” on page 199.



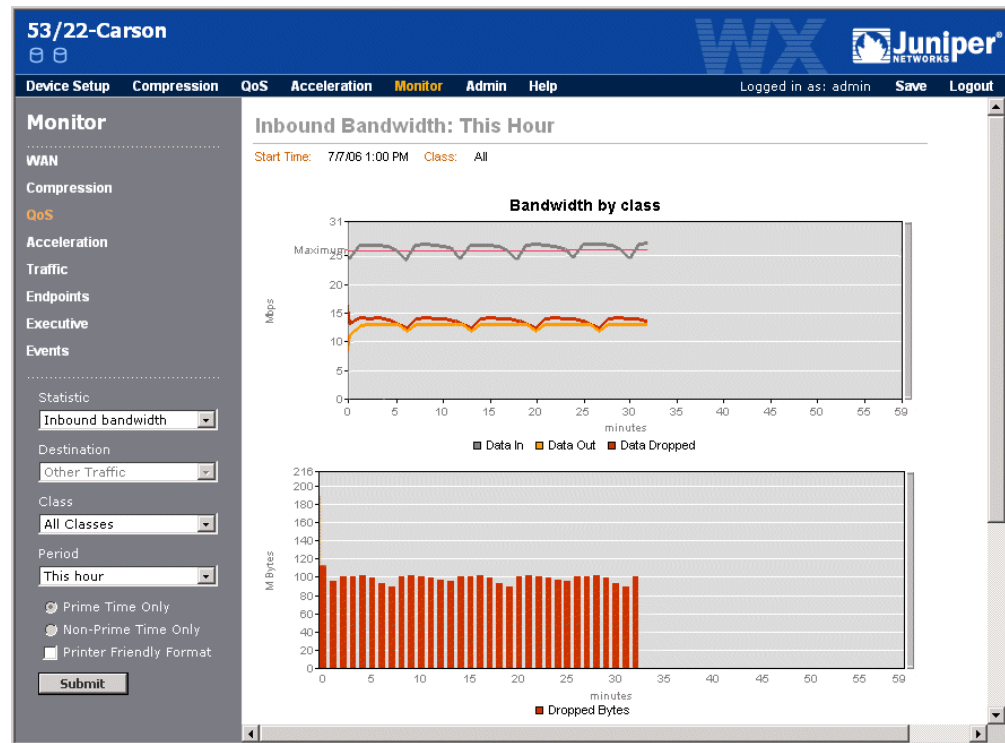
NOTE: Inbound bandwidth management is not supported for off-path WX devices.

To view inbound bandwidth statistics:

1. Click **Monitor** in the menu frame, and **QoS** in the navigation frame.
2. Select **Inbound bandwidth** from the **Statistic** menu, change one or more of the following report parameters, and click **Submit**.
 - Select a traffic class from the **Class** 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 155: 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 acceleration statistics displayed in the WXOS Web console:

- “TCP Acceleration Statistics” in the next section
- “Fast Connection Setup Statistics” on page 268
- “Forward Error Correction Statistics” on page 269
- “CIFS and Exchange Acceleration Statistics” on page 271
- “HTTP Acceleration Statistics” on page 272

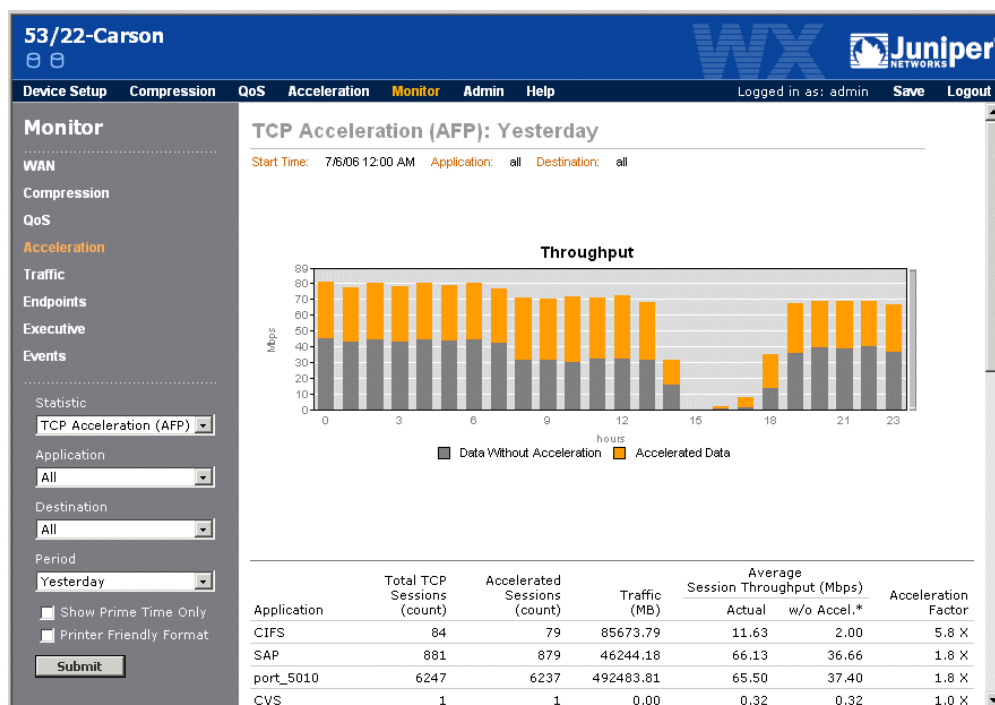
TCP Acceleration Statistics

If TCP Acceleration is enabled for one or more endpoints and applications, the TCP Acceleration report shows the session statistics and the average throughput improvements due to TCP Acceleration. To configure TCP Acceleration for specific endpoints and applications, refer to “Accelerating WAN Traffic” on page 203.

To view TCP Acceleration statistics:

1. Click **Monitor** in the menu frame, and then click **Acceleration** in the left-hand navigation frame.
2. Select **TCP Acceleration** from the **Statistic** 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 156: TCP Acceleration 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 TCP Acceleration enabled.
 - **Accelerated Data** (orange bars). Average increase in data throughput as a result of TCP Acceleration.
- 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 TCP Acceleration was disabled.

- **Acceleration Factor.** The performance increase for the accelerated sessions due to TCP Acceleration (actual throughput divided by the estimated throughput without acceleration). This value indicates the overall impact of TCP Acceleration.


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 compression 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 203.

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 Setup** from the **Statistic** 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 157: Fast Connection Setup Statistics



3. 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 compression 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

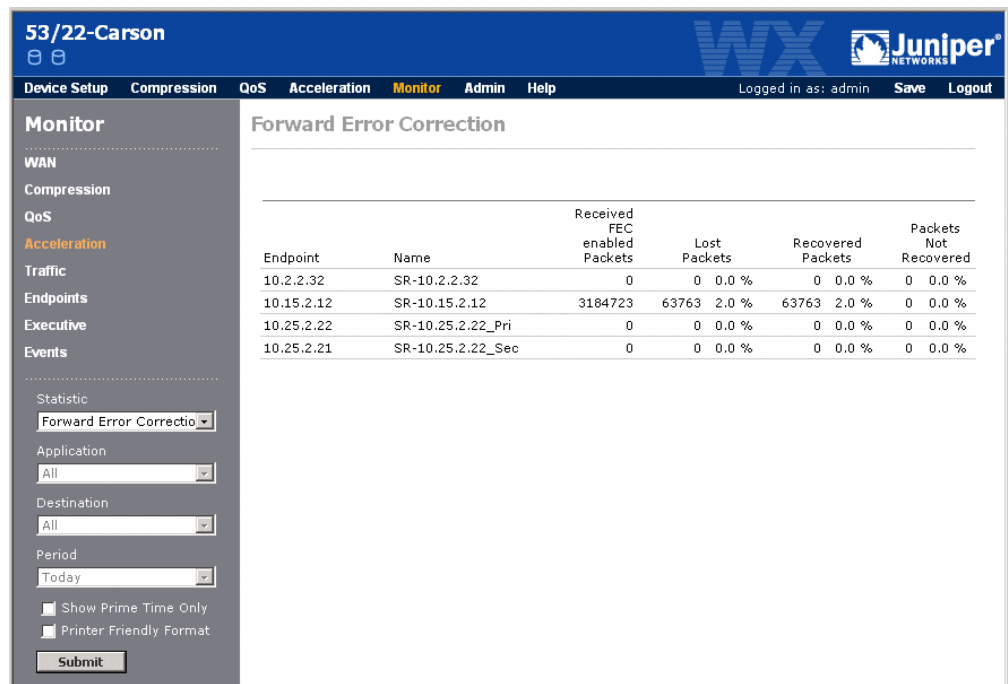
The Forward Error Correction report shows the number of packets received from each remote endpoint that has forward error correction enabled. The report also shows the percentage of received packets that were lost, recovered, and retransmitted. The statistics are cumulative since the last time the counters 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 203. 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** menu, and click **Submit**.

Figure 158: 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 214.

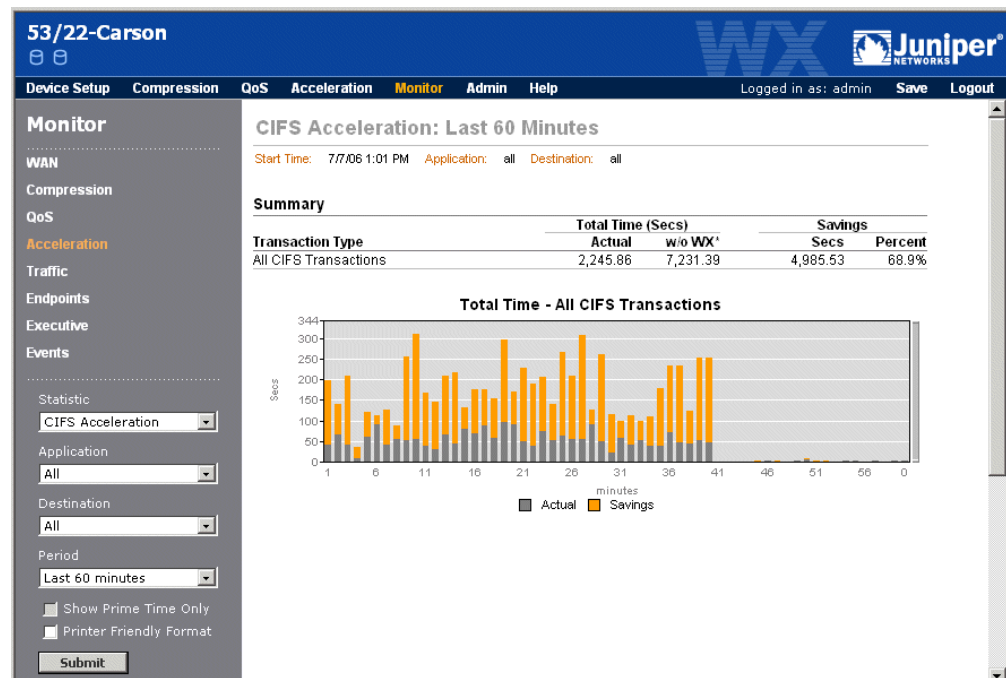


NOTE: View CIFS and Exchange acceleration reports on the client-side WX device. The acceleration statistics apply to the traffic in both directions. However, depending on the nature of the traffic flow, compression statistics may need to be viewed on the server-side device.

To view CIFS or Exchange acceleration statistics:

1. Click **Monitor** in the menu frame, and then click **Acceleration** in the left-hand navigation frame.
2. Select **CIFS Acceleration** or **Exchange Acceleration** from the **Statistic** menu, 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 159: 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 223.




NOTE: View HTTP acceleration reports on the client-side WX device. The acceleration statistics apply to the traffic in both directions. However, compression statistics should probably be viewed on the server-side device.

To view HTTP acceleration statistics:

1. Click **Monitor** in the menu frame, and then click **Acceleration** in the left-hand navigation frame.
2. Select **HTTP Acceleration** from the **Statistic** menu, 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 160: 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.



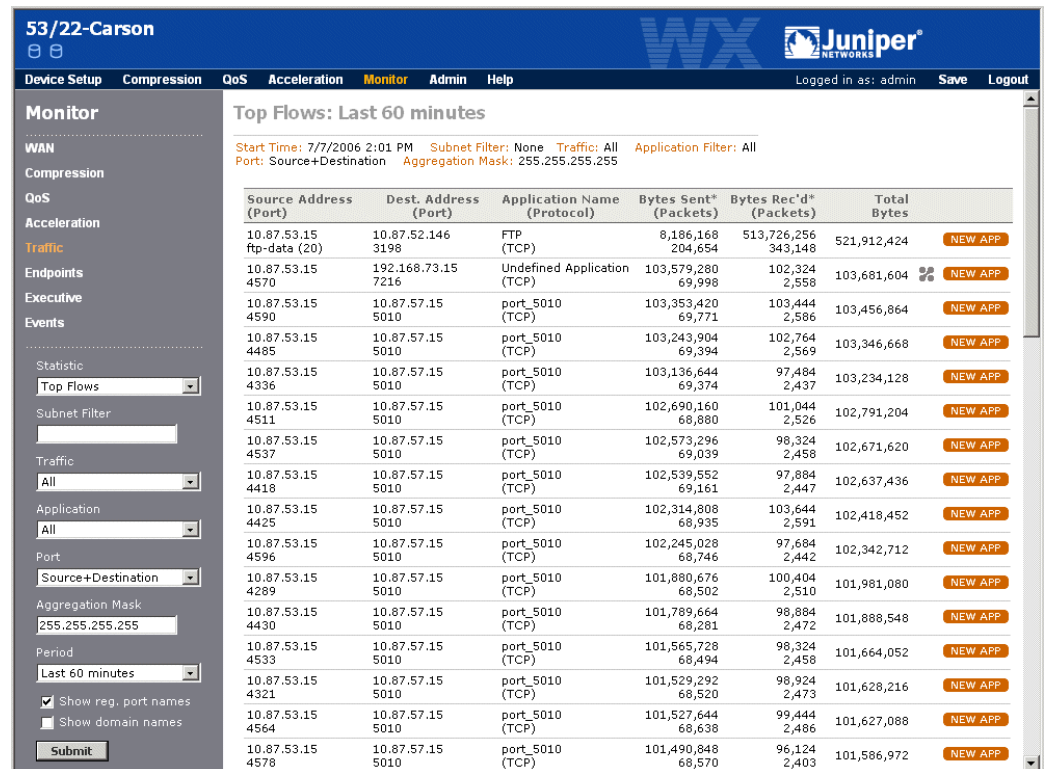
NOTE: A flow constitutes data sent and/or received from a single source IP address and port number, to a single destination IP address and port number over the same protocol. Only the traffic flows that started in the selected time period are shown.


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 446).

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 161: Traffic Statistics



- 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 95. Note that an  is shown next to flows for undefined applications.
- 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">■ All. All traffic for the selected statistic.■ All Compressed. Compressed traffic only.■ Compressed Undefined Apps. Compressed traffic for undefined applications only.■ Passthrough Only. Traffic sent from the WAN to the LAN that was not compressed. Does not apply to off-path WX devices or to in-line devices that use tunnel switching.
Application	<p>Select an application to limit the traffic to a specific application.</p>
Port	<p>If you select the top flows, you can select a view of the port information.</p> <ul style="list-style-type: none">■ Ignore Port. Traffic is consolidated across all ports for each pair of source and destination addresses.■ Source Only. Traffic is consolidated across the same source ports for each pair of source and destination addresses.■ Destination Only. Traffic is consolidated across the same destination ports for each pair of source and destination addresses.■ Source + Destination. Traffic is shown for each combination of source and destination port.
Aggregation mask	<p>If you select the top flows, sending addresses, or receiving addresses, you can enter a subnet mask to view all traffic from the same subnet as one consolidated entry. The default mask is “255.255.255.255”, which shows a separate flow for each host. (This was the “Subnet Mask” field in previous versions of WXOS.)</p>
Period	<p>Select the time period (last 60 minutes, last 24 hours, or all). Note that if you select Last 60 minutes or Last 24 hours, only the traffic flows that started in the selected time period are shown.</p>
Show reg. port names	<p>If you select the top flows, click the check box to view the registered names for all ports in the collected data. Clear the check box to view the names only for port numbers up to 1024.</p>
Show domain names	<p>If you select the top flows, click the check box to view the domain names for each IP address. To specify the DNS servers to be queried, refer to “Configuring Device Address and Contact Information” on page 63. The IP address is displayed if its domain name cannot be resolved (the DNS queries may take a few seconds).</p>

Endpoints Summary

The Endpoints Summary report shows the status of each tunnel, Network Sequence Caching, IPSec encryption, Multi-Path, and 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 162: Endpoints Summary

Endpoint	Name	Compression (Outbound)	Compression (Inbound)	Network Sequence Caching	Encryption	Circuit Speed (Kbps)	QoS (Outbound)	TCP Acceleration (AF)	Fast Connection Setup	Forward Error Correction	Multi-Path
10.15.2.12	SR-10.15.2.12	OK	OK	OK	OK	1500	OK	OK	OK	OK	OK
10.25.2.22	SR-10.25.2.22_Pri	Warning	OK	OK	OK	256	OK	OK	OK	OK	Down
10.25.2.21	SR-10.25.2.22_Sec	OK	OK	OK	OK	256	OK	OK	OK	OK	OK
10.2.2.32	SR-10.2.2.32	OK	OK	OK	OK	128	OK	OK	OK	OK	OK

2. The following icons are used to indicate the status of each connection:

Icon	Description
	OK — Indicates a connection between this device and the remote device for the following features: <ul style="list-style-type: none"> ■ Compression (outbound or inbound) ■ Network Sequence Caching (shown on WXC devices only) ■ Encryption ■ Multi-Path
	Warning — 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.

Icon	Description
	<p>Down — Indicates no connection between this device and the remote device for the following features:</p> <ul style="list-style-type: none"> ■ Compression. The outbound or inbound service tunnel is down (the remote device may be inaccessible). ■ Encryption. 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 238). ■ Network Sequence Caching. A problem exists or NSC is enabled on the local device, but disabled on the remote device.
	<p>Configured — 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):</p> <ul style="list-style-type: none"> ■ Outbound QoS ■ TCP Acceleration (must also be configured on the remote device for acceleration to occur) ■ Fast Connection Setup ■ Forward Error Correction
	<p>Not configured — The feature is not fully configured between this device and the remote device. However, in the Compression (Inbound) column, this icon indicates that the remote device does not have a service tunnel enabled to the local device.</p>
	<p>Unknown — The connection is in a transitory state.</p>
Blank	<p>A blank in the Network Sequence Caching column indicates that the remote device is not a WXC or the service tunnel to the device is down. Also, the circuit speed is blank if outbound QoS is not configured.</p>

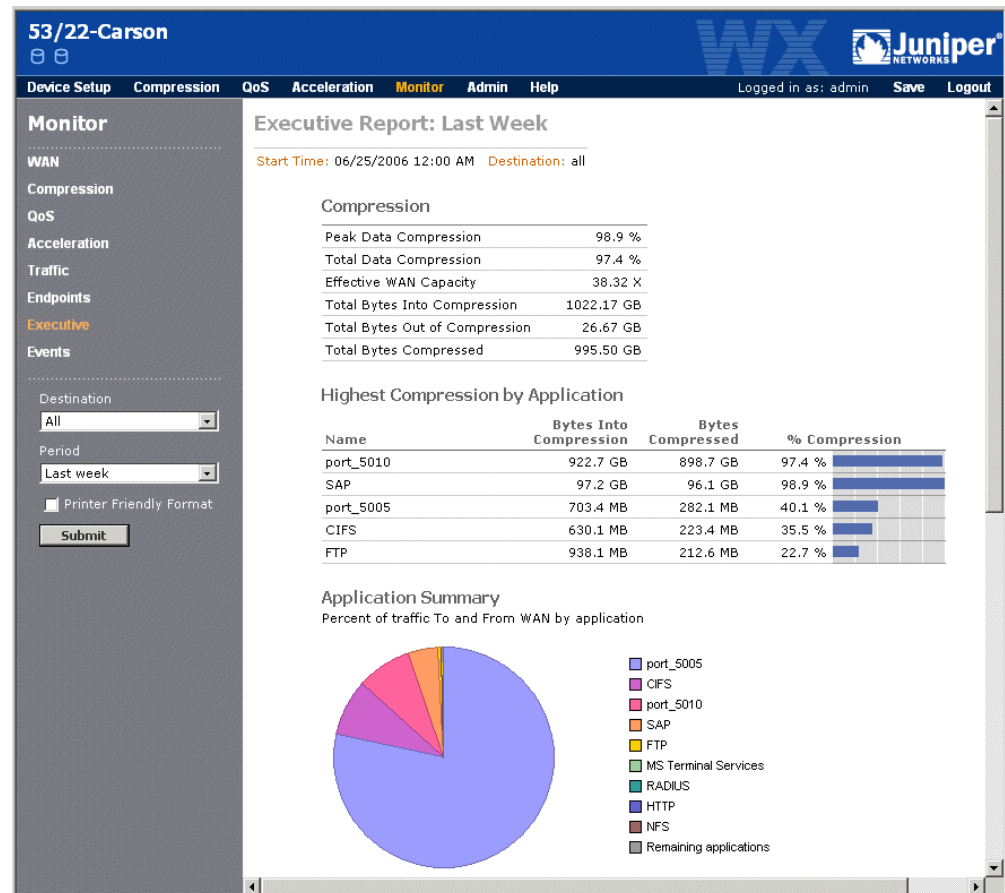
Executive Summary

The Executive report summarizes compression 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 163: Executive Summary



3. Review the following information. Keep in mind that all values are for the selected destination, and time period.
 - The Compression Summary table shows the following:
 - **Peak Data Compression.** Highest percentage of data compression 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 Compression.** Percentage of compressed data for the selected time period.
 - **Total Bytes Into Compression.** Number of bytes into the data compression engine.
 - **Total Bytes Compressed.** Number of bytes compressed.
 - **Total Bytes Out of Compression.** Number of bytes of traffic output after data compression.
 - **Effective WAN Capacity.** Factor increase in WAN capacity resulting from the total data compression. For example, this value is 2.00 if total data compression is 50 %.

- The Highest Compression by Application table has the following columns.
 - **Application Name.** Names of the top five monitored applications with the highest compression percentage. The **Others** category indicates the traffic for compressed applications that are undefined or unmonitored.
 - **Bytes Into Compression.** Number of bytes into the device's compression engine for each application.
 - **Bytes Compressed.** Number of bytes compressed for each application.
 - **Percent Compression.** Percentage of data compression 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 138).

Events Summary

The Events report lists the performance and system events that have occurred since the last time the WX device was restarted (up to 200 performance events and 200 system events). To define thresholds for performance events or to enable or disable the generation of specific system events, refer to “Configuring Events” on page 140.

To view the Events report:

1. Click **Monitor** in the menu frame, and click **Events** in the left-hand navigation frame.
2. Optionally, change the following report parameters, and click **Submit**.
 - Select the type of events displayed (performance or system) from the **Event Type** menu. The default is All.
 - Select the minimum severity level of the events displayed (OK, Warning, Major, or Critical), which correspond to the event color codes of green, yellow, orange, and red. The default is OK (the lowest level).
 - Select an application or traffic class from the **Application/Class** menu to view performance events for a specific application or traffic class. The default is All. This option applies only to performance events.

- Select a device from the **Destination** menu to view performance events for a specific device. The default is All. This option applies only to performance events.

Figure 164: Events Summary

WXC 10.88.8.200

Device Setup Compression QoS Acceleration **Monitor** Admin Help

Logged in as: admin Save Logout

Monitor

WAN

Compression

QoS

Acceleration

Traffic

Endpoints

Executive

Events

Event Type: All

Minimum Severity: OK

Application/Class: All

Destination: All

☐ Show Prime Time Only

☐ Printer Friendly Format

Submit

Events

Minimum Severity: OK Application/Class: All Destination: All

Event	Destination	App/Class	Value	Threshold	Date/Time
Compression (%)	10.88.9.100	All (Aggregated)	95	96	JUL 13 2006, 06:00:PM
TCP Acceleration Throughput (Kbps)	All (Aggregated)	All (Aggregated)	32	1000	JUL 13 2006, 06:00:PM
Decompressor Session Closed	--	--	--	--	JUL 06 2006, 02:22:PM
Compressor Session Closed	--	--	--	--	JUL 06 2006, 02:22:PM
Decompressor Session Closed	--	--	--	--	JUL 06 2006, 02:22:PM
Primary Self Registration	--	--	--	--	JUL 06 2006, 02:22:PM
Compressor Session Closed	--	--	--	--	JUL 06 2006, 02:22:PM
Compressor Session Closed	--	--	--	--	JUL 06 2006, 02:22:PM
Security Login Success	--	--	--	--	JUL 06 2006, 02:19:PM
Compressor Session Opened	--	--	--	--	JUL 06 2006, 02:15:PM

3. Review the following information.

- **Performance Events.** A performance event shows the destination, application or traffic class, and performance threshold specified in the event definition, along with the value that violated the threshold, and the date and time the event occurred. Click the event name to view the evaluation time period and whether the evaluation is limited to prime time days and hours.

For a description of how each performance metric is calculated, refer to Table 4 on page 142.

- **System Events.** A system event shows the event name and the date and time the event occurred. Click the event name to view additional details. For example, for a self-registration event, click the event name to view the IP address of the registration server. For a description of each system event, refer to “SNMP Traps and Syslog Messages” on page 427.

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 291.



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 285.
- “Loading a Device Configuration File” on page 286.
- “Loading a Boot Image” on page 287.
- “Clearing Application Monitoring Statistics” on page 288.
- “Setting the Device to the Factory Default Configuration” on page 288.
- “Rebooting the Device” on page 290.

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 165: Saving the Configuration

2. Select one of the following:

- | | |
|------------------|---|
| Flash memory | Save the current configuration to <i>startup.cfg</i> in Flash memory or click Save to the filename 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 “/juniper/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 “/juniper/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.

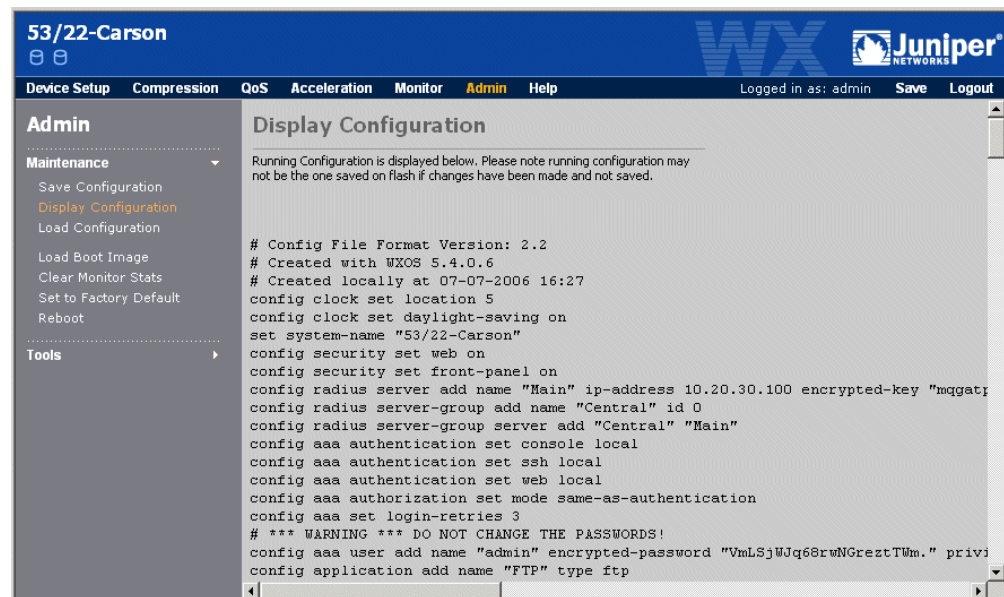
Displaying the Running Configuration

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 166: 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 305).

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 167: 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 290.

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 168: 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 290 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 169: Clearing Application Monitoring Statistics



2. To clear the application monitoring statistics, click **Clear**.

Setting the Device to the Factory Default Configuration

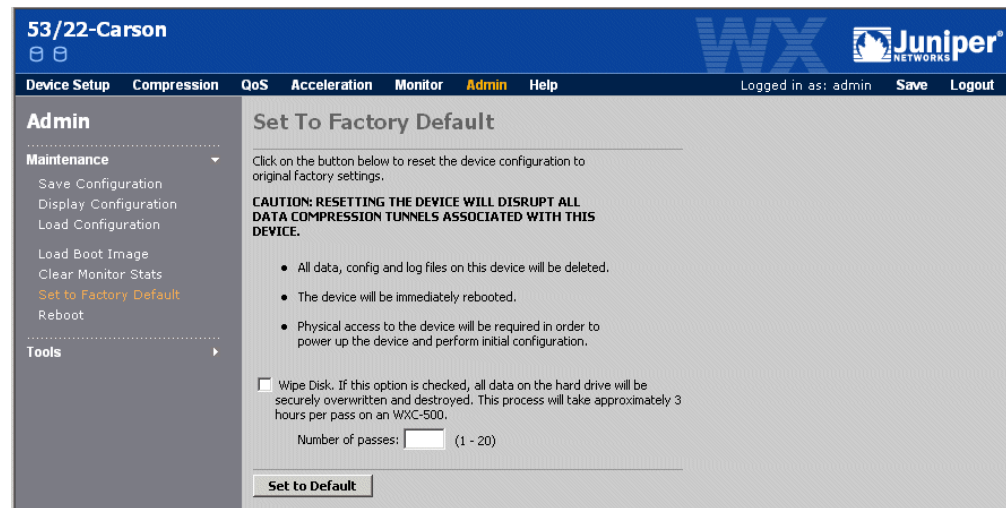
You can erase all device configuration information, including compression statistics and network address information, by restoring the factory default configuration. This is useful when you must move the device to another location.



NOTE: Restoring the factory default configuration removes all data, configuration information and log files. It also disrupts the service 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 283). 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 170: 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

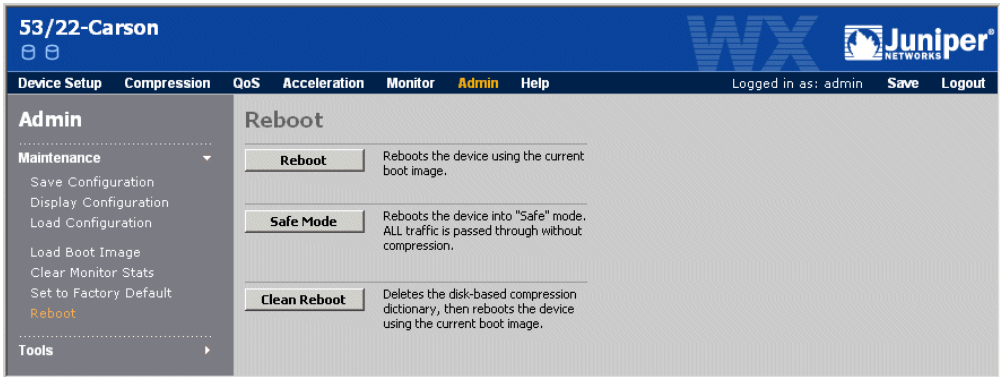
If you load a new boot image of the WXOS software on a device, you must reboot the device to activate the new software. During a reboot, the current boot image (*srs.os*) and the device configuration file (*startup.cfg*) are loaded from Flash memory into main memory. For a WX 15, you can use the CLI to select a specific boot image to be loaded (refer to “reboot” on page 319).

In addition, you can reboot a WX device in Safe Mode so that the power stays on, but traffic is passed through without compression.

To reboot the device:

- 1. Click **Admin** in the menu frame, and then click **Reboot** in the navigation frame.

Figure 171: 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 compression. Note the following: <ul style="list-style-type: none">■ If IPSec is enabled, the default policy may cause traffic to be dropped (refer to “Defining the Default IPSec Policy” on page 238).■ The warning “ERROR SW Passthru” is displayed in the front panel.■ To exit Safe Mode, click Reboot to do a standard reboot.
Clean Reboot	Reboots the device and clears the compression dictionary used for Network Sequence Caching. 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 292
- “Running a Packet Capture” on page 293
- “Generating NetFlow Records” on page 294
- “Entering CLI Commands from the Web Console” on page 295
- “Viewing and Saving System Logs” on page 296
- “Viewing and Saving the Access Control Log” on page 297
- “Exporting Performance Data” on page 298
- “Creating a Diagnostic File” on page 299
- “Viewing Flow Diagnostics” on page 300
- “Viewing the WX 100 Server/Client Summary” on page 303

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 172: Using the Ping Utility

The screenshot shows the Juniper WX Web Console interface. At the top, there is a blue header bar with the text "53/22-Carson" and the Juniper Networks logo. Below the header is a navigation bar with tabs: "Device Setup", "Compression", "QoS", "Acceleration", "Monitor", "Admin" (highlighted), and "Help". On the right side of the navigation bar, it says "Logged in as: admin" and has "Save" and "Logout" buttons. On the left side, there is a "Tools" menu with options: "Ping", "Traceroute", "Packet Capture", "NetFlow", and "Command Line Interface". The "Ping" option is selected. The main content area is titled "Ping" and contains the following fields: "Destination IP address:" with a text input field, "Data Size:" with a dropdown menu showing "32" and the unit "bytes", and "Number of times:" with a dropdown menu showing "3". At the bottom of the form are "Submit" and "Reset" buttons.

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).

4. 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.

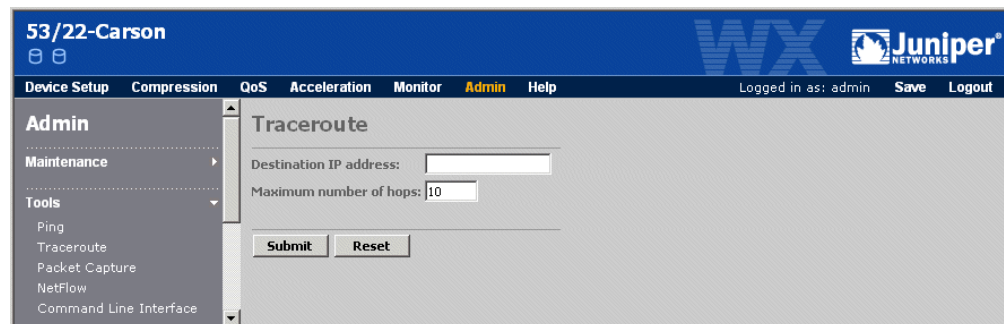
Running a Traceroute to a Network Device

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:

5. Click **Admin** in the menu frame, click **Tools** in the left-hand navigation frame, and then click **Traceroute**.

Figure 173: Using Traceroute to Determine the Network Path to a Device



6. Enter the IP address of the destination device, and the maximum number of router hops (1 to 30) to search for that device.
7. 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". Note the following:

- If tunnel switching is enabled, intermediate decompressed packets are captured that have zeros for the source and destination, and may have checksum errors. These packets are internal to the device and can be ignored.
- If IPsec is enabled, encrypted packets are captured, but decrypted packets are not.

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 174: Using the Packet Capture Utility

The screenshot displays the Juniper WX-10.90.68.200 Packet Capture utility interface. The top navigation bar includes links for Device Setup, Compression, QoS, Acceleration, Monitor, Admin, and Help. The Admin section is active, showing a left-hand navigation menu with options like Maintenance, Tools, and various system logs. The main content area is titled 'Packet Capture' and contains the following sections:

- Instructions:** A text block explaining how to start and stop a packet capture, and how to save or delete the captured data.
- Interface:** A dropdown menu set to 'Local'.
- Size (Bytes):** A text input field with a range of 4096-1499996160.
- Maximum Packets:** Radio buttons for 'All' and a text input field for a specific number of packets.
- Snap Length:** Radio buttons for 'All' and a text input field set to '1514' Bytes.
- Filtering:** Radio buttons for 'Off' and 'On'. Below this are input fields for 'Source' and 'Destination' IP addresses, and a dropdown for 'IP Protocol' set to 'Any'.
- TCP Flags:** Checkboxes for FIN, SYN, RST, PUSH, ACK, URG, ECE, and CWR.
- Storage Format:** A dropdown menu set to 'libpcap'.
- Delete After:** A text input field set to '1' hours.

At the bottom of the main content area, there are buttons for 'Start', 'Stop', 'Save...', and 'Delete'.

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.
Maximum 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.
Filtering	<p>Optionally, select On to limit the packet capture to any combination of the following:</p> <ul style="list-style-type: none"> ■ Enter a source and/or destination IP address and port number ■ Select the TCP or UDP protocol, or select Enter and enter a protocol number (0 to 255) ■ Select one or more TCP flags (applied only to TCP traffic) <p>To populate the filter settings from a current traffic flow, refer to “Viewing Flow Diagnostics” on page 300.</p>
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).

3. To start the packet capture, click **Start**. The status is displayed on the left side 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 175: Generating NetFlow Records

53/22-Carson

Device Setup Compression QoS Acceleration Monitor **Admin** Help

Logged in as: admin Save Logout

Admin

Maintenance

Tools

Ping

Traceroute

Packet Capture

NetFlow

Command Line Interface

Display System Log

Save System Log

Display Access Control Log

Save Access Control Log

Top Traffic > NetFlow

In order to use the NetFlow Export feature the IP Address and Port of the NetFlow server must be entered below.

Enable NetFlow ☐ Yes

IP Address

Port

Submit Reset Cancel

NetFlow(TM) - NetFlow is a Trademark of Cisco Systems, Inc.

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 437.

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 305.

Note the following when entering commands from the Web console:

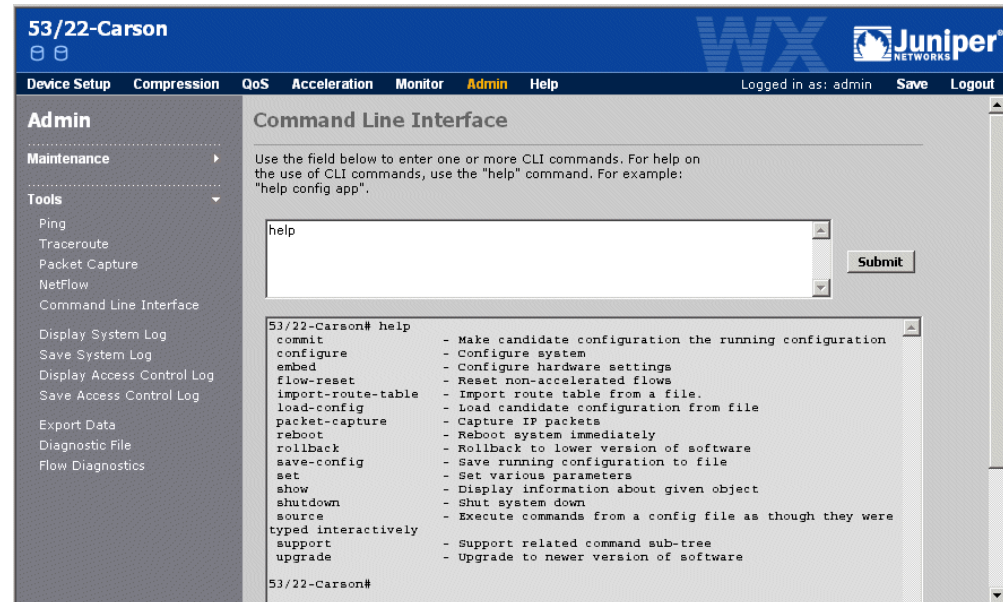
- CLI configuration commands are applied to the candidate configuration when you click **Submit**. Use the “commit” command to apply changes to the running configuration.
- Only “show,” “configure,” and “commit” commands are supported, as well as “ls” and “pwd”. The “cd” command is not supported, nor 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 305.

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 176: 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

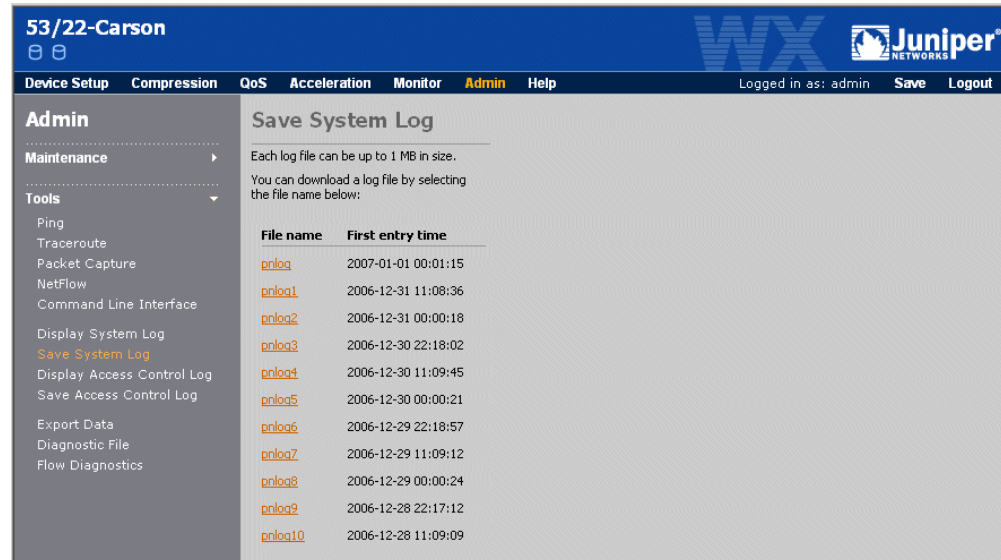
The system log files can be displayed in the Web console. You can also download these log files to a local machine for use by third-party applications. If your network has dedicated syslog servers, you can configure the WX device to send log messages to up to five syslog servers, as described in “Enabling Syslog Reporting” on page 72.

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 177: Saving the System Log file for the Running Configuration



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.

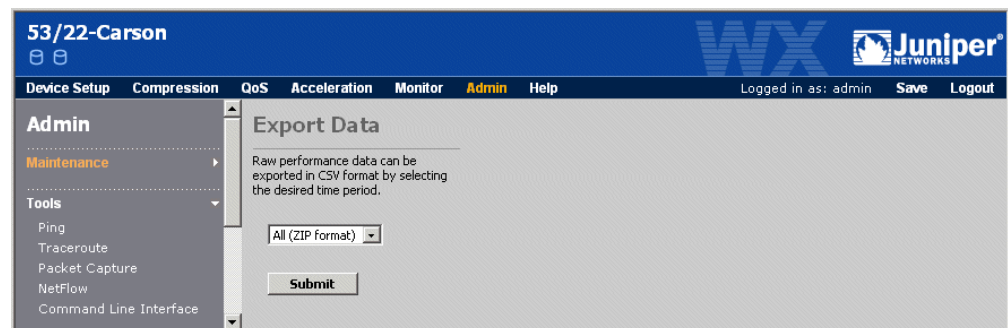
Exporting Performance Data

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 245).

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 178: 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 “Performance Statistics Export” on page 438 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.

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 179: Creating a Diagnostic File

The screenshot shows the Juniper WX 53/22-Carson web interface. The top navigation bar includes links for Device Setup, Compression, QoS, Acceleration, Monitor, Admin (selected), and Help. The left-hand navigation frame shows the Admin menu expanded, with Tools selected. The main content area is titled "Diagnostic File" and contains the following text:

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 Juniper 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@juniper.net.

For faster response, fill out the information below.

The form includes the following fields:

- Name:
- Company:
- Phone:
- Email:
- Problem description:

At the bottom of the form are two buttons: **Submit** and **Reset**.

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 a WX 100 also includes information for the client devices (if any).

Email the diagnostic file as an attachment to support@juniper.net. A support representative will contact you.

Viewing Flow Diagnostics

You can view diagnostic details for up to 50 of the most recently started active traffic flows. You can also initiate a packet capture for a specific flow, and download the top 50 flows to a file in CSV format (for a description of the exported diagnostics, refer to “Flow Diagnostics Export” on page 447).



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.

To view the flow diagnostics:

- 1. Click **Admin** in the menu frame, click **Tools** in the left-hand navigation frame, and then click **Flow Diagnostics**.
- 2. To export the diagnostics for the 50 most recent traffic flows to a file in CSV format, click Download, click **Save**, and specify the name of the . To erase the current traffic statistics, click **Clear**.
- 3. To view the top 50 most recent traffic flows, click **Go**.

Figure 180: Traffic Flow Diagnostics

WX 10.88.8.55Backup server

Device Setup

Compression

QoS

Acceleration

Monitor

Admin

Help

Logged in as: adminSaveLogout

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
- Dagnostic File
- Flow Diagnostics

Flow Diagnostics

This feature allows you to view diagnostic information for a specific flow. To view a list of recent flows, click **Go**. To view flows for a specific source or destination subnet or IP address, enter the subnet (e.g., 12.0.0.0/8) or IP address (e.g., 192.168.1.2/32) in the appropriate fields.

Source Subnet

Destination Subnet

Application

Display

Source Port

Destination Port

Protocol

☒ Show reg. port names

☐ Show domain names

Go

Download...

50 most recent flows (of 65521 found).
To view diagnostic information for a specific flow, click the icon.
To execute a packet capture for a specific flow, click the icon.

Source Address (Source Port)	Destination Address (Destination Port)	Application Name (Application Type)	Protocol	Start Time	Last Active	
10.84.0.10 (domain (53))	10.88.9.30 (4577)	DNS (Default)	TCP	10:33:54.76 JAN 01, 2007	10:33:54.302 Jan 01, 2007	
10.88.8.200 (1033)	10.88.9.100 (3578)	Undefined Applicat... (Default)	TCP	13:16:00.16 DEC 21, 2006	10:33:54.183 Jan 01, 2007	
10.84.0.10 (domain (53))	10.88.9.30 (4575)	DNS (Default)	TCP	10:33:53.297 JAN 01, 2007	10:33:54.23 Jan 01, 2007	
10.84.0.10 (domain (53))	10.88.9.30 (4573)	DNS (Default)	TCP	10:33:53.16 JAN 01, 2007	10:33:53.243 Jan 01, 2007	

4. To view specific traffic flows, specify the following information and click **Go**.

Source Subnet	<p>Enter a subnet to view just the traffic flows from that subnet. The format is:</p> <p>< IP address > / < subnet mask ></p> <p>Where < <i>subnet mask</i> > is the number of bits used for the network portion of the address (such as “10.10.20.0/24”).</p>
Source Port	<p>Enter the source port number of the flows you want to view. An asterisk indicates any port. For a list of common application ports, refer to Appendix , “Common Application Port Numbers”.</p>
Destination Subnet	<p>Enter a subnet to view just the traffic flows sent to that subnet. The format is:</p> <p>< IP address > / < subnet mask ></p> <p>Where < <i>subnet mask</i> > is the number of bits used for the network portion of the address (such as “10.10.20.0/24”).</p>
Destination Port	<p>Enter the destination port number of the flows you want to view. An asterisk indicates any port.</p>
Application	<p>Select an application to view just the traffic flows for the selected application (the default is All).</p>
Protocol	<p>Select an application protocol or select Any to indicate TCP or UDP. You can also type in a protocol number (0 to 134).</p>
Show reg. port names	<p>Click the check box to view the registered names for all ports. Clear the check box to view the names only for port numbers up to 1024.</p>
Show domain names	<p>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 63. The IP address is displayed if its domain name cannot be resolved (the DNS queries may take a few seconds).</p>



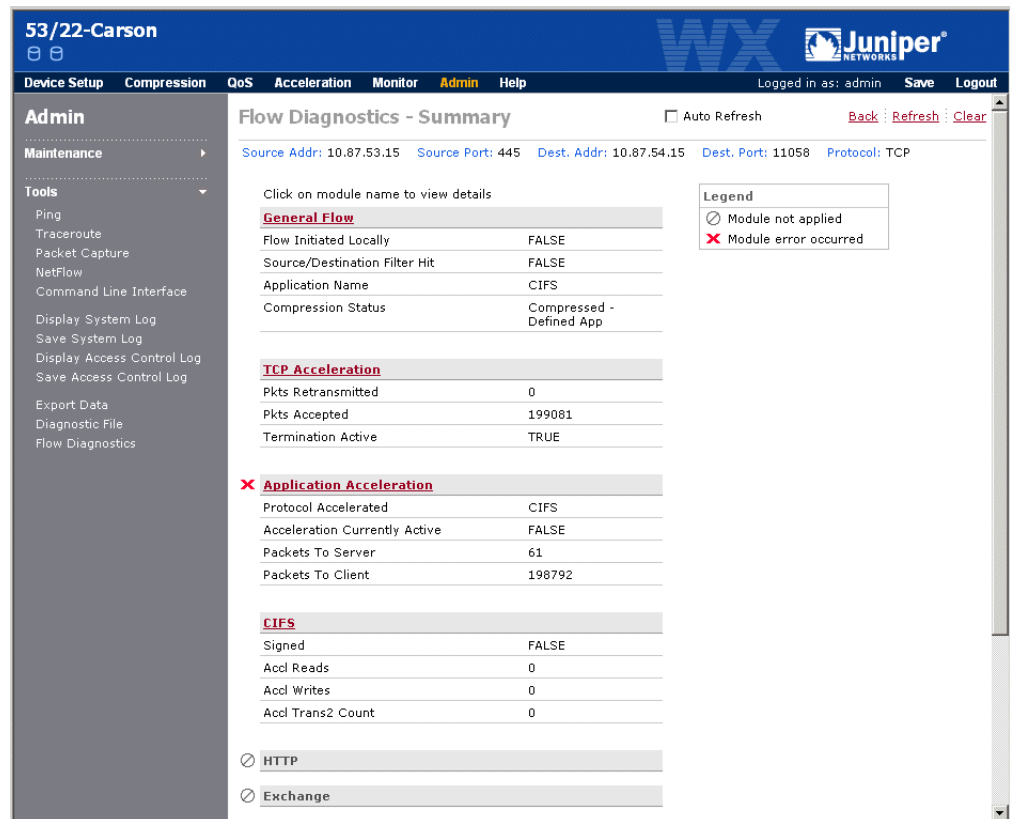
5. To view a summary of the diagnostic details for a traffic flow, click  next to the flow.
6. To open the Packet Capture page with the filtering criteria for a traffic flow, click  next to the flow.

Figure 181: Traffic Flow Diagnostics Summary

The summary details are grouped into the following sections. Note that a **✗** next to any section indicates a problem in that area.

- General Flow
 - TCP Acceleration
 - Application Acceleration
 - CIFS, HTTP, or Exchange
 - Compression
 - Network Sequence Caching
 - QoS
7. Click the **Auto Refresh** check box to update the summary every five seconds.
 8. To view more details related to a specific section, click the section name. Refer to “Flow Diagnostics Export” on page 447 for a description of the flow details provided.

Viewing the WX 100 Server/Client Summary

A 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 a WX 100:

1. Click **Admin** in the menu frame, click **Tools** in the left-hand navigation frame, and then click **Server/Client Summary**.

Figure 182: Viewing the WX 100 Server/Client Summary

55/22-SR100					
Number of Active Clients -- 2					
<div>Device Setup</div> <div>Compression</div> <div>QoS</div> <div>Acceleration</div> <div>Monitor</div> <div>Admin</div> <div>Help</div>					
<div>Admin</div> <div>Maintenance</div> <div>Tools</div> <div>Display System Log</div> <div>Save System Log</div> <div>Display Access Control Log</div> <div>Save Access Control Log</div> <div>Export Data</div> <div>Diagnostic File</div> <div>Flow Diagnostics</div> <div>Server/Client Summary</div>					
Server/Client Summary					
Port	Status	Model	Disk(s)	No. of Tunnels	
				OUT	IN
1	Active	WXC-500	8 GB	0	2
2	Active	WXC-500	8 GB	0	1
Server					
	Active	SR-100		0	0

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"> ■ Active. Client connected and processing traffic. ■ Passive. Client connected, but idle. ■ Not Connected. No client installed.
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 305
- “Logging In Using the CLI” on page 307
- “CLI Basics” on page 307
- “CLI Command Summary” on page 309
- “System-Level Commands” on page 313
- “Configuration Commands” on page 324
- “Show Commands” on page 399



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 306
- “Using a Terminal Connected to the Serial Port” on page 306

You can also access the CLI from the Web console, as described in “Entering CLI Commands from the Web Console” on page 295.

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 4.1 (protocol versions SSHv1 and SSHv2) with DES/3DES encryption. Up to four connections are supported. Multiple channels, port forwarding, and X11 forwarding are not supported.

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 (Figure 183). The serial port is of type RS-232 (AT-compatible) with a male, DB-9 connector. The WX 15 and WX 20 include a crossover cable.

Figure 183: Connecting a Terminal to a WX Device



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 **juniper** 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 388.

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 10 summarizes the keyboard shortcuts.

Table 10: 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.
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 11 provides a summary of the available CLI commands. Note that the “ping” and “traceroute” commands can be entered at any level

Table 11: CLI Command Summary

Command	Description
System-Level Commands	
“activate” on page 313	Set various logging, help, and display options.
“commit” on page 313	Apply configuration changes to the running configuration.
“configure” on page 313	Enter configuration mode.
“copy” on page 314	Copy files to or from a WX device.
“embed” on page 314	Enable or disable hardware passthrough.
“flow-reset” on page 314	Reset CIFS traffic flows so that acceleration can be applied.
“import-route-table” on page 315	Import a Cisco routing table from an FTP or TFTP server.
“list” on page 315	List the files in a specific directory on the device.
“load-config” on page 316	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 317	Capture raw network traffic on the Local or Remote interface.
“ping” on page 318	Verify connectivity with other network devices.
“reboot” on page 319	Reload the WXOS software on the device.
“remove” on page 320	Delete files from the device.
“reset” on page 320	Reset tunnel balancing information for WX 100 clients
“rollback” on page 320	Revert to an earlier version of WXOS.
“save-config” on page 321	Save the running configuration to Flash memory to preserve committed changes the next time the device is rebooted.
“set” on page 322	Change the device name, location, or administrator contact information.
“shutdown” on page 322	Turn off the device in preparation for disconnecting the power.
“source” on page 323	Execute CLI commands stored in a file.
“support” on page 323	Generate a diagnostic file for Technical Support.
“traceroute” on page 324	Trace the network path to a remote device.
“upgrade” on page 323	Load a new (later) version of WXOS from an FTP or TFTP server.
Configuration Commands	
“configure aaa” on page 324	Define user accounts locally, and specify how users are authenticated.
“configure acceleration” on page 326	Configure the various methods of acceleration WAN traffic.
“configure application” on page 332	Define application definitions, assign them to traffic classes, and enable or disable compression and acceleration by application.
“configure arp” on page 335	Add static ARP entries or clear dynamic entries.
“configure backup” on page 336	Configure a WX device as a backup for one or more primary devices.
“configure clock” on page 338	Set the device time, time zone, and enable or disable daylight savings time.
“configure console” on page 339	Set the baud rate of the DB9 console port.

Command	Description
“configure dns” on page 339	Specify the DNS servers and domain name used to resolve IP addresses on the Traffic report.
“configure event” on page 340	Specify the performance and system events that can be generated.
“configure filter” on page 342	Exclude applications or address pairs from data compression, and specify whether fragmented packets are compressed.
“configure flow-reset” on page 343	Reset CIFS traffic flows after each reboot so that acceleration can be applied.
“configure interface” on page 344	Specify interface speeds and duplex modes, run a test to detect a mode mismatch, enable the compression 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 345	Change the device IP address, subnet mask, or default gateway.
“configure ipsec” on page 346	Configure IP security (IPSec) to authenticate and encrypt traffic between WX devices in the same community.
“configure license” on page 349	Enter a new license key.
“configure log” on page 350	Set logging levels and output log messages to the system console.
“configure mon-apps” on page 351	Specify the applications to be monitored on reports.
“configure multi-path” on page 352	Configure primary and secondary paths between WX devices.
“configure ospf” on page 356	Use OSPF to discover routes on the LAN side of the WX device.
“configure packet-interception” on page 357	Configure RIP, WCCP, or external routing to route traffic to an off-path device.
“configure path-mtu-discovery” on page 361	Configure automatic discovery of the MTU size for each tunnel.
“configure prime-time” on page 361	Specify the days and times during the week when network performance is most important.
“configure profile-mode” on page 362	Enable Demo Mode and configure virtual devices for non-intrusive testing.
“configure qos inbound” on page 363	Configure inbound QoS settings.
“configure qos outbound” on page 365	Configure outbound QoS settings.
“configure radius” on page 370	Specify RADIUS servers and server groups used to authenticate WX users.
“configure reduction” on page 371	Configure data compression settings.
“configure reduction-subnet” on page 378	Specify the subnets advertised to other WX devices for compression.
“configure reg-server” on page 380	Configure the WX device that acts as the registration server.
“configure remote-routes” on page 382	Specify how often remote routes are fetched from other devices, and whether each remote route is validated.
“configure rip” on page 383	Use RIP to discover routes on the LAN side of the WX device.
“configure route” on page 384	Add static routes to the routing table, enable router load balancing, and specify the ICMP age-out interval.
“configure route-poll” on page 387	Obtain dynamic routes by periodically polling a Cisco router
“configure security” on page 388	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 389	Enable or disable SNMP, generate traps, and specify community strings.
“configure sntp” on page 390	Specify primary and secondary SNTP servers to maintain the device time.
“configure ssl certificate” on page 391	Import SSL certificates to optimize SSL traffic.
“configure ssl optimization” on page 392	Enable SSL optimization and reset optimization statistics.
“configure stack-group” on page 392	Configure a WX 100 to support client devices.
“configure syslog” on page 394	Send syslog messages to one or more syslog server

Command	Description
“configure system” on page 394	Specify the topology type and community size.
“configure tacplus” on page 395	Specify TACACS+ servers used to authenticate WX users.
“configure top-talker” on page 396	Export top traffic statistics to a file and/or send to a NetFlow server.
“configure wan-performance-monitor” on page 397	Configure WAN performance monitoring to measure latency and loss between the current device and one or more remote WX devices.
Show Commands	
“show aaa” on page 399	Display authentication methods for the Console, Web, and CLI.
“show acceleration” on page 399	Display application acceleration configuration.
“show access-log” on page 400	Display management access log.
“show all” on page 400	Display all system configuration information.
“show application” on page 401	Display application definition.
“show arp” on page 401	Display ARP entries.
“show backup-sr” on page 401	Display backup mode configuration.
“show clock” on page 401	Display time related parameters.
“show connection” on page 402	Display list of current connections to remote WX devices.
“show console” on page 402	Display console (serial) port parameters.
“show contact” on page 402	Display contact information for the device.
“show dns” on page 402	Display DNS server addresses and default domain name.
“show event” on page 403	Display performance event definitions and the enabled system events.
“show filter” on page 403	View compressed applications and excluded address pairs.
“show flow-details” on page 404	View details of a specific traffic flow.
“show flow-reset” on page 405	View configuration and status of a flow reset.
“show import-route-table” on page 405	Display import route table information.
“show interface” on page 405	Display network interface parameters.
“show ip” on page 406	Display the IP parameters.
“show ipsec” on page 406	Display IPSec configuration and security associations.
“show license” on page 407	Display license information.
“show location” on page 407	Display location description for this system.
“show log” on page 407	Display system log.
“show mon-apps” on page 407	Display list of monitored applications.
“show multi-path” on page 408	Display multi-path configuration.
“show ospf” on page 408	Display OSPF settings.
“show packet-capture” on page 408	Display packet capture settings.
“show packet-interception” on page 409	Display packet interception parameters for off-path devices.
“show path-mtu-discovery” on page 410	Display MTU discovery settings.
“show prime-time” on page 410	Display prime time settings.
“show profile-mode” on page 410	Display Demo Mode settings.
“show qos inbound” on page 411	Display QoS settings.
“show radius” on page 411	Display RADIUS configuration.

Command	Description
“show reduction” on page 412	Display compression status.
“show reduction-subnet” on page 413	Display compression subnet status.
“show reg-detail” on page 413	Display detailed information about the registration database.
“show reg-server” on page 414	Display registration server parameters.
“show reg-server” on page 414	Display summary information about the registration database.
“show remote-routes” on page 414	Display remote route information.
“show rip” on page 415	Display RIP parameters.
“show route” on page 415	Display routing table.
“show route-poll” on page 415	Display routing poll table.
“show security” on page 416	Display security related parameters.
“show snmp” on page 416	Display SNMP related parameters.
“show sntp” on page 416	Display Sntp related parameters.
“show ssl certificate” on page 416	Display list of imported SSL certificates or a specific certificate.
“show ssl optimization” on page 416	Display SSL optimization configuration and statistics.
“show stack-group” on page 417	Display model and status of clients connected to a WX 100.
“show syslog” on page 417	Display syslog parameters.
“show system” on page 418	Display general system information.
“show system-name” on page 418	Display the system name.
“show tacplus” on page 418	Display TACACS + server settings.
“show top-talker” on page 419	Display top-talker parameters.
“show uptime” on page 419	Display system uptime.
“show version” on page 419	Display version information.
“show wan-performance-mon” on page 419	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.

activate

The `activate` command enables or disables various logging, help, and display options.

`activate [no] <dashes | errorhelp | log | more | | retries | usedParams | verbose>`

Where:

- **dashes.** Displays dashes on error lines, rather than spaces (disabled by default).
- **errorhelp.** Displays command help even if help is invoked incorrectly, such as by entering “help” after the command (disabled by default).
- **log.** Enters CLI commands in the system log (enabled by default).
- **more.** Displays one page of output at a time and prompts you to press any key to continue (enabled by default).
- **retries.** Redisplays a command entered incorrectly and displays a caret (^) below the error (enabled by default).
- **usedParams.** When viewing online help for a command, all options already entered are ignored (enabled by default).
- **verbose.** Displays help for all command options (enabled by default).

Use the `no` keyword to disable an option. For example, to disable the pause between each page of output:

`activate no more`

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 WXC 590 or the latest 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. On a WX 100, this command is available only on non-fiber versions that have the network connectors on the front panel.

flow-reset

A traffic flow cannot be accelerated unless the WX sees the start of the flow. By default the flow-reset mode is on so that eligible CIFS traffic flows are reset if a packet is received for the flow within 900 seconds (15 minutes) of the tunnel establishment time. This allows established CIFS traffic flows to be accelerated each time the WX is restarted. To enable or disable flow-reset mode or change the default duration, refer to the “configure flow-reset” on page 343.

To manually start a flow reset for a specified number of seconds (5 to 86400) independent of the tunnel start time:

```
flow-reset start duration <seconds>
```

To stop the flow reset at any time:

```
flow-reset stop
```

To view the configuration and/or status of the flow reset:

```
show flow-reset [configuration | status]
```

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).

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 286.



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 321
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 service 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 283).

- **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.



NOTE: If tunnel switching is enabled, running a packet capture will capture intermediate decompressed 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. You cannot run a packet capture until the previous one is deleted. To view the status of the last packet capture (if any):

```
show packet-capture
```

To delete the previous packet capture:

```
packet-capture delete
```

2. To start a packet capture:

```
packet-capture start interface <local | remote | both> size <number>
[packets <number>] [format <libpcap | snoop>] [snaplen <max size>] [savetime
<time>] [hosts <source_IP>,<dest_IP>] [ports <source_port> , [<dest_port>]]
[ip-proto <number | tcp | udp>] [tcp-flags <tcp-fin, tcp-syn, tcp-rst,tcp-psh,
tcp-ack,tcp-urg,tcp-ece,tcp-cwr>]
```

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.

- **hosts <source_IP, dest_IP>** . Source and destination IP addresses of the traffic to be captured (default is all IP addresses). If just the source IP is specified, the destination address defaults to “any”.
 - **ports <source_port, dest_port>** . Source and destination port numbers of the traffic to be captured (default is all ports). If just the source port is specified, the destination port defaults to “any”.
 - **ip-proto <number | tcp | udp>** . IP protocol of the traffic to be captured. Specify a protocol number (0 to 255), “tcp”, or “udp” (default is all protocols).
 - **tcp-flags <tcp-fin, tcp-syn, tcp-rst, tcp-psh, tcp-ack, tcp-urg, tcp-ece, tcp-cwr>** . TCP flags required on packets to be captured (applies only to TCP traffic). Multiple flags must be separated by commas.
3. To stop a packet capture:
- ```
packet-capture stop
```
4. 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.

## ***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 compression.

1. To immediately reboot the system:

```
reboot [-all | -client-id <0-6>] [-safe-mode] [-no-resync]
```

Where:

- **-all.** Reboots a WX 100 stack server and all of its client devices. Does not support the safe-mode option.
- **-client-id <0-6> .** Reboots all of a 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 any processing. Note that the warning “ER SW Passthru” is displayed in the front panel of devices that have an LCD.
- **-no-resync.** Reboots the device without saving the compression dictionary used for Network Sequence Caching. Available only on WXC devices.

Type “y” to confirm that you want to reboot.

2. On a WX 15, WXC 250, or WXC 500 (version 3.0), 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 323). 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.

## **reset**

The **reset** command lets you reset the information used to distribute the tunnels across the client devices of a WX 100 stack server, and to reset the status of a new disk drive on a WXC device.

1. To reset the general load balancing information:

```
reset stack-group tunnel-lb-info
```

To reset the tunnel-preference information for WXC clients:

```
reset stack-group tunnel-pref-info
```

Use these commands only as part of the procedure to rebalance the tunnels assigned to each client (refer to “Distributing Tunnels Across Client Devices” on page 53).

2. After replacing a disk drive on a WXC device, enter the following command to activate the new drive:

```
reset disk status
```

## **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. Before you rollback to a previous version of WXOS, note the following device requirements:

- WX 20s require WXOS 3.0 or greater.
- WX 15s, WX 100s, and WXC 500s require WXOS 5.0 or greater.
- WX 60s require WXOS 5.0.10 or greater.
- WXC 250s require WXOS 5.0.8 or greater.
- WXC 590s require WXOS 5.2.4 or greater.



- 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 283.



**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 system information, such as 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>
```

5. On a WX 100 stack server, to enable or disable the use of packet counts to load-balance the tunnels across the client devices (enabled by default):

```
set stack-group tunnel-lb-pkt-count <on | off>
```

## 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 a WX 100 stack server and all of its client devices. Does not support reset options (everything is reset).
- **-client-id <0-6>**. Shuts down all of a 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 299.

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 320).

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 370.

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 indicate whether a user can run packet capture (disallowed by default):

```
user packet-capture name <name> <allow | disallow>
```

To change a user's idle timeout and/or access level (a zero indicates no idle timeout):

```
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 203.

- **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.
- **TCP Acceleration.** 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 WX, between the two WX devices, and between the receiving WX 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, TCP Acceleration 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 or 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 TCP Acceleration to be enabled.

To accelerate traffic between two WX devices, the following conditions must be met:

- The selected applications must be compressed, and a service 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 365).
- For TCP Acceleration, clustering must be enabled if the outbound and return traffic does not always traverse the same two WX devices.



**NOTE:** Acceleration is most effective in networks with high-speed connections and high latency. It may have no effect if traffic must traverse high-latency or low-speed connections that are 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) compressed (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 TCP Acceleration 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 FEC, FCS, and TCP Acceleration, 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 Packet Flow Acceleration 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 acceleration statistics for CIFS, Exchange, or both (default is all):</p> <pre>reset-stats [all   cifs   exchange]</pre> <p>To clear the counters shown on the Forward Error Correction report:</p> <pre>forward-error-correction clear-counters</pre> |

#### Endpoints

You can enable the same Packet Flow Acceleration methods for all endpoints or add each endpoint and specify its methods (all methods are disabled by default).

#### Enabling Specific Endpoints

To add an endpoint, a service tunnel must exist for the device and outbound QoS must be configured correctly:

```
endpoint add ip-address <IP address> [mode <on | off>] [active-flow-pipelining
<on | off>] | {[fast-connection-setup <on | off>] [forward-error-correction <on |
off>] [data-pkts <4-25>] [recovery-pkts <0-5>]]}
```

Where:

- **mode <on | off>**. Enables or disables the endpoint for Packet Flow Acceleration (disabled by default).
- **<method> <on | off>**. Enables or disables a specific method. TCP Acceleration cannot be used with Fast Connection Setup.
- **data-pkts <4-25> recovery-pkts <0-5>**. 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.

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.

To change the settings for an endpoint:

```
endpoint set ip-address <IP address> [mode <on | off>] [active-flow-pipelining
<on | off>] | {[fast-connection-setup <on | off>] [forward-error-correction <on |
off>] [data-pkts <4-25>] [recovery-pkts <0-5>]]}
```

To remove an endpoint from the list shown by the “show acceleration” command:

```
endpoint remove <IP address>
```

In the Web console, the endpoint is disabled, but is not deleted.

#### Enabling All Endpoints

To use the same 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 service tunnel exists and outbound QoS is configured correctly. The specified methods are applied to all qualifying endpoints, and to all qualifying endpoints added to the same community in the future.




| Setting               | Commands                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Applications          | <p>To specify the applications that use Fast Connection Setup:<br/> <b>fast-connection-setup application add &lt;name&gt;</b></p> <p>To add an application that is included or excluded from TCP Acceleration:<br/> <b>active-flow-pipelining application add &lt;name&gt;</b></p> <p>To indicate whether the specified applications are included or excluded from TCP Acceleration (included by default):<br/> <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:<br/> <b>active-flow-pipelining application remove &lt;name&gt;</b><br/> <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 (enabled by default):<br/> <b>cifs set mode &lt;on   off&gt;</b></p> <p>To enable SMB signing when SMB signing is required (disabled by default):<br/> <b>cifs set apply-signing &lt;on   off&gt;</b></p> <p>To disable SMB signing when it is not required (enabled by default):<br/> <b>cifs set disable-signing &lt;on   off&gt;</b></p> <p>To downgrade SMB2 to SMB on flows between Vista and non-Vista devices so the flows can be accelerated (enabled by default):<br/> <b>cifs set downgrade-smb2 &lt;on   off&gt;</b></p> <p>To accelerate a CIFS application (must be an application of type CIFS):<br/> <b>cifs application add &lt;name&gt;</b></p> <p>Note that acceleration is enabled for all remote endpoints for which TCP Acceleration is enabled.<br/> To disable acceleration for a CIFS application:<br/> <b>cifs application remove &lt;name&gt;</b></p> <p>If you enable SMB signing, specify the user name used to access the Windows server, and enter a password at the prompt:<br/> <b>cifs apply-signing add username &lt;name&gt;</b></p> <p>To enter a Windows domain name for the account used to access the server (optional):<br/> <b>cifs apply-signing add domain &lt;name&gt;</b></p> <p>To delete the account used to access the Windows server:<br/> <b>cifs apply-signing delete username &lt;name&gt;</b></p> <p>For more information about CIFS acceleration, refer to “Microsoft CIFS and Microsoft Exchange Acceleration” on page 215.</p> |
| Exchange Acceleration | <p>To enable or disable Exchange acceleration (disabled by default):<br/> <b>exchange set mode &lt;on   off&gt;</b></p> <p>To accelerate an Exchange application (must be an application of type Exchange):<br/> <b>exchange application add &lt;name&gt;</b></p> <p>Note that acceleration is enabled for all remote endpoints for which TCP Acceleration is enabled.<br/> To disable acceleration for a Exchange application:<br/> <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 215.</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 TCP Acceleration is enabled. 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 WX 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 216.</p> |

TCP Acceleration Clusters (asymmetric routing support)

For TCP Acceleration 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:
- All devices in the same cluster must be the same model, such as all WXC 500s, and they must all have the same version of WXOS. WX 100s need not have the same number of clients.
  - 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, service 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 decompressors 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 <IP-address1 IP-address2 ...></pre> <p>To remove one or more devices from the cluster:</p> <pre>cluster remove <IP-address1 IP-address2 ...></pre>
TCP Acceleration buffers	<p>For optimum performance of TCP Acceleration, 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"> ■ small. 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. ■ medium. 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. ■ large. Recommended for links with speeds of 512 Kbps (or higher) and round-trip-times from 600 to 1200 ms. Not recommended for slower links. ■ huge. Recommended for links with speeds of 512 Kbps (or higher) and round-trip-times above 1200 ms. Not recommended for slower links. <p>NOTE: Larger buffer sizes use more memory, but smaller buffer sizes may restrict throughput. The recommended buffer sizes allow up to about 100 Mbps of uncompressed throughput.</p>
Heartbeats	<p>On a high-loss link, compression may be terminated if heartbeat packets are lost. By default, when TCP Acceleration or FEC is enabled for a remote endpoint, the local device stops compressing data for the remote device if it fails to respond to 15 consecutive heartbeats (passthrough mode). If 30 consecutive heartbeats get no response, the service 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 compression and disconnect the tunnel:</p> <pre>set heartbeat-misses passthru <number default> disconnect <number default></pre> <p>This setting applies only to remote endpoints for which TCP Acceleration or FEC are enabled. To change the heartbeat settings for all other endpoints, refer to “configure reduction” on page 371.</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 TCP Acceleration sessions (“default” is 1460):</p> <pre>active-flow-pipelining set mss-override <64-1460 default></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 compression. To conserve system capacity, you should disable compression for applications whose traffic is encrypted or already compressed.
- Enable TCP acceleration (if data compression 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 compression 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 96.

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 326), 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 indicate whether an application uses SSL encryption (default is false, but must be true to enable SSL optimization):
- set name <name> ssl-encapsulated <true | false>**
- Note that the application type must be **Default** for applications that use SSL.
5. 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.juniper*.net

<uri> is up to eight strings separated by slashes. An asterisk (*) by itself indicates any string. For example:

www.juniper*.net/*/index.htm

Note that an asterisk is treated as a single character when it is part of another string, such as “www.juniper*.net”.

- **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).

6. 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>
```

7. To delete an application definition:

```
remove <application name>
```

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 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. If you have multiple static routes with different gateways (up to four), you can lower the ARP timer to reduce the number of minutes before a failover occurs to the next gateway when the current gateway becomes unavailable.

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 change the ARP gateway timer for all static routes (default is 20 minutes):

```
config arp set gateway-refresh-frequency all <1-9 or 10>
```

Enter 1 to 9 minutes or 10 to reset the timer to the default. This is useful if you are using load balancing with static routes or have multiple redundant routes via different gateways. Enter the **show route** command to view the current timer setting (a 10 indicates the default).

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 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 compression 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 184 shows a deployment where the backup can advertise the same compression subnets as the primary, even though it is not in the same data path.

Figure 184: Sample Backup Deployment



- To back up an off-path device, the compression 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, verify that all devices belong to the Default community (the backup, primary, and all of the primary's remote endpoints). When backup mode is enabled on a device, the device is automatically removed from all communities, except the Default community.
3. On the primary devices to be backed up:
 - a. Verify that the primary and backup devices have the same topology setting (click Device Setup > Advanced > Topology). Hub is recommended.
 - b. Do not enable compression for "ALL discovered WX devices" on the Endpoints page (click Compression > Endpoints).
 - c. Do not enable link failure propagation on the Local or Remote interface of a primary device unless these settings also apply to the backup device (click Device Setup > Interfaces).
4. On each remote WX device, if compression, 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 compression is enabled for "ALL discovered WX devices", then an outbound service tunnel to the backup device is formed automatically when the backup becomes active.
5. Disable load balancing on all devices that compress data for a primary device.
6. On the backup device:
 - a. Verify that the primary and backup devices have the same topology setting (click Device Setup > Advanced > Topology). Hub is recommended.
 - b. Do not advertise any subnets for compression (click Compression > Compression Subnets, and clear the check box for the local subnet).

- c. Specify a primary device supported by the backup device:

```
config backup-sr remote-sr add <IP address>
```

- d. To enable or disable backup mode (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 compression, this step is unnecessary if the new device enables compression for “ALL” devices.

- e. To remove a primary device:

```
config backup-sr remote-sr remove <IP address>
```

- f. To view the current backup configuration:

```
show -run backup-sr
```

- g. 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 ntp” on page 390. 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, type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

configure event

You can configure performance thresholds for compression, acceleration, throughput, and dropped traffic so that events are triggered when the average performance for the previous hour or day exceeds (or drops below) the specified threshold. You can also enable or disable the generation of SNMP traps and syslog messages for system events, such as login failures.

To view the performance and system events that have occurred, refer to “Events Summary” on page 280. Performance events are also sent to any SNMP trap destinations and syslog servers that you have defined.

1. To view the current performance event definitions (if any) and/or the system events that are currently enabled:

```
show -run event [configuration <all | performance | system>]
```

2. Type the following command to enter the configure event mode:

```
config event
```

3. To add a performance event definition:

```
performance add metric <name> type <type> value <number> [reference <name>]
[destination <IP_address | <non-WX>] [period <daily | hourly>] [severity <level>]
[mode <enabled | disabled>] [app-mode <all | any | default>] [dest-mode <all | any
| default>] [prime-time-mode <enabled | disabled>]
```

Where:

metric <name>	<p>Name of the metric. You can create multiple event definitions for each metric. Table 4 on page 142 describes how each metric is calculated.</p> <ul style="list-style-type: none"> ■ comp. Compression (%) ■ comp-thru-out. Outbound compressed throughput (Kbps) ■ tcp-accel. TCP Acceleration (%) ■ tcp-accel-thru-in. Inbound TCP acceleration throughput (Kbps) ■ app-accel. Application acceleration (%) for CIFS, Exchange, HTTP ■ wan-thru-in. Inbound WAN throughput (Kbps) ■ wan-thru-out. Outbound WAN throughput (Kbps) ■ qos-thru-out. Outbound QoS throughput (Kbps) ■ bytes-dropped-out. Number of outbound bytes dropped ■ packet-dropped-out. Number of outbound packets dropped
type <type>	<p>Indicates whether the threshold value is an absolute value or relative to the average performance for the past seven days (absolute-above, absolute-below, relative-above, relative-below).</p>
value <number>	<p>Event threshold value (the units depend on the specified metric). For example, to generate an event if compression falls below 80 % of the average for the past seven days, specify a relative-below type for the comp metric, and enter “80” for the threshold value.</p>
reference <name>	<p>Name of a specific application or traffic class to be monitored. Use this option only if you specify “default” for the app-mode. Note that all metrics apply to applications, except for QoS Throughput, Bytes Dropped, and Packets Dropped:</p> <p>For the Application Acceleration metric, be sure to specify a CIFS, Exchange, or HTTP application (all others have no effect).</p>

destination < IP_address < non_WX_name >	IP address of a specific WX to be monitored, or the name of a non-WX endpoint, such as “Other Traffic” (non-WX endpoints apply only to the WAN and QoS Throughput, Bytes Dropped, and Packets Dropped metrics). Use this option only if you specify “default” for the dest-mode .
period < daily hourly >	Indicates whether performance is evaluated at the end of each hour or once a day at midnight (default is daily). A new event is triggered for each hour or day that violates the threshold.
severity < level >	Severity of events triggered by this definition. The corresponding severity level used on syslog events is shown in parentheses. <ul style="list-style-type: none"> ■ ok (Notice) ■ warning (Info) ■ major (Error) ■ critical (Critical)
mode < enabled disabled >	Enable or disable the generation of events for this definition (enabled by default).
app-mode < all any default >	Indicates whether one or all applications can trigger an event: <ul style="list-style-type: none"> ■ all. An event occurs based on the overall performance of all applications or traffic classes (the default). ■ any. An event occurs if any application or traffic class violates the specified threshold. ■ default. An event occurs only if the application or traffic class specified by the reference keyword violates the threshold.
dest-mode < all any default >	Indicates whether one or all endpoints can trigger an event: <ul style="list-style-type: none"> ■ all. An event occurs based on the overall performance of all remote WX endpoints (the default). ■ any. An event occurs if any WX endpoint violates the specified threshold. ■ default. An event occurs only if the endpoint specified by the destination keyword violates the threshold.
prime-time-mode < enabled disabled >	Limit the generation of events for this definition to prime time days and hours (disabled by default). No events are generated if you enable prime-time mode, but prime time is not defined (refer to “Defining the Prime Time” on page 115).

To change a performance event definition, specify the definition ID number and the other event settings you want to change. Use the **show event** command to view the definition IDs.

```
performance set id <number> [metric <name>] [type <type>] [value <number>]
[reference <name>] [destination <IP_address | <non-WX>] [period <daily | hourly>]
[severity <level>] [mode <enabled | disabled>] [app-mode <all | any | default>]
[dest-mode <all | any | default>] [prime-time-mode <enabled | disabled>]
```

If you change the metric name, verify that the threshold value and other settings are appropriate for the new metric.

To delete a performance event definition:

```
performance remove id <number>
```

4. To enable or disable the generation of a system event (all system events are enabled by default):

```
system set name <name> mode <enabled | disabled>
```

To view the system event names, enter a “?” after the **name** keyword. For a description of each system event, refer to “SNMP Traps and Syslog Messages” on page 427.

configure filter

By default, all applications running over TCP or UDP (except Groupwise, HTTPS, SMTP, SSH, and Traceroute) are enabled for data compression. The **filter** command lets you specify the applications, protocols, or source and destination address pairs to be compressed. You can also disable the compression 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 compression will be minimal. Note that applications must be defined before they can be filtered. To create application definitions, refer to “configure application” on page 332. Undefined applications are compressed by default.

Note the following:

- If you disable data compression 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 compression:

```
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 compression:

```
set mode-applications <off | include | exclude>
```

Set the mode to “off” to compress 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 compression:

```
set mode-address-pair <off | include | exclude>
```

Set the mode to “off” to compress 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 compression of packet fragments:

```
set ip-fragments {on | off}
```

All packet fragments are compressed by default. Fragments may not be associated with the correct application, but disabling compression may cause fragments to arrive before the compressed 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 flow-reset

A traffic flow cannot be accelerated unless the WX sees the start of the flow. To reset eligible CIFS traffic flows after each reboot so that acceleration can be applied (enabled by default):

```
configure flow-reset [mode <on | off>] [duration <seconds>]
```

Eligible traffic flows are reset if a packet is received for the flow within the specified number of seconds (5 to 86400) from the time the tunnel is initially established. The default is 900 (15 minutes).

To view the configuration and/or status of the flow reset:

```
show flow-reset [configuration | status]
```

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 compression 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, WXC 500, and WXC 590):

```
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 compression 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 compressed 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 compression subnets, enable the WAN compression subnet option (refer to “configure reduction-subnet” on page 378).

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 380).

configure ipsec

IP security (IPSec) can be used to authenticate and encrypt traffic between a 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 352). 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).



NOTE: IPSec is NOT supported on a WX 100 stack server with one or more clients.

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 229). The Setup Wizard updates the Wizard template.

To change the IPSec settings:

1. To view the current IPSec settings:

```
show -run ipsec [application-filter | sa [<ip-address>]]
```

Where:

- **application-filter.** Displays the applications that require or never use IPSec.

- **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 specify applications that require or never use IPsec (by default, all applications use IPsec when IPsec is available):

```
application-filter <app_name> tunnel-in-ipsec <if-configured | no | required>
```

If IPsec is required for an application, traffic is dropped if IPsec is unavailable.

5. 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.

6. 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.

7. 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]
```

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 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 compression. 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:

https://www.juniper.net/generate_license

If you have any problems with the licensing process, open a case with the Juniper Case Manager at <http://www.juniper.net/cm>. To call from the United States, Canada, or Mexico, dial + 1-888-314-JTAC. To call from other locations, check the list of local support centers at http://www.juniper.net/support/support_contacts.html or dial + 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 https://www.juniper.net/generate_license to obtain a license key.
 - b. On the WX device, use the following command to enter the license key:

`config license set license-key <new license key>`

configure log

Log messages can be output to the terminal connected to the WX console port, and default logging levels can be changed.

1. To enable or disable the output of log message to the system console (disabled by default):

`config log set console-output <on | off>`

2. To set the severity of the messages written to the system log:

`config log set severity {default | <module>=ceidovan | default}`

Where:

- **default.** The default logging levels are “cei” (see below).
- **< module > .** Setting logging levels by module is intended primarily for diagnostic purposes.
- **c.** Critical error messages about software or hardware malfunctions.
- **e.** Error messages, such as License expired.
- **i.** Informational messages, such as reload requests and low-process stack messages.
- **d.** Debug messages.

- o. Notices about unusual events that are not errors.
 - v. Verbose mode, which provides additional details the above messages.
 - a. All messages.
 - n. None.
3. To commit the changes to the running configuration, type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

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 performance monitoring reports (refer to “configure wan-performance-monitor” on page 397). For more information about monitoring statistics, refer to “Monitoring and Reporting” on page 245.

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 332.



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-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 service 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 compression must be enabled on each WX device that supports Multi-Path (refer to “configure reduction” on page 371). Encryption can be enabled for one or both paths (refer to “configure ipsec” on page 346).

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 137).

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 365)
- WX control packets (refer to “configure reduction” on page 371)

Also, multi-path DSCP values override ToS type-of-service settings used for Cisco router balancing (refer to “configure route” on page 384)

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 365.
- **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).

```
===== 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 compression.

- **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 compressed and uncompressed traffic is sent through the service 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 123.
 - The WX accepts any combination of GRE and Layer 2 (L2) encapsulation for forwarded and return traffic. L2 takes precedence if offered by the router, provided the WX is directly connected to the router at Layer 2.
 - For high-availability environments, service groups using multicast are supported where one or more routers can load-balance traffic across multiple off-path WX devices.
- **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 127.

In each case, the redirected traffic is compressed (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. Enter the following command to view the current packet interception settings. In this example of a WCCP service group, the “Assigned Hash” is used by the router to load-balance traffic across the three WX devices in the group.

```
show -run packet-interception
```

```
Packet interception: WCCP
Router IP: 224.1.1.100
Priority: 230
WCCP auth type: none
IP Protocol  Service ID  Forward  Return  Changes  Designated  State

TCP          51      L2      GRE      3        Yes  Connected
UDP          55      GRE     GRE      1        No   Connecting
```

WCCP Service Groups:

```
Service Identifier: 51
Number of routers: 1
Router 0 IP:      10.87.122.254
MAC address:      00:d0:bc:f4:04:7c
Forwarding:       L2
Packet Return:    GRE
Number of WXs:    3
WX 0 IP:          10.87.119.200
Assigned Hash:    92492492492492492492492492492492492
                  49249249249249249249249249249249
WX 1 IP:          10.87.120.200
Assigned Hash:    24924924924924924924924924924924924
                  92492492492492492492492492492492
WX 2 IP:          10.87.121.200
Assigned Hash:    49249249249249249249249249249249249
                  24924924924924924924924924924924
```

```
Service Identifier: 55
Number of routers: 0
Number of WXs:     0
```

WCCP Statistics:

```
'Here I Am' Send Count:    15744
'Here I Am' Rcv Count:     0
'I See You' Rcv Count:     9198
'Redirect Assignment' Send Count: 2
'Redirect Assignment' Rcv Count: 0
'Removal Query' Rcv Count:  0
'I See You' Error Count:    0
'Removal Query' Error Count: 0
Auth Mismatch Count:       0
Auth Failure Count:        0
```

```

Send Failure Count:      0
Receive Failure Count:   0
Select Failure Count:    0
No Gateway for Router ID Count: 0
RARP for Gateway Failure Count: 0
Restart Count:          1

```

WCCP GRE Statistics:

```

In Packet Count:      0
Bad Packet Count:     0
Non-WCCP Packet Count: 0
Fragmented Packet Count: 0
No Memory Count:      0
IP Send Error Count:  0
Process Packet Count: 0
Return GRE as GRE Count: 0
Return L2 as GRE Count: 2362
Return GRE as L2 Count: 0

```

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 all passthrough traffic is included in the tunnel. However, if an appropriate 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 path-mtu-discovery

In environments where the Maximum Transmission Unit (MTU) can vary by network path, MTU discovery automatically adjusts the size of the compressed meta-packets sent from the local WX to each remote WX endpoint.

If MTU discovery is enabled, three UDP probes with the do-not-fragment flag are sent to each remote endpoint, starting with the maximum MTU size (1500). If there is no response within two seconds, the next highest MTU size is tried (1492). Two lower sizes are attempted if necessary (1276 and 576).

For each tunnel, probes with the current MTU size are sent periodically to test whether the maximum MTU has decreased (unless the current size is 576). Also, if the discovered MTU for a tunnel is less than 1500, “upward” probes are sent periodically with the next highest MTU size to test whether the MTU for the path has increased.

1. To view the current MTU discovery settings:

```
show -run path-mtu-discovery
```

2. Type the following command to configure MTU discovery:

```
config path-mtu-discovery
```

3. To enable or disable MTU discovery (enabled by default):

```
set mode {on | off}
```

4. To change the number of minutes between probes (default is 10):

```
set probe-interval {default | <1 - 1440>}
set upward-probe-interval {default | <1 - 1440>}
```

5. To exclude a remote WX endpoint from receiving the MTU probes:

```
endpoint exclude ip-address <IP address>
```

To remove one or all WX endpoints from the exclusion list:

```
endpoint remove ip-address {all | <IP address>}
```

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 compression 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 Demo 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 acceleration.

On the compression 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 245).

1. To view the current Demo Mode settings:

```
show -run profile-mode
```

2. Type the following command to configure Demo Mode:

```
config profile-mode
```

3. To enable or disable Demo Mode (disabled by default):

```
set mode <on | off>
```



CAUTION: Enabling Demo 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 traffic classes for inbound bandwidth management:

Class	Description
Compressed	Compressed traffic from other WX devices.
Intranet	Uncompressed 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 274).
TCP	TCP traffic that is not in the Compressed or Intranet class.
Default	All traffic that is not in the Compressed, 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 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 168.

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 179). 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 <8-1000000></pre> <p>To specify the local WAN as oversubscribed (disabled by default):</p> <pre>set oversubscribed-mode <on off></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 <name></pre> <p>To move an application to a new class (an application can belong to only one class):</p> <pre>class application move <class> <application></pre> <p>To change the name of a class:</p> <pre>class set name <oldname> <newname></pre> <p>To delete a class (any applications in the class are moved to the Default class):</p> <pre>class remove <name></pre>
Templates	<p>To add a new template (up to 20 characters):</p> <pre>template add <name></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 <template> <class> <percent></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 <template> <class> <percentage></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 <template> <class> <priority></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 <name></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 <template> <class> <packets></pre> <pre>template set age-out <template> <class> <2-5000 no-limit></pre> <p>NOTE: Please contact Technical Support for assistance before changing queue lengths or age-out times.</p>

Settings	Commands
Endpoints	<p>To enable or disable bandwidth management for all current and future WX and non-WX endpoints:</p> <pre>set enable-all-endpoints <on off></pre> <p>The remote WAN circuit speed defaults to 1 Mbps for the WX 15, and 1 Gbps for all other WX endpoints. When this setting is off, QoS is not enabled for future endpoints, and you can use the tunnel remove command to disable QoS for individual endpoints.</p> <p>To enable bandwidth management to a single 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 <IP address> <speed></pre> <p>CAUTION: Unless bandwidth detection 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 170). 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 (in Kbps):</p> <pre>tunnel set <address other-traffic> <8-1000000></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 <address other-traffic> <template> tunnel set non-prime-time <address other-traffic> <template></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 <IP address></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>
Bandwidth detection	<p>If the WAN bandwidth to a remote WX device is variable, such as for Frame Relay or shared satellite links, you can enable bandwidth detection 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 compressed meta packet.</p> <p>To enable or disable bandwidth detection (disabled by default):</p> <pre>set congestion-control-mode <on off></pre> <p>To enable bandwidth detection for all QoS-enabled endpoints or a list of included endpoints (default is all endpoints):</p> <pre>set congestion-control-endpoint-policy <all include></pre> <p>To enable or disable bandwidth detection for a specific endpoint:</p> <pre>tunnel set congestion-control-mode <IP address> <on off></pre> <p>To specify the minimum bandwidth (in Kbps) for an endpoint (optional):</p> <pre>tunnel set min-bandwidth <IP address> <bandwidth></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 <tcp-fast-backoff scps-slow-backoff default></pre>

Settings	Commands
Tunnel passthrough	<p>When bandwidth detection is enabled, it applies only to the tunneled traffic sent between WX devices. For remote endpoints that have bandwidth detection enabled, you can enable or disable the inclusion of passthrough traffic in the service tunnel (disabled by default):</p> <pre>set tunnel-sr-passthrough <on off></pre> <p>Bandwidth detection is not applied to passthrough traffic sent to non-WX endpoints.</p>
Non-WX 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 depends on the device type (refer to “WX Device Specifications” on page 421).</p> <p>You can also view the WAN Throughput and WAN Application Summary reports for each virtual endpoint (refer to “WAN Statistics” on page 246).</p> <p>To add virtual endpoints:</p> <pre>non-sr add name <name> bandwidth <8-1000000> [prime-time <template>] [non-prime-time <template>] [mode <on off>]</pre> <p>Where:</p> <ul style="list-style-type: none"> ■ name <name> . Virtual endpoint name (up to 20 characters). If the name includes spaces, enclose the name in quotation marks. ■ bandwidth <8-1000000> . WAN circuit speed associated with this endpoint (in Kbps). ■ prime-time <template> . 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 %). ■ non-prime-time <template> . Optional name of the template used for nonprime-time hours. Default is none. ■ mode <on off> . 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. <p>NOTE: 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 <endpoint> subnets <ip-address>[/mask],<ip-address>[/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 <name> [new-name <name>] [bandwidth <8-1000000>] [prime-time <template>] [non-prime-time <template>] [mode <on off>]</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 <ip-address>[/mask]...</pre> <p>To delete a virtual endpoint:</p> <pre>non-sr remove <name></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 compressed 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 <tos dscp off></pre> <p>For applications whose traffic is compressed, specify whether the ToS/DSCP value is restored to its original value after the traffic is decompressed by the remote WX device (enabled by default):</p> <pre>tos-dscp set restore <on off></pre> <p>To enable or disable ToS/DSCP changes for a traffic class (disabled by default):</p> <pre>tos-dscp class set mode <class> <on off></pre> <p>To set a 6-bit DSCP value (0 to 63) for a traffic class:</p> <pre>tos-dscp class set dscp <class> <0-63></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 <class> <0-255></pre> <p>To set a ToS IP precedence value (0 to 7) for a traffic class:</p> <pre>tos-dscp class set ip-precedence <class> <0-7></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 <on off></pre> <p>To exclude additional LAN/WAN address or subnet pairs from outbound bandwidth management:</p> <pre>add <LAN address>[/mask]--<WAN address>[/mask]</pre> <p>Use an asterisk (*) by itself to indicate any LAN/WAN address or subnet.</p> <p>NOTE: 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 {<LAN address>[/mask]--<WAN address>[/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 324.

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 compression and decompression engines.

1. To view the current compression settings:

```
show -run reduction [all | network-sequence-mirroring | pre-sync status]
```

The “all” option includes compression statistics since the last time the device was reset. For example:

```
===== Reduction Statistics =====
Packets: Total=79837075 - Accept=43450309
Overflow=0 FilterPassthru=75 Default Decompressor=0 No Decompressor=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 compression settings and statistics:

Keyword	Description
Total	Number of uncompressed packets into the device.
Accept	Number of packets into the compression engine.
Overflow	Packets not compressed because the compression queue is full (the device is too busy or the WAN link is too slow).
FilterPassthru	Packets not compressed due to application or address filter settings.
Default Decompressor	Packets compressed and sent to the default decompressor.
No Decompressor	Packets not compressed because of no remote WX.
The following statistics are shown only if they are non-zero.	
Reject Protocol	Packets for IP protocols that are not compressed.
Accept Protocol	Packets for additional IP protocols that are enabled for compression (does not include TCP and UDP packets, which are compressed by default).
Exclude Address	Packets not compressed due to source/destination filter settings.
TTL Expired	Packets not compressed because the Time to Live value was zero.
Accept Fragmented	Fragmented packets compressed (fragment compression is enabled by default).
Reject Fragmented	Fragmented packets not compressed (fragment compression is disabled).
Malformed	Malformed packets not compressed.
SR Traffic	Management packets sent to other WX devices (not compressed).
Local	Packets destined for the local subnet (not compressed).
Mid Watermark packets	Packets that received less compression processing because the compression queue exceeded the optimum level (the device is busy or the WAN link is slow).
Mid Watermark reached	Number of times that the compression queue exceeded the optimum level.
Hi Watermark reached	Number of times the compression queue became full. Packets received while the queue is full are counted as overflow (not compressed).

2. Type the following command to enter the configure compression mode:

```
config reduction
```

3. The following table describes the compression settings.

Settings	Commands
Compression and Decompression	<p>You can disable decompression to stop other devices in the community from sending compressed data to this device: To enable or disable decompression (enabled by default):</p> <pre>set assembler <on off></pre> <p>To enable or disable data compression (enabled by default):</p> <pre>set reducer <on off></pre>
Endpoints	<p>By default, each WX device attempts to form service tunnels with all other devices in the community. To form tunnels with only specific devices:</p> <pre>add assembler-list <IP addresses></pre> <p>Multiple IP addresses must be separated by spaces. To replace the current list of decompressors with a new list:</p> <pre>set assembler-list <IP addresses></pre> <p>To specify which devices to form tunnels with (the default is all):</p> <pre>set assembler-mode <all list hub-only></pre> <p>To remove devices from the current list of decompressors:</p> <pre>remove assembler-list <IP addresses></pre>
Network Sequence Caching	<p>On WXC devices, you can enable or disable the use of Network Sequence Caching (NSC) for data compression. NSC uses disk storage to identify longer patterns of repeated traffic, and retains those patterns for longer periods of time (even when a service 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 compression must be enabled, and you must enable TCP Acceleration for the appropriate devices and applications. TCP Acceleration also requires outbound QoS (refer to “configure acceleration” on page 326).</p> <p>To enable NSC on a WXC device (disabled by default):</p> <pre>network-sequence-mirroring set mode <on off></pre> <p>To specify how heavily the disk is used for compression:</p> <pre>network-sequence-mirroring set disk-access-policy <0-3></pre> <p>The number (0 to 3) indicates how heavily the disk (and CPU) is used during data compression. 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 <on off></pre> <p>NSC Endpoints</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 <all list></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 <IP addresses> network-sequence-mirroring endpoint remove <IP addresses></pre> <p>NSC Applications</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 <include exclude></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 <name> network-sequence-mirroring application remove <name></pre>

Settings	Commands
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 compression 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 <NSC-device-address> ftp://<host:port>[:<username>:<password>]/<path and file name></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 <NSC-device-address> ftp://<host>/<path and file name></pre>
Default decompressors	<p>To create a list of up to six default decompressors (for more information about default decompressors, refer to “Defining Default Decompressors” on page 157):</p> <pre>set def-assembler-list <IP address></pre> <p>Multiple IP addresses must be separated by a space.</p> <p>To add subnets to be excluded from the default decompressors:</p> <pre>add excl-subnet-list <IP address>/<subnet mask></pre> <p>To remove all or specific subnets from the exclude list:</p> <pre>remove excl-subnet-list <all IP address>/<subnet mask></pre>
Preferred decompressors	<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 decompressor. To designate one or more (up to 80) preferred decompressors:</p> <pre>set pref-assembler-list <space separated list of WX IP addresses></pre>
Load balancing	<p>The load balancing policy enables two or more WX devices to share the transmission of compressed data to a common destination with equal cost paths. To enable or disable load balancing:</p> <pre>set lb-policy <off per-packet per-destination per-flow></pre> <p>Where:</p> <ul style="list-style-type: none"> ■ off. All traffic is routed to one of the available tunnels. No load balancing (default). ■ per-packet. Traffic is distributed on a per-packet basis (round robin). ■ per-destination. Traffic is distributed based on destination IP address. ■ per-flow. Traffic is distributed based on source and destination IP addresses and ports.
Tunnel mode	<p>To change how traffic is sent in a service tunnel to a remote device:</p> <pre>set tunnelmode <udp multi-flow visibility ipcomp></pre> <p>Where:</p> <ul style="list-style-type: none"> ■ udp. Uses UDP (port 3577) to send meta packets as a single traffic flow (default). ■ multi-flow. 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 <integer between 256 and 1024></pre> ■ visibility. 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 service tunnel. Verify that your tools are configured to monitor UDP traffic. ■ ipcomp. Uses the IP payload compression protocol (protocol number 108) to send meta packets as a single traffic flow. Provides optimum compression in most environments.

Settings	Commands																		
Tunnel switching	<p>To enable tunnel switching on selected devices, such as to send compressed traffic between communities (disabled by default):</p> <pre>set tunnel-switching {on off}</pre> <p>NOTE: This feature must be implemented carefully to avoid unnecessary compression. For more information, refer to “Configuring Tunnel Switching” on page 163. Tunnel switching is limited to compression and QoS, and cannot be used on a WX 100 in stack mode.</p>																		
Heartbeat packets and retry settings	<p>By default, UDP keep-alive “heartbeat” packets are sent every 5 seconds to confirm the operability of the service tunnels between WX devices. To change the heartbeat frequency:</p> <pre>set heartbeat-frequency <1-300></pre> <p>NOTE: 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 compressing data for the device (passthrough mode). If 10 consecutive heartbeats get no response, each remote device disables the tunnel. To change the number of missed heartbeats that stops compression and disconnects tunnels:</p> <pre>set heartbeat-misses passthru <number default> disconnect <number default></pre> <p>Note that the number of missed heartbeats allowed is higher for remote endpoints for which TCP Acceleration or Forward Error Correction is enabled (refer to “configure acceleration” on page 326).</p> <p>After a tunnel is disabled, each remote WX attempts to reestablish the tunnel every minute for the first hour, every 15 minutes for the second hour, every hour for the next 22 hours, and once a day thereafter. To change the frequency of attempts during the first hour:</p> <pre>set retry-frequency <1-5></pre> <p>To enable or disable the backoff mechanism which determines whether remote devices attempt to reestablish disconnected tunnels:</p> <pre>set retry-backoff <on off></pre>																		
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 <0-255></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> <thead> <tr> <th>IP Precedence</th><th>CLI value</th></tr> </thead> <tbody> <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> </tbody> </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>NOTE: These values are overridden by the IP precedence or DSCP settings defined for Multi-Path (refer to “Configuring Policy-Based Multi-Path” on page 129).</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 service 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 <on off></pre>																		

Settings	Commands
Dynamic Resource Allocation (DRA)	<p>DRA enhances compression on low-speed WAN links (such as 128 Kbps). During good network conditions (such as low CPU load), the WX device attempts to further compress the data without compromising latency or packet loss. To enable or disable DRA (enabled by default):</p> <pre>set modes DRA <on off></pre> <p>NOTE: 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 365.</p>
Fast compression tunnels	<p>Fast compression tunnels increase WAN throughput to remote WX devices by decreasing the resources devoted to MSR compression. This feature may lower compression percentages, but it allows Application Flow Acceleration, TCP Acceleration, 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 a WX 15, WX 20, or WXC 250). To enable fast compression tunnels for one or more remote endpoints (multiple IP addresses must be separated by spaces):</p> <pre>set fast-reduction-tunnel <IP-addresses> none</pre> <p>To disable one or more fast-compression tunnels, you must disable all tunnels and restart the compression engine:</p> <pre>config reduction set fast-reduction-tunnel none config reduction set reducer off config reduction set reducer on commit</pre> <p>You can then enable fast-compression for specific endpoints, as needed</p> <p>CAUTION: If you disable fast-compression tunnels for endpoints that have bandwidth detection enabled, all tunnelled traffic to those endpoints will be blocked. Before disabling fast-compression tunnels, always disable bandwidth detection to all endpoints that have fast-compression tunnels.</p>
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 compression 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 <number></pre> <p>The valid symbol sizes are: 1-8, 10, 12, 16, 32, and 64 (use -1 to restore the default). On a WX 15, do not exceed a symbol size of 8.</p>

Settings	Commands
Meta-packets	<p>Multiple packets of compressed 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 <number between 576 and 1500></pre> <p>By default, the amount of time each meta-packet of compressed 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 <mode></pre> <p>Where <mode> 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 compressed packets into meta packets so that compressed 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 compressed traffic based on the application. To add a policy route:</p> <pre>policy-route add <application name> <gateway_IP_address></pre> <p>To remove a policy route:</p> <pre>policy-route remove <application name></pre>
MAC addresses	<p>By default, the source hardware (MAC) address of a decompressed packet is the MAC address of the WX device that decompressed the packet. To change the source MAC address of decompressed packets:</p> <pre>set assembly-source-mac-mode [default copy-source user-defined <mac>]</pre> <p>Where:</p> <ul style="list-style-type: none"> ■ default. Uses the MAC address of the WX decompressor. ■ copy-source. Uses the source MAC address received in the compressed packet. ■ user-defined <mac>. Specify a MAC address (the format is xx:xx:xx:xx:xx:xx).
Compression tradeoff for speed	<p>Under high traffic loads, compression is scaled back to increase throughput. To specify the relative tradeoff of compression for speed (“default” indicates the “standard” tradeoff):</p> <pre>set reduction-tradeoff-for-speed <minimum standard maximum default></pre> <p>Use “minimum” to ensure optimum compression under all traffic loads. On high-speed WAN links (over 20 Mbps), “maximum” is recommended for optimum throughput.</p>

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 reduction-subnet

Compression subnets are the subnets on the LAN side of a WX device that you can selectively advertise to the other WX devices in the community. The other devices can then compress and accelerate traffic sent to the advertised subnets. Initially, the only compression 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 384)
- 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 356 and “configure rip” on page 383
 - Periodically poll the routing table of a Cisco router (refer to “configure route-poll” on page 387)
 - Import a file of routes from an FTP server (refer to “import-route-table” on page 315)
- Enable the WAN compression 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 compression 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 compression subnet mode:

```
config reduction-subnet
```

3. To add entries to the Enabled list of compression subnets:

```
add enable <IP address/subnet mask>
```

To add entries to the Disabled list of compression 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 compression 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 compression 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 Compression Subnets page and cannot be advertised for compression.

To include routes discovered on the Remote interface as potential compression 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 357).



NOTE: Allow up to one minute for the remote routes to be added to the list of compression 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. To dynamically adjust advertised subnets to exclude (carve out) any hosts or gateways that become unreachable (disabled by default as of WXOS 5.5):


```
set carveout <on | off>
```
 6. To verify the mode and compression 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 compression can occur only between devices in the same community. You can define separate communities to control how data is compressed, and you can add the same device to multiple communities for backup and redundancy. A device that belongs to multiple communities can compress and decompress 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 (compressors) or assemblers (decompressors):

```
show -run reg-detail [<IP address> | -assemblers | -reducers]
```

```
Number of registered nodes: 4
Number of compressors: 4
Number of decompressors: 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 2006 WX1
192.168.53.22   SA/SR    0    7    0    JAN 07 09:56:43 2006 WX2
192.168.54.22   SA/SR    0    6    0    JAN 07 14:41:21 2006 WX3
192.168.55.22   SA/SR    R    0    7    0    JAN 07 09:51:42 2006 WX4
Key for 'Duty': H=Hub R=RegServer S=SecondaryRegServer
Key for 'Type': SA=Decompressor SR=Compressor
```

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 <registration server IP address></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, juniper.</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>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.</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 <secondary registration server IP address></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 <3-hours 8-hours 24-hours 7-days once-only></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 <1-5></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 <2-60></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 <days></pre>

Settings	Commands
Communities	<p>To add the name of a new community (up to 31 characters): <code>community add <name></code></p> <p>To change a community name: <code>community set name <old name> <new name></code></p> <p>To add a WX device to a community: <code>community remote-sr add <community> <IP address></code></p> <p>To delete a WX device from a community: <code>community remote-sr remove <community> <IP address></code></p> <p>To delete a community: <code>community remove <name></code></p> <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>
Devices	<p>To delete a WX device from all communities: <code>delete-entry <IP address></code></p>
Database	<p>If you add or change communities on a secondary registration server, you can export the community database to the primary server: <code>community export-to-primary</code></p> <p>To prevent the primary registration server from overriding the community database on the secondary server, enter the following command on the secondary server: <code>community set disable-import</code></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 remote-routes

Remote routes are the compression subnets advertised by other WX devices in the community. You can specify how often remote routes are fetched from other devices, whether multi-cost routes are supported, and whether each remote route is validated.

- 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.

- Type the following command to enter the configure remote route mode:

```
config remote-routes
```

- 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.

4. To enable or disable support for environments where multiple WX devices advertise the same subnet with different costs (disabled by default):

```
set multi-cost <on | off>
```

If this option is enabled, the lowest-cost remote WX is used until it becomes unavailable, and then the next lowest-cost WX is used. Note that WX load balancing and preferred decompressors will be ignored and should be disabled.

5. 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. Route validation is not supported for off-path devices using packet interception or when load balancing is enabled.

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 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 357).

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 (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>
```

To enable or disable the generation of a TTL expiration response for packets that have a time-to-live of zero (enabled by default). Enabling this option allows the WX to appear as a hop in a traceroute (the packets are dropped). If this option is disabled, packets with a TTL of zero are passed through without any processing.

```
set icmp-ttl-response <on | off>
```


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 185, 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 185: 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 compressed 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 compressed meta-packet in round robin fashion.
- **per-destination.** ToS values are assigned to each compressed meta-packet by applying a hash function to the destination IP address.
- **per-flow.** ToS values are assigned to each compressed 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 compression 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 324. To define the RADIUS servers, refer to “configure radius” on page 370.

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 enable or disable the lock on the front panel keypad on devices that have an LCD (disabled by default):

```
set front-panel <on | off>
```

4. 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 >
```

To accept all SSH connections or just SSH version 2 (the default is all):

```
set ssh-protocol <all | v2-only>
```

5. 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>]}
```

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 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 community string has spaces, enclose it in double quotation marks.

5. To enable or disable the generation of SNMP traps (disabled by default):

```
set trap <on | off>
```

To enable or disable traps for authentication failures (disabled by default):

```
set auth-failure-trap <on | off>
```

6. To add a trap destination and community string (limit the community string to 25 characters):

```
traps add destination <IP address> community <string>
```

If the community string has spaces, enclose it in double quotation marks. To change a community string:

```
traps set destination <IP address> community <string>
```

To delete a trap destination:

```
traps delete destination <IP address>
```



NOTE: For a description of each trap, refer to “SNMP Traps and Syslog Messages” on page 427.

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 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 338).



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 ssl certificate

For each application that you enable for SSL optimization, the SSL certificates and private keys for the application server must be imported on the WX device closest to the server. You can import up to 100 certificates. To enable applications for SSL optimization, refer to “configure ssl optimization” on page 392.



NOTE: Imported SSL certificates and private keys are NOT copied to a backup WX device.

To load SSL certificates and private keys:

1. To view the list of imported SSL certificates or the details of a specific certificate, use the following command:

```
show -run ssl certificate <friendly_name>
```

2. To import an SSL certificate from an FTP or TFTP server:

```
config ssl certificate add friendly-name <name> certificate-file <[path]name>
[privkey-file <[path]name>] server <server_address[<path>]> [passphrase]
```

Where:

- **friendly-name <name>** . Unique name for the certificate to be imported (up to 15 characters). Use only ASCII letters, numbers, and underscores. To update the private key of an existing certificate, this name must match the existing name.
- **certificate-file <[path]name>** . Name of the certificate file on the server. Include the absolute path (such as “/tmp/certs/mycert.crt”) if the path is not specified with the server address. The supported certificate formats are PKCS12, PEM, and DER.
- **privkey-file <[path]name>** . If the private key is not included in the certificate, specify the name of the private key file on the server. Include the absolute path (such as “/tmp/certs/mypriv.key”) if it is not specified with the server address.
- **server <address>** . Address of the FTP or TFTP server where the certificate file is stored. Include the path if it is not specified with the certificate or private key file name. The formats are:

```
ftp://<IP address>[:<user>:<pass>]/<path>]
tftp://<IP address>[/<path>]
```
- **passphrase** . If the private key is encrypted, include the **passphrase** keyword. You will be prompted to enter the password needed to access the key. Encrypted certificates are not supported. However, if you import a PKCS12 file, you must enter the password used to create the PKCS12 file.

3. To delete outdated certificates:

```
config ssl certificate delete friendly-name <name>
```

configure ssl optimization

Applications that use SSL version 3 (or later) for encryption can be decrypted and optimized (compressed, accelerated, and managed by QoS). In addition to enabling applications for optimization, the SSL certificates for each application server must be imported on the WX closest to the server (refer to “configure ssl certificate” on page 391). For more information about SSL optimization, refer to “Optimizing SSL Traffic” on page 240.

To enable SSL optimization or reset the current SSL statistics:

1. To view the current SSL configuration and status, use the following command:

```
show -run ssl optimization <configuration | status>
```

2. Type the following command to enter the configure SSL optimization mode:

```
config ssl optimization
```

3. To enable or disable SSL optimization (disabled by default):

```
set mode <on | off>
```

To indicate whether only optimized SSL traffic flows are sent over the IPsec tunnel (enabled by default):

```
set ipsec-optimized-only <on | off>
```

4. To optimize an SSL application (must have SSL encryption enabled in the application definition):

```
application add <name>
```

To disable optimization for an SSL application:

```
application remove <name>
```

5. To reset the current SSL optimization statistics:

```
reset-stats
```

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 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, WXC 500s, WXC 590s (refer to “Connecting Client Devices to the Server” on page 52). 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 service tunnels between two WX 100 stack 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 or WXC 590s, 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 service tunnels between two WX 100 stack servers with the same number of clients, specify the same number of tunnels (one for each client) on both WX 100 stack servers:

```
config stack-group endpoint set ip-address <IP address> max-tunnels <1-6>
```



NOTE: Be sure to disable bandwidth detection on the WX 100 (bandwidth detection reduces throughput if multiple tunnels are enabled). Also, configuring SSL optimization or Policy-Based Multi-Path overrides multiple service tunnels, and the command to establish multiple tunnels will be ignored.

To disable multiple service tunnels, set the maximum tunnels to one:

```
config stack-group endpoint set ip-address <IP address> max-tunnels 1
```

4. By default, tunnels are hosted only on the clients (not on the WX 100). To run the WX 100 in standalone mode (no clients) enter the following CLI command to allow tunnels to be hosted on the WX 100:

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

To return to server mode (all tunnels must be hosted on the clients):

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

5. 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.

6. 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.

1. To view the current syslog settings:

```
show -run syslog
```

2. To enable or disable syslog (disabled by default):

```
config syslog set syslog <on | off>
```

3. To enter the IP address of a syslog server:

```
config syslog set destination <IP address | none>
```

Multiple syslog addresses must be separated by spaces. Up to five syslog servers can be defined.

4. To specify the facility of the syslog messages (default is **local10**):

```
config syslog set facility local<1-7>
```

5. To set the severity of messages uploaded to the syslog server, specify any combination of “c”, “e”, “i” and “o” (do not include spaces between the letters):

```
config syslog set severity <ceio>
```

Where:

- **c.** Critical error messages about software or hardware malfunctions.
- **e.** Error message, such as License expired.
- **i.** Informational messages, such as reload requests and low-process stack messages.
- **o.** Informational messages about unusual events that are not errors.



NOTE: For a description of syslog messages, refer to “SNMP Traps and Syslog Messages” on page 427.

6. To commit the changes to the running configuration, type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

configure system

The topology settings determine whether the WX attempts to form a tunnel with all other WX devices in the same community, and affects the maximum number of tunnels the device can support. For guidelines about specifying the topology settings, refer to “Configuring Topology Settings” on page 108.

1. To view the current topology settings:

```
show -run system
```

2. To specify the topology type:

```
configure system topology type <hub | mesh | spoke | point-to-point>
```

To specify the community size for a hub or mesh topology:

```
configure system topology community-size <large | small>
```

3. To commit the changes to the running configuration, type **commit**. To retain the changes the next time the device is restarted, type **save-config**.

configure tacplus

You can define up to five TACACS+ servers to authenticate WX users. The servers are accessed in the order specified. WX devices conform to the TACACS+ protocol specification 1.78 (draft-grant-tacacs-02.txt). To specify how the servers are used to authenticate users, refer to “configure aaa” on page 324.

1. To view the current TACACS+ server settings:

```
show -run tacplus
```

2. Type the following command to enter the configure TACACS+ mode:

```
config tacplus
```

3. You can define up to 5 TACACS+ servers. To add a TACACS+ server:

```
server add name <name> ip-address <address> auth-port <number>  
timeout <seconds> retransmit <number>
```

Where:

- **name <name>** . TACACS+ server name (up to 31 characters). If the name includes spaces, enclose the name in quotation marks.
- **ip-address <address>** . IP address of the server.
- **auth-port <number>** . UDP port number (1 to 65535) on the server used for authentication (default is 49).
- **timeout <seconds>** . Number of seconds (1 to 60) that the WX waits for the server to respond (default is 10).
- **retransmit <number>** . Number of times (1 to 100) that a request is sent to the server before trying the next server, if any (default is three).

After you 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 TACACS+ 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 TACACS+ 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>
```

To remove a server definition:

```
server remove <name>
```

4. To change the source IP address used in TACACS+ packets (defaults to the device's IP address):

```
set client-source <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**.

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 <tftp://<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 249). 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 351). Note that data compression 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 352).

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) compressed (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 TCP Acceleration 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 FEC, FCS, and TCP Acceleration, 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	Timer
10.87.242.1	00:02:b3:8e:97:bb	dynamic	REMOTE	19:42 20m
1.1.1.1	00:00:00:00:00:00	static	LOCAL	perm
10.87.242.2	00:0e:0c:4a:dc:e6	dynamic	LOCAL	perm
10.87.242.3	00:0c:f1:fd:e3:85	dynamic	LOCAL	10:32 20m

The Timer column indicates the amount of time before a non-permanent entry expires, and the timer reset value when the address is revalidated. The default timer is 20 minutes (20m), but it can be set from 1 to 9 minutes for static entries (refer to “configure arp” on page 335).

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	Compression	Decompression	RTT(ms)	Tunnel Uptime	Description
10.87.52.22	M	Failed	Disallowed			Remote compressor off
10.87.53.22	M	Failed	Disallowed			Remote compressor off
10.87.54.22	M	Failed	Disallowed			Remote compressor off
10.87.56.22	M	Connected	Connected	400	P1 00:01:02	

Role: H-Hub, S-Spoke, M-Mesh, B-Backup, A-Active Backup

Tunnel Uptime: dd:hh:mm

The “Pn” value indicates the client port number (shown on WX 100 stack servers only).

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 event

The view the performance event definitions (if any) and/or the system events that are enabled for generation:

```
show -run event [configuration <all | performance | system>]
```

===== Performance Events =====

ID Metric	Threshold Type	Event Threshold	App/Class	Endpoints
x 46 Compression	absolute	< 96	all	10.88.9.100
x 47 TCP Acceleration	absolute	< 90	all	all
x 49 Bytes Dropped	absolute	> 1	all	all

===== System Events =====

ID	Name	Enabled	Severity
1207	if-speed-mode-mismatch	enabled	major
1208	if-speed-mode-ok	enabled	ok
1209	if-duplex-mismatch	enabled	major
1802	fail-safe-mode-active	enabled	critical
.			
.			
.			

show filter

To view the applications and addresses whose traffic is not compressed, the protocols enabled for compression, and whether fragmented packets are compressed:

```
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 compression 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 274).
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 flow-reset

A traffic flow cannot be accelerated unless the WX sees the start of the flow. If flow reset is enabled, eligible CIFS traffic flows are reset if a packet is received for the flow within the specified number of seconds (5 to 86400) from the time the tunnel for the flow was initially established. The default is 900 (15 minutes).

To view the configuration and/or status of the flow reset:

```
show flow-reset [configuration | status]
```

```
===== Flow reset configuration information =====
```

```
Mode:      on
```

```
Duration: 900
```

```
===== Flow reset status information =====
```

```
End time:      FRI AUG 25 18:14:09 2006
```

```
Last reset time: FRI AUG 25 18:13:41 2006
```

```
Flow reset count:      460
```

```
Flow reset failures:    0
```

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 [application-filter | sa [<ip-address>]]
```

Where:

- **application-filter.** Displays the applications that require or never use IPsec.
- **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, Encryption
```

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-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
```



```

Fragmented Packet Count: 0
No Memory Count:      0
IP Send Error Count:   0
Process Packet Count:  0
Return GRE as GRE Count: 0
Return L2 as GRE Count: 2362
Return GRE as L2 Count: 0

```

show path-mtu-discovery

To view the current MTU discovery settings:

```

show -run path-mtu-discovery
Mode:                on
Probe-interval:      10 (default)
Upward-probe-interval: 10 (default)

```

*** PathMTU Excluded Endpoints ***

===== PathMTU Probe State =====

DC-address	Tunnel MTU	Probe State	Probe MTU	Seconds To Next Probe	Up-Probe State	Up-Probe MTU	Seconds To Next Probe	Up-MTU Count
10.87.57.22	1500	ACK	1500	451	INIT	1500	na	0
10.87.52.22	1500	MIN-MTU	576	na	SENT-LAST	576	367	0
10.87.54.22	1500	ACK	1500	179	INIT	1500	na	0

Note that if the discovery Probe MTU is 576 (the minimum), the Seconds to Next Probe is “na” to indicate that no more discovery probes are sent. Similarly, if the Up-Probe MTU is 1500 (the maximum), no more “upward” probes are sent. Note that the “config reduction set max-meta-pkt-size” command can be used to set the maximum MTU size to less than 1500.

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 Demo Mode settings:

```

show -run profile-mode

Demo Mode: off
Remote WX 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 decompression: 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 compression settings:

```
show -run reduction [all | network-sequence-mirroring | pre-sync status]
```

The “all” option includes compression statistics since the last time the device was reset. For example:

```
===== Compression Statistics =====
Packets: Total=79837075 - Accept=43450309
Overflow=0 FilterPassthru=75 Default Decompressor=0 No Decompressor=45723

Reject Protocol=57
Accept Protocol=58040
WX Traffic=1107
Local=36902
Mid Watermark packets=1351
Mid Watermark reached=5
Hi Watermark reached=1
```

The following table describes the compression settings and statistics:

Keyword	Description
Total	Number of uncompressed packets into the device.
Accept	Number of packets into the compression engine.
Overflow	Packets not compressed because the compression queue is full (the device is too busy or the WAN link is too slow).
FilterPassthru	Packets not compressed due to application or address filter settings.
Default Decompressor	Packets compressed and sent to the default decompressor.
No Decompressor	Packets not compressed because of no remote WX device.
The following statistics are shown only if they are non-zero.	
Reject Protocol	Packets for IP protocols that are not compressed.
Accept Protocol	Packets for additional IP protocols that are enabled for compression (does not include TCP and UDP packets, which are compressed by default).
Exclude Address	Packets not compressed due to source/destination filter settings.
TTL Expired	Packets not compressed because the Time to Live value was zero.
Accept Fragmented	Fragmented packets compressed (fragment compression is enabled by default).
Reject Fragmented	Compression of fragmented packets is disabled).
Malformed	Malformed packets not compressed.
SR Traffic	Management packets sent to other WX devices (not compressed).
Local	Packets destined for the local subnet (not compressed).
Mid Watermark packets	Packets that received less compression processing because the compression queue exceeded the optimum level (the device is busy or the WAN link is slow).
Mid Watermark reached	Number of times the compression queue exceeded the optimum level.
Hi Watermark reached	Number of times the compression queue was full. Packets received while the queue is full are counted as overflow (not compressed).

show reduction-subnet

To view the current compression 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 (compressors) or assemblers (decompressors):

```
show -run reg-detail [<IP address> | -assemblers | -reducers]
```

```
Number of registered nodes: 4
```

```
Number of compressors: 4
```

```
Number of decompressors: 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 2006	WX1
192.168.53.22	SA/SR		0	7	0	JAN 07 09:56:43 2006	WX2
192.168.54.22	SA/SR		0	6	0	JAN 07 14:41:21 2006	WX3
192.168.55.22	SA/SR	R	0	7	0	JAN 07 09:51:42 2006	WX4

```
Key for 'Duty': H=Hub R=RegServer S=SecondaryRegServer
```

```
Key for 'Type': SA=Decompressor SR=Compressor
```

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 compression 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
Multi-cost mode: off
```

Destination	Netmask	Cost	Decompression-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

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>]

WXC-10.87.247.2# show route
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
ICMP TTL Response: on
ICMP Redirect Ignore: off

Number of routes: 3
  Destination    Netmask      Gateway      Type    Timer
  0.0.0.0        0.0.0.0      10.87.247.1  static   10
  10.87.247.0    255.255.255.0  10.87.247.2  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: on
Trap destination      Trap community string
10.87.52.146          xxxxx
10.87.240.3           xxxxx
10.87.240.240         xxxxx
Authentication-failure trap status: on
```

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 ssl certificate

To view the list of imported SSL certificates or the details of a specific certificate:

```
show -run ssl certificate [<friendly_name>]
```

show ssl optimization

To view the current SSL optimization configuration:

```
show -run ssl optimization configuration

Global mode: on
IPSec optimized only: on
Application Name  Mode
HTTPS            on
```


To view the current SSL optimization statistics:

```
show -run ssl optimization status

Total client-side flows      = 0
Active client-side flows    = 0
Client PDUs(client-side flows) = 0
Server PDUs(client-side flows) = 0
Total server-side flows     = 332
Active server-side flows    = 0
Client PDUs(server-side flows) = 1334
Server PDUs(server-side flows) = 1336
Maximum renegotiate count   = 0
Maximum certificate depth   = 1
Flows with soft error       = 0
Flows with hard error       = 0
Total alerts                = 664
Session cache hits          = 330
Session cache misses        = 0
Session cache checksum matches = 0
Session cache add failures  = 0
Maximum session item lifetime(sec) = 0
Maximum PDU length in forward direction = 139
Maximum PDU length in reverse direction = 1108
Maximum flow duration (sec) = 1
```

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, memory size, disk space, and topology type and size:

```
show -run system
```

```
System name: SR-10.87.72.10
```

```
Location:
```

```
Contact:
```

```
Software version: 5.5.0.8
```

```
Model No.: WXC-500 - v1.5
```

```
System started at: FRI JUL 14 10:52:06 2006
```

```
System up for: 2 days 22 hours 5 minutes
```

```
Info for flash file system: /ata0
```

```
File system size: 255836160 bytes
```

```
Free space: 214167552 bytes
```

```
File system block size: 4096
```

```
Hard drive 0: operational
```

```
Serial number   : L505D67H
```

```
Firmware revision: BANC1G10
```

```
Model number    : Maxtor 7L250S0
```

```
Hard drive 1: operational
```

```
Serial number   : L505D74H
```

```
Firmware revision: BANC1G10
```

```
Model number    : Maxtor 7L250S0
```

```
Total physical memory: 2139090944 (0x7F7FF000) bytes
```

```
Topology type: mesh
```

```
Community size: small
```

show system-name

To view the system name:

```
show -run system-name
```

```
System name: WX-10.87.72.10
```

show tacplus

To view the current TACACS+ server settings:

```
show -run tacplus
```

```
TACACS+ Server Settings:
```

Order	Name	IP Address	Auth Port	Timeout	Retransmits
1	TACACS1	10.10.20.30	49	10	3

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 WX devices, and the pin-outs for the DB9 console port:

- “General Specifications—All Platforms” in the next section
- “WX Family Specifications” on page 424
- “WXC Family Specifications” on page 425
- “DB9 Console Port Pin-Outs” on page 426

General Specifications—All Platforms

Table 12 describes the specifications that apply to all WX and WXC platforms.

Table 12: General Specifications for All WX and WXC Platforms

Product Features	Description
Traffic Services	IP payload compression, protocol acceleration, QoS, traffic visibility, application identification, route optimization, IPSec encryption, packet aggregation
Protocols Supported	Any IP-based traffic (TCP, UDP, GRE, ICMP, L2TP, etc.)
Applications Supported	All IP-based applications, including Microsoft Office applications, Oracle E-Business Suite, Sharepoint, Microsoft Exchange, Citrix, SAP, web-based applications, etc.
Network Integration	
Installation	In-line between aggregation switch and edge router, or off WAN router using route injection (RIP), WCCPv2, or policy-based routing
Auto-deployment	No-touch auto-configuration available out of the box through WX CMS software
Transparency	Transparent bridge mode operation, configurable DSCP and IP port transparency
Topology Support	Point to point, hub and spoke, full mesh
Network Discovery	Via RIP v1/v2, OSPF, and router polling
Tunnel Creation	Automatic or manual
Asymmetric Routing Support	Supported for both inline and off-path
Load Balancing	Active/active or active/passive, with passive in hot standby
Fault Tolerant Non-stop Operation	10/100/1000 BaseT auto switch-to-wire on any power, hardware, or software failure condition
High Availability	A backup device can support multiple primary devices; automatically fail-to-wire
Quality of Service (QoS)	
Honor, Preserve and/or Set ToS/DSCP	Retain settings or prioritize using ToS/DiffServ values by application
Bandwidth Allocation	Create traffic classes for bandwidth allocation with time of day option
Application Identification	Automatic, based on source/destination IP address/port, ToS/DSCP, IP protocol, L7 identification for HTTP and Citrix; follows port hopping applications (FTP, Exchange)
Rate Optimization	Multipath: application level path selection based on link SLA
Traffic Acceleration	
Packet Flow Acceleration	TCP Acceleration, Fast Connection Setup, and Forward Error Correction
Application Flow Acceleration	Microsoft CIFS, Linux file services, Microsoft Exchange, and HTTP
Device Management	
SNMP, Syslog	SNMPv2c, MIB II, WX Enterprise MIB and local syslog
Secure Remote Access	SSHv1, SSHv2, and HTTPS (SSL)
Reports	26 device-level reports available through WebView; 36 network-wide reports available with WX CMS
Authentication, Authorization, and Accounting	AAA local database and RADIUS and TACACS+ support
Network Upgradeable	Via FTP, HTTP and TFTP; dual software images and configurations

Product Features	Description
Monitoring	
Compression Statistics	Per device, per application, and per destination; both real-time and historical
WAN Performance Statistics	Network latency, loss, and availability for SLA monitoring and enforcement
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	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
Event/Performance Monitoring	Generate automatic alerts (SNMP traps, email, console) for up to 200 administrator definable performance or system events
Operating Environment	
Temperature	41° to 104° F (5° to 40° C)
Relative Humidity	10% to 85%, non-condensing at 95° F (35° C)
Maximum Altitude	10,000 feet (3048 meters)
Non-Operating Environment	
Temperature	-40° to 158° F (-40° to 70° C)
Relative Humidity	5% to 95%, non-condensing at 95° F (35° C)
Maximum Altitude	40,000 feet (12,192 meters)
Regulations	
Emissions	FCC Class A, EN 55022 Class A, EN 55024 Immunity, EN 61000-3-2, VCCI Class A
Safety	CAN/CSA-C22.2 No. 60950-1-03 - UL 60950-1 and EN 60950-1
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 Family Specifications

Table 13 describes the specifications for the nondisk-based WX family of platforms.

Table 13: WX Family Specifications

	WX 15	WX 20	WX 60	WX 100
Performance				
Total compression throughput speed	64 Kbps to 1 Mbps	64 Kbps to 2 Mbps	512 Kbps to 20 Mbps	1 Mbps to 20 Mbps (standalone) 1 Mbps to 155 Mbps (with up to six WX clients)
Tunnels supported	Up to 6 with all features enabled	Up to 10 with all features enabled	Up to 110 with all features enabled	Up to 105 with all features enabled (standalone) Up to 660 with all features enabled and six WX clients
Virtual endpoints	2	5	110	105
Application definitions	Up to 100	Up to 256	Up to 256	Up to 256
Routes	Up to 1000	Up to 8K	Up to 8K	Up to 24K (OSPF), 8K other
Client of WX 100?	No	No	Yes	No
Connections				
Network interfaces	Two copper fail-to-wire 10/100 Ethernet ports	Two copper fail-to-wire 10/100 Ethernet ports	Two copper fail-to-wire 10/100/ 1000 Ethernet ports	Two copper fail-to-wire 10/100/ 1000 Ethernet ports or Two fiber 1000 Ethernet ports
Power				
Power	100-240VAC, 50-60Hz, 50 Watts Max or 170 BTU/hr	100-240VAC, 50-60Hz, 150 Watts Max or 510 BTU/hr	100-240VAC, 50-60Hz, 150 Watts Max or 510 BTU/hr	Dual 100-240VAC, 50-60Hz, 300 Watts Max or 1025 BTU/hr. Designed to work with IT power systems.
Dimensions and Weight				
(W x H x D)	15.3 x 1.8 x 9.1 in (38.9 x 4.5 x 23 cm) 1 rack unit	17.1 x 1.8 x 14.3 in (43.4 x 4.5 x 36.3 cm) 1 rack unit	17.1 x 3.4 x 16.7 in (43.4 x 8.7 x 42.4 cm) 2 rack units	17.1 x 3.4 x 16.7 in (43.4 x 8.7 x 42.4 cm) 2 rack units
Weight	4 lbs (1.8 kg)	19 lbs (8.6 kg)	20.2 lbs (9.2 kg)	30 lbs (13.6 kg)

WXC Family Specifications

Table 14 describes the specifications for the disk-based WXC family of platforms.

Table 14: WXC Family Specifications

	WXC 250	WXC 500	WXC 590	WX 100 w/up to Six WXC Clients
Performance				
Total compression throughput speed	128 Kbps to 2 Mbps	512 Kbps to 20 Mbps	2 Mbps to 45 Mbps	34 Mbps to 155 Mbps
Tunnels supported	Up to 10 with all features enabled	Up to 50 with all features enabled	Up to 140 with all features enabled	Up to 840 with all features enabled
Virtual endpoints	5	60	120	105
Disk Capacity	40 GB	500 GB (redundant 250 GB drives)	500 GB (redundant, field-serviceable 250 GB drives)	Up to 3 TB
Application definitions	Up to 256	Up to 256	Up to 256	Up to 256
Routes	Up to 8K	Up to 8K routes	Up to 8K routes	Up to 24K (OSPF), 8K other
Client of WX 100?	No	Yes	Yes	No
Connections				
Network interfaces	Two copper fail-to-wire 10/100 Ethernet ports	Two copper fail-to-wire 10/100/1000 Ethernet ports	Two copper fail-to-wire 10/100/1000 Ethernet ports	Two copper fail-to-wire 10/100/1000 Ethernet ports or Two fiber 1000 Ethernet interfaces
Power				
Power	100-240VAC, 50-60Hz, 150 Watts Max or 510 BTU/hr	100-240VAC, 50-60Hz, 150 Watts Max or 510 BTU/hr	Dual 100-240VAC, 50-60Hz, 300 Watts Max or 1025 BTU/hr. Designed to work with IT power systems.	Dual 100-240VAC, 50-60Hz, 300 Watts Max or 1025 BTU/hr. Designed to work with IT power systems.
Dimensions and Weight				
(W x H x D)	17.1 x 1.8 x 14.3 in (43.4 x 4.5 x 36.3 cm) 1 rack unit	17.1 x 3.4 x 16.7 in (43.4 x 8.7 x 42.4 cm) 2 rack units	17.1 x 3.4 x 16.7 in (43.4 x 8.7 x 42.4 cm) 2 rack units	17.1 x 3.4 x 16.7 in (43.4 x 8.7 x 42.4 cm) 2 rack units
Weight	21 lbs (9.5 kg)	25 lbs (11.3 kg)	25 lbs (11.3 kg)	30 lbs (13.6 kg)

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 15: 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 16: 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 and syslog messages for the system events generated by WX devices. System events are also displayed on the Events monitoring report along with performance events (refer to “Events Summary” on page 280).

Severity Levels

Syslog messages and the Events report severity filter use different terms for some severity levels, as shown below. In the following sections, the syslog term is shown first, followed by the filter term in parentheses. WX performance events use the same severity terms as the report filter.

Syslog Severity	Events Report Severity Filter
Notice	OK
Information	Warning
Error	Major
Critical	Critical

To view the Events report, refer to “Events Summary” on page 280.

System Events and SNMP Traps

Table 17 describes each system event, and provides the SNMP trap name and its associated object ID (OID).

Table 17: System Events and SNMP Traps

Event Name	Severity	Description	SNMP Trap/OID
Cold Start	Notice (OK)	The device was restarted.	Cold Start 1.3.6.1.6.3.1.1.5.1
Compressor Buffer Overflow	Error (Major)	The compression input buffer is approaching full capacity.	jnxWxEvtCompressionBufferOverflow 1.3.6.1.4.1.8239.2.2.1.3.2.0.2
Compressor Passthrough Traffic	Information (Warning)	Compression stopped for a remote endpoint that is unavailable (no heartbeats received).	None
Compressor Session Closed	Information (Warning)	Compressor session to the device noted in jnxWxCommonEventDescr was terminated.	jnxWxEvtCompressionSessionClosed 1.3.6.1.4.1.8239.2.2.1.3.2.0.3
Compressor Session Opened	Notice (OK)	Compressor session to the device noted in jnxWxCommonEventDescr was opened.	jnxWxEvtCompressionSessionOpened 1.3.6.1.4.1.8239.2.2.1.3.2.0.5
DC Queue Above High Watermark	Information (Warning)	The incoming decompression queue is becoming full, which signals remote WxS to temporarily reduce the compression load. Occasional High Watermark messages indicate normal traffic bursts. Continuous High Watermark messages may indicate that the decompressing WxS should be upgraded to a higher capacity device or that the load into the WxS should be decreased.	jnxWxEvtDCQAboveHiWatermark 1.3.6.1.4.1.8239.2.2.1.3.2.0.19
DC Queue Below High Watermark	Notice (OK)	Decompression queue dropped below the high-watermark level.	jnxWxEvtDCQBelowHiWatermark 1.3.6.1.4.1.8239.2.2.1.3.2.0.20
Decompressor Session Closed	Information (Warning)	Decompressor session to the device noted in jnxWxCommonEventDescr was terminated.	jnxWxEvtDecompressionSessionClosed 1.3.6.1.4.1.8239.2.2.1.3.2.0.4
Decompressor Session Opened	Notice (OK)	Decompressor session to the device noted in jnxWxCommonEventDescr was opened.	jnxWxEvtDecompressionSessionOpened 1.3.6.1.4.1.8239.2.2.1.3.2.0.6
Disk Failure	Error (Major)	A hard disk on a WXC device is down or failed to initialize. Operation continues with some loss in the percentage of compression.	jnxWxEvtDiskFailure 1.3.6.1.4.1.8239.2.2.1.3.2.0.17
Fail Safe Mode Active	Critical	Indicates that the device was restarted in Safe Mode. Safe Mode operation keeps the device powered on, but all traffic is passed through without compression.	jnxWxCommonEventInFailSafeMode 1.3.6.1.4.1.8239.2.1.3.2.0.1

Event Name	Severity	Description	SNMP Trap/OID
Interface Duplex Mismatch	Error (Major)	A possible duplex mismatch exists between the local or remote interface and the device attached to that interface. The interface is identified by jnxWxCommonEventDescr.	jnxWxCommonEventInterfaceDuplexMismatch 1.3.6.1.4.1.8239.2.1.3.2.0.14
Interface Speed Mode Mismatch	Error (Major)	A speed or duplex mismatch exists between the local and remote interface on the WX device.	jnxWxCommonEventInterfaceSpeedMismatch 1.3.6.1.4.1.8239.2.1.3.2.0.12
Interface Speed Mode Ok	Notice (OK)	A previously detected mismatch between the local and remote interface is now resolved. The local and remote interface speed and mode are matched.	jnxWxCommonEventInterfaceSpeedOk 1.3.6.1.4.1.8239.2.1.3.2.0.13
Invalid Route Removed	Information (Warning)	Route that could not be validated was deleted.	None
Ipssec Security Association Added	Notice (OK)	An IPSec security association (SA) was negotiated and added to the SA database.	jnxWxCommonEventIpssecSecurityAssociationAdded 1.3.6.1.4.1.8239.2.1.3.2.0.15
Ipssec Security Association Deleted	Information (Warning)	An IPSec security association has been deleted from the SA database.	jnxWxCommonEventIpssecSecurityAssociationDeleted 1.3.6.1.4.1.8239.2.1.3.2.0.17
Ipssec Security Association Expired	Information (Warning)	An IPSec security association has expired.	jnxWxCommonEventIpssecSecurityAssociationExpired 1.3.6.1.4.1.8239.2.1.3.2.0.16
Ipssec Throughput Limit Exceeded	Error (Major)	Exceeded licensed throughput for IPSec-encrypted traffic. Please contact Juniper Networks to obtain a new license with a higher speed.	jnxWxCommonEventIpssecThruputLimitExceeded 1.3.6.1.4.1.8239.2.1.3.2.0.22
LAN Link Down	Information (Warning)	Indicates the Local interface link has failed. Verify that the link state change was not due to a network error.	LAN Link Down 1.3.6.1.6.3.1.1.5.4
LAN Link Up	Notice (OK)	Indicates the Local interface link has been established.	LAN Link Up 1.3.6.1.6.3.1.1.5.3
License Expired	Error (Major)	Software license has expired and all processing is disabled. Please contact Juniper Networks for a permanent license.	jnxWxCommonEventLicenseExpired 1.3.6.1.4.1.8239.2.1.3.2.0.4
License Will Expire	Information (Warning)	Software license will expire soon. If you are using an evaluation license, contact Juniper Networks to obtain a permanent license.	jnxWxCommonEventLicenseWillExpire 1.3.6.1.4.1.8239.2.1.3.2.0.6
Login Failure	Error (Major)	An attempt to log in has failed. Verify the user is authorized to administer the device.	jnxWxCommonEventLoginFailure 1.3.6.1.4.1.8239.2.1.3.2.0.7
Management Config Save Failure	Error (Major)	An attempt to save the configuration failed.	None
Management Startup Config Saved	Information (Warning)	Startup configuration was saved successfully.	None

Event Name	Severity	Description	SNMP Trap/OID
Multi Path Status Change	Information (Warning)	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.	jnxWxEvtMultiPathStatusChange 1.3.6.1.4.1.8239.2.2.1.3.2.0.16
Performance Threshold Crossed	Information (Warning)	A performance threshold was violated. Review the Event report for more details (refer to “Events Summary” on page 280). If this system event is disabled, no SNMP traps or syslog messages are generated for performance events.	jnxWxEvtPerformanceThreshCrossed 1.3.6.1.4.1.8239.2.2.1.3.2.0.21
Power Supply Failure	Error (Major)	One or more sources of power to the system has failed. A redundant power-supply has presumably taken over.	jnxWxCommonEvtPowerSupplyFailure 1.3.6.1.4.1.8239.2.1.3.2.0.2
Power Supply Ok	Notice (OK)	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.	jnxWxCommonEvtPowerSupplyOk 1.3.6.1.4.1.8239.2.1.3.2.0.3
Primary Down Backup Engage Failed	Error (Major)	The primary WX is unreachable, and the backup WX has failed to engage (generated by the backup device).	jnxWxEvtPrimaryDownBackupEngageFailed 1.3.6.1.4.1.8239.2.2.1.3.2.0.14
Primary Down Backup Engaged	Notice (OK)	The primary WX is unreachable, and the backup WX has successfully engaged (generated by the backup device).	jnxWxEvtPrimaryDownBackupEngaged 1.3.6.1.4.1.8239.2.2.1.3.2.0.13
Primary Reg Server Unreachable	Error (Major)	The primary registration server is currently unreachable.	jnxWxEvtPrimaryRegServerUnreachable 1.3.6.1.4.1.8239.2.2.1.3.2.0.7
Primary Self Registration Done	Notice (OK)	A WX registered successfully with the primary registration server.	None
Primary Up Backup Disengaged	Notice (OK)	The primary WX is now reachable, and the backup WX has disengaged (generated by the backup device).	jnxWxEvtPrimaryUpBackupDisengaged 1.3.6.1.4.1.8239.2.2.1.3.2.0.15
RIP Auth Failure	Error (Major)	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.	jnxWxEvtRipAuthFailure 1.3.6.1.4.1.8239.2.2.1.3.2.0.1
Registration Password Mismatch	Error (Major)	A WX failed to register due to an incorrect registration server password.	None
Secondary Reg Server Unreachable	Error (Major)	The secondary registration server is currently unreachable.	jnxWxEvtSecondaryRegServerUnreachable 1.3.6.1.4.1.8239.2.2.1.3.2.0.8
Secondary Registration IP Failed	Error (Major)	The Multi-Path secondary IP address failed to register with the registration server.	None
Secondary Registration IP OK	Notice (OK)	The Multi-Path secondary IP address registered with the registration server.	None
Secondary Self Registration Done	Notice (OK)	A WX registered successfully with the secondary registration server.	None

Event Name	Severity	Description	SNMP Trap/OID
Security Login Success	Notice (OK)	User logged in successfully.	None
Stack Client Disk Failure	Error (Major)	A hard disk on a WXC client of a WX 100 stack server is down or failed to initialize. Operation continues with some loss in compression.	jnxWxEventDiskFailure 1.3.6.1.4.1.8239.2.2.1.3.2.0.17
Stack Client Disk OK	Notice (OK)	Disk on a WXC client device is working properly.	None
Stack Client Link Down	Error (Major)	On a WX 100 stack server, the link to a client device is down.	None
Stack Client Link Up	Notice (OK)	On a WX 100 stack server, the link to a client device is working properly.	None
Throughput Limit Exceeded	Error (Major)	Exceeded licensed throughput. Please contact Juniper Networks to obtain a new license with a higher speed.	jnxWxCommonEventThruputLimitExceeded 1.3.6.1.4.1.8239.2.1.3.2.0.5
WAN Link Down	Information (Warning)	Indicates the Remote interface link has failed. Verify that the link state change was not due to a network error.	WAN Link Down 1.3.6.1.6.3.1.1.5.4
WAN Link Up	Notice (OK)	Indicates the Remote interface link has been established.	WAN Link Up 1.3.6.1.6.3.1.1.5.3
WAN Perf Status Change	Information (Warning)	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.	jnxWxEventWanPerfStatusChange 1.3.6.1.4.1.8239.2.2.1.3.2.0.18
WAN Performance Acceptable	Notice (OK)	WAN performance no longer violates the current loss or latency thresholds.	None
WAN Performance Unacceptable	Information (Warning)	WAN performance has violated the current loss or latency threshold.	None

Syslog Messages

Table 18 lists the syslog messages generated by WX devices.

Table 18: Syslog Messages

Message ID	101: PN_LIC_LICENSE_WILL_EXPIRE_SOON_ID
Message	License will expire on < date >
Severity	Information (Warning)
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
Severity	Error (Major)
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 compression/decompression has been disabled
Severity	Error (Major)
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
Severity	Error (Major)
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 > .
Severity	Information (Warning).
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 > .
Severity	Information (Warning)
Recommended Action	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: Compressor buffer is reaching full capacity
Severity	Error (Major)
Recommended Action	If the situation persists, reduce the traffic entering the compressor. Appropriate traffic filter may also be used to reduce the amount of packets to be processed by the compressor.

Message ID	1002: PN_ASSEMBLER_END_SESSION_INFO_ID
Message	SA: Session closed - ip = < ip address > sesid = < id > .
Severity	Information (Warning)
Recommended Action	Decompressor session to device < ip address > has ended. If this is not a 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 > .
Severity	Error (Major)
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 > .
Severity	Error (Major)
Recommended Action	Check the network connectivity to secondary registration server < ip address > .
Message ID	1104: PN_REGISTER_PRIMARY_SELFREG_SUCCESS_ID
Message	REG: Primary self registration done. IP = < ip address > .
Severity	Notice (OK)
Recommended Action	None
Message ID	1105: PN_REGISTER_SEC_SELFREG_SUCCESS_ID
Message	REG: Secondary self registration done. IP = < ip address > .
Severity	Notice (OK)
Recommended Action	None
Message ID	1106: PN_REGISTER_PASSWORD_MISMATCH_ERROR_ID
Message	REG: Registration failed. Password mismatch. IP = < ip address >
Severity	Error (Major)
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	1107: PN_REGISTER_SEC_IP_ERROR_ID
Message	REG: Failed in setting the secondary IP for the primary reg server = < primary ip address > , secIP = < secondary ip address > .
Severity	Error (Major)
Recommended Action	Check the network connectivity to secondary registration server < secondary ip address > .
Message ID	1108: PN_REGISTER_SEC_IP_SUCCESS_ID
Message	REG: Registration of secondary IP address completed with primary reg server IP = < primary ip address >
Severity	Notice (OK)
Recommended Action	None

Message ID	1202: PN_BRIDGE_GENERIC_HARDENING_ERROR_ID
Message	Health monitor detected anomalous system condition
Severity	Error (Major)
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 >
Severity	Information (Warning)
Recommended Action	None
Message ID	1204: PN_BRIDGE_LOCAL_LINK_DOWN_INFO_ID
Message	Local interface: Link Down
Severity	Information (Warning)
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 >
Severity	Notice (OK)
Recommended Action	None
Message ID	1206: PN_BRIDGE_REMOTE_LINK_DOWN_INFO_ID
Message	Remote interface: Link Down
Severity	Information (Warning)
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.
Severity	Notice (OK)
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.
Severity	Error (Major)
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.
Severity	Notice (OK)
Recommended Action	None. The connectivity to primary device is restored.
Message ID	1702: PN_MGMT_STARTUP_CONFIG_SAVED_ID
Message	SaveStartupConfig: Saved successfully
Severity	Notice (OK)
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 >
Severity	Error (Major)
Recommended Action	Contact Technical Support.
Message ID	1802: PN_INIT_IN_SAFE_MODE_ID
Message	Safe-mode suspend: case 2
Severity	Critical
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
Severity	Notice (OK)
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 >
Severity	Error (Major)
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 >
Severity	Notice (OK)
Recommended Action	Please verify that the person who logged in was someone authorized to configure the system.
Message ID	2302: PN_PFA_DCQ_ABOVE_HWM_ID
Message	Decompression queue above high watermark.
Severity	Information (Warning)
Recommended Action	None
Message ID	2303: PN_PFA_DCQ_BELOW_HWM_ID
Message	Decompression queue below high watermark.
Severity	Notice (OK)
Recommended Action	None
Message ID	2501: PN_IPSEC_GENERIC_ERROR_ID
Message	One of the following: <ul style="list-style-type: none"> ■ IPsec Added SA < source IP address > -> < destination IP address > s : SPI < SPI number > < encryption algorithm > < authentication algorithm > Hours Remaining 24.00 MBytes Remaining 100.00 ■ IPsec Expired SA < source IP address > -> < destination IP address > s : SPI < SPI number > due to < "life time" or "data life time" > ■ IPsec Deleted SA < source IP address > -> < destination IP address > s : SPI < SPI number >

Severity	Added SA – Notice (OK) Expired/Deleted SA – Information (Warning)
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
Severity	Error (Major)
Recommended Action	NSC continues without this disk, but with some loss in the percentage of data compression. 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"> ■ WP: ***** Unacceptable Performance detected due to LOSS on Path, Path Ip = < remote IP address > ***** ■ WP: ***** Unacceptable Performance detected due to LATENCY on Path, Path Ip = < remote IP address > ***** ■ WP: ***** Acceptable Performance detected on Path, Path Ip = < remote IP address > *****
Severity	Acceptable – Notice (OK) Unacceptable – Information (Warning)
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 437
- Performance Statistics Export on page 438
- Top Traffic Export on page 446
- Flow Diagnostics Export on page 447

NetFlow Version 5 Export

Traffic data can be exported to a Cisco NetFlow server in Version 5 format (refer to “Traffic Statistics” on page 274).

Table 19 describes the NetFlow packet header.

Table 19: 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 20 describes each traffic flow entry in a NetFlow packet (up to 30 entries per packet).

Table 20: 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 298).

- “General Device Information” in the next section
- “Data Section Information” on page 439
- “System Session Statistics” on page 440
- “Compression Session Statistics” on page 442
- “Application Session Statistics” on page 442
- “Bandwidth Management Statistics” on page 444
- “Inbound Traffic By Port Statistics” on page 445

General Device Information

Table 21 describes the exported general device information.

Table 21: 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.
TCP Acceleration applications	Names of the applications using TCP Acceleration.
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 Demo Mode.

Data Section Information

Table 22 describes the data section information that precedes the set of statistic tables for each exported time range.

Table 22: Data Section Information

Parameter	Description
< time > data section	Indicates the time range for the statistics tables that follow: <ul style="list-style-type: none"> ■ This hour ■ Last hour ■ Today ■ Yesterday ■ This week ■ Last week
ip =	IP address of the 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 23 describes the exported system session statistics.

Table 23: 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 Decompression Engine.
Bytes Out AE	Number of bytes out of the Decompression Engine.
Packets Into AE	Number of packets into the Decompression Engine.
Packets Out AE	Number of packets out of the Decompression Engine.
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 compression due to no remote WX.
Packets PT NO AE	Number of packets that passed through without compression due to no remote WX.
Bytes PT By Filter	Number of bytes that passed through without compression due to a manually configured filter (such as an application filter).
Packets PT By Filter	Number of packets that passed through without compression due to a manually configured filter (such as an application filter).
OfPt Bytes (Overflow Pass-through)	Number of bytes that passed through without compression due to device buffer overflow.
OfPt Packets (Overflow Pass-through)	Number of packets that passed through without compression due to device buffer overflow.
Bytes PT NO SR	Number of bytes that passed through without compression due to a disabled compression engine on this device.
Packets PT NO SR	Number of packets that passed through without compression due to a disabled compression engine on this device.
Bytes PT NON-IP	Number of non-IP bytes that passed through without compression (e.g., IPX, etc.).
Packets PT NON-IP	Number of non-IP packets that passed through without compression (e.g., IPX, etc.).
Bytes PT IP-Other	Number of IP bytes that passed through without compression because the protocols were not configured for compression.
Packets PT IP-Other	Number of IP packets that passed through without compression because the protocols were not configured for compression.
Bytes PT SR	Number of bytes that passed through without compression 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 compression 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 compression because the device is part of a compression cluster and the data will be processed by another WX device.
Packets PT SR-Hash	Number of packets that passed through without compression because the device is part of a compression cluster and the data will be processed by another WX device.
Bytes PT IpFrag	Number of bytes that passed through without compression because the device is not enabled to compress IP fragments.
Packets PT IpFrag	Number of packets that passed through without compression because the device is not enabled to compress IP fragments.

Parameter	Description
Bytes PT License	Number of bytes that passed through without compression because the throughput level exceeded the device's license.
Packets PT License	Number of packets that passed through without compression because the throughput level exceeded the device's license.
Bytes PT Tunneled Only	Number of bytes that passed through without compression.
Packets PT Tunneled Only	Number of packets that passed through without compression.
Bytes PT VLAN	Number of bytes of VLAN traffic that passed through without compression.
Packets PT VLAN	Number of packets of VLAN traffic that passed through without compression.
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 Compression 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 Compression 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 Compression Engine at the peak five-second interval of data input.
TP Packets Out (throughput)	Number of packets out of the Compression 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	<p>Maximum data compression rate for any five second interval within the selected time period. Peak percentage compression is calculated by the following formula:</p> $10^5 \times \left(\frac{\text{Bytes In} - \text{Bytes Out}}{\text{Bytes In}} \right) = \text{Peak \% Reduction}$
Rsv H1 through Rsv H20	Reserved
PkIn1 to PkIn6	<p>Six fields that show the number of packets in each of six packet-size ranges for traffic into the WX device, as follows:</p> <ul style="list-style-type: none"> ■ PkIn1 Less than 64 bytes ■ PkIn2 64 to 127 ■ PkIn3 128 to 255 ■ PkIn4 256 to 511 ■ PkIn5 512 to 1023 ■ PkIn6 More than 1023 bytes
PkOut1 to PkOut6	<p>Six fields that show the number of packets in each of six packet-size ranges for traffic out of the WX device, as follows:</p> <ul style="list-style-type: none"> ■ PkOut1 Less than 64 bytes ■ PkOut2 64 to 127 ■ PkOut3 128 to 255 ■ PkOut4 256 to 511 ■ PkOut5 512 to 1023 ■ PkOut6 More than 1023 bytes

Compression Session Statistics

Table 24 describes the exported compression statistics for each session.

Table 24: Compression 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 compressed and/or encrypted data from this device.
Packets In	Number of packets into this compression engine that were intended for the destination IP address.
Packets Out	Number of compressed packets sent to the destination IP address.
Packets Into Ipsec	Number of packets that were identified for encryption and intended for the destination IP address.
Packets Out of Ipsec	Number of encrypted packets sent to the destination IP address.
Packets Dropped by Ipsec	Number of packets intended for the destination IP address that were dropped according to the default IPsec policy.
Ipsec Overhead	Number of bytes added by IPsec processing.

Application Session Statistics

Table 25 describes the exported application session statistics.

Table 25: 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 compressed data from this device.
Bytes In	Number of bytes into the device that have been identified for compression, 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 compression, and addressed for the WX device listed with the destination IP address and application ID.
Acc Bytes In	Number of bytes eligible for TCP Acceleration.
Est Boost Bytes	Estimated number of bytes accelerated by TCP Acceleration.
Active Session time	Number of milliseconds during which data was sent for all TCP Acceleration 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 TCP Acceleration 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 compression using NSC (WXC devices only), and sent to the WX device listed with the destination IP address and application ID.

WAN Statistics

Table 26 describes the exported WAN statistics.

Table 26: 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 27 describes the exported Application Flow Acceleration statistics.

Table 27: 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 28 describes the bandwidth management statistics collected per application class for each service tunnel.

Table 28: Bandwidth Management Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Tunnel	Outbound bandwidth management: The IP address of the destination decompressor or the default allocation. Inbound bandwidth management: The parameter is Inbound.
Class	Outbound bandwidth management: The bandwidth class ID, which is a collection of applications that a user has mapped to the class. Inbound bandwidth management: One of the four pre-defined classes (i.e., Compressed, Intranet, TCP or Default).
Bytes In	Outbound bandwidth management: The total number of application bytes into the WX device. Inbound bandwidth management: The total number of bytes into the Remote interface of the WX device by class.
Bytes Out	Outbound bandwidth management: The total number of application bytes out of outbound bandwidth management. Inbound bandwidth management: the total number of bytes out of inbound bandwidth management.
Bytes Dropped	Outbound bandwidth management: The total number of application bytes dropped by the bandwidth management feature. Inbound bandwidth management: The total number of bytes dropped by the bandwidth management feature.
Packets In	Outbound bandwidth management: The total number of application packets into the WX device. inbound bandwidth management: The total number of packets passed into the device by inbound bandwidth management.
Packets Out	Outbound bandwidth management: The total number of application packets transmitted by the device. (The total number does not include meta packetization.) Inbound bandwidth management: The total number of packets out of inbound bandwidth management.
Packets Dropped	Outbound bandwidth management: The total number of application packets dropped by the bandwidth management feature. Inbound bandwidth management: The total number of packets dropped by the bandwidth management feature.

WAN Performance Statistics

Table 29 describes the exported WAN performance statistics.

Table 29: 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 30 describes the Inbound traffic by port statistics.

Table 30: Inbound Traffic by Port Data

Parameter	Description
Src Port	Inbound data's source port number.
Bytes In	Number of compressed bytes from the source port for unmonitored applications.
Packets In	Number of compressed packets from the source port for unmonitored applications.
Dst Port	Inbound data's destination port number.
Bytes In	Number of compressed bytes to the destination port for unmonitored applications.
Packets In	Number of compressed packets to the destination port for unmonitored applications.

Top Traffic Export

Table 31 describes the Traffic statistics exported to the *ip-flow.csv* file.

Table 31: 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.
BytSent	Number of bytes sent by the source.
PktsSent	Number of packets sent by the source.
BytRcv	Number of bytes received by the source.
PktsRcv	Number of packets received by the source.
TotalSendDelay	Cumulative delay between packets sent (in milliseconds).
TotalRcvDelay	Cumulative delay between packets received (in milliseconds).
Type	Indicates the traffic type: <ul style="list-style-type: none"> ■ RA. Compressed application (matched an application definition) ■ RO. Compressed undefined application ■ PT. Passed through due to policy setting ■ U. Unknown passthrough traffic, such as non-TCP/UDP traffic
StartTime	Start date and time of traffic flow.
EndTime	End date and time of traffic flow.

Flow Diagnostics Export

Table 32 describes the flow diagnostics data exported to the *rtdiag-flow.csv* file.

Table 32: Flow Diagnostics

Parameter	Description
SrcIp	IP address of the flow source.
SrcPort	Source port number.
DstIp	IP address of the flow destination.
DstPort	Destination port number.
Application	Traffic flow application name.
Type	Application type (Default, CIFS, Citrix, Exchange, FTP, or HTTP).
Proto	Traffic flow protocol (TCP, UDP, or protocol number).
Start Time	Date and time the flow started.
General Flow	
Flow Initiated Locally	The WX was the first to see the SYN packet (TRUE or FALSE). Note that application acceleration diagnostics (including CIFS, Exchange, and HTTP) are available only for flows initiated locally.
Source/Destination Filter Hit	The flow matches a source/destination filter (TRUE or FALSE).
Idle	No packets have been sent for the last 10 seconds (TRUE or FALSE).
Application Name	Flow application.
Bytes From LAN	Number of uncompressed bytes received from the LAN.
Bytes From WAN	Number of decompressed bytes received from the WAN.
Bytes To Wan	Number of compressed bytes sent to the WAN.
Stack Client Id	Number of the WX 100 client (1 to 6) hosting the tunnel (zero for non-WX 100 clients).
Throughput From WAN	Decompressed throughput received from the WAN for the last five seconds (shown as “n Kbytes/sec”).
Compression Status	Indicates “Compressed - Defined App”, “Compressed - Undefined App”, or “Passthrough - < reason >”.
TCP Acceleration	
SYN Error (Decompressor)	The decompressor received a SYN for an invalid session (TRUE or FALSE).
SYN-ACK Error (Decompressor)	The decompressor detected a SYN-ACK error, caused by bad sequence numbers, invalid TCP options, no SYN packet seen, or no compression session (TRUE or FALSE).
SYN-ACK Error(Compressor)	The compressor detected a SYN-ACK error, caused by bad sequence numbers, unexpected sequence numbers, , or no decompression session (TRUE or FALSE).
Qos Not Active	QoS was not enabled when the flow started (TRUE or FALSE).
Bad TCP Options	Some TCP options were unrecognized (TRUE or FALSE).
Slowdown (RDTP)	Number of times full RDTP compression buffers slowed the transmission rate.
Slowdown (Decompressor TCP Slow)	Number of times the transmission rate slowed due to a busy remote decompressor. When this occurs, packets are sent, but they are not ACKed locally.
Slowdown (Flow)	Number of times TCP slowed the transmission rate.
Slowdown (Age Update)	Number of times the flow has been quiet for more than 60 seconds. New packets are not ACKed locally until packets are received from the remote WX.
Duplicate SYN seen (Compressor)	The compressor detected two SYN packets (TRUE or FALSE).

Parameter	Description
Duplicate SYN-ACK seen (Compressor)	The compressor detected two SYN-ACK packets (TRUE or FALSE).
Duplicate SYN seen (Decompressor)	The decompressor detected two SYN packets (TRUE or FALSE).
Duplicate SYN-ACK seen (Decompressor)	The compressor detected two SYN-ACK packets (TRUE or FALSE).
Session Switched	The AFP session was switched to another member of the cluster (TRUE or FALSE).
Dropped Count (MaxFreedBuf)	Number of packets dropped due to a slow remote decompressor (Decompressor TCP Slow).
Pkts Retransmitted	Number of packets retransmitted by RDTP.
Retransmission Errors	Number of times retransmission failed.
Pkts Accepted	Number of compressed packets accepted by RDTP.
Syn From Cluster Peer	The SYN packet was received by another WX in the cluster (TRUE or FALSE).
Termination Active	TCP termination is active (TRUE or FALSE).
LAN Retransmits	Number of packets retransmitted to the LAN.
LAN Out Of Order	Number of out-of-order packets received and sent to the LAN.
LAN Duplicate Acks	Number of duplicate acknowledgements received from the LAN-side endpoint.
WAN Out of Order	Number of out-of-order packets received from the WAN.
Application Acceleration	
Protocol Accelerated	Name of the accelerated application (CIFS, Exchange, or HTTP)
Acceleration Currently Active	The flow is being accelerated now (TRUE or FALSE).
AAP Sync Version	Indicates the version of Application Flow Acceleration.
Flow Initialized	The flow has been initialized (TRUE or FALSE).
Flow State	Indicates the flow state (ACTIVE or PASSIVE).
Hard Quit	The application ended the flow (TRUE or FALSE).
Soft Quit	The application disabled acceleration for the flow (TRUE or FALSE).
Unknown Quit	An unknown agent disabled acceleration (TRUE or FALSE).
Sequence Overlap	Packets received with overlapping sequence numbers (TRUE or FALSE).
SYN Timer Scheduled	SYN packet received, waiting for SYN-ACK (TRUE or FALSE).
SYN Timer Unscheduled	SYN-ACK packet received (TRUE or FALSE).
SYN Timer Unscheduled (Clean)	Flow removed before SYN-ACK packet received (TRUE or FALSE).
SYN Timer Action	Timer expired before SYN-ACK packet received, and a hard quit occurs (TRUE or FALSE).
Packets To Server	Number of packets sent to the application server.
Packets To Client	Number of packets sent to the client.
CIFS	
Signed	SMB signing is enabled, and the flow cannot be accelerated (TRUE or FALSE).
Client OS	Client operating system, such as "Windows 2003".
Server OS	Server operating system.
Accl Reads	Number of accelerated read requests.
Total Reads	Total number of read requests.

Parameter	Description
Accl Writes	Number of accelerated write requests.
Total Writes	Total number of write requests.
Accl Trans2 Count	Number of accelerated Trans2 packets.
Total Trans2 Count	Total number of Trans2 packets.
Free Disk Space Positive	Number of times the WX determined that the server had enough disk space to satisfy a write request or a Trans2/SetEndOfFile request. These requests are accelerated.
Free Disk Space Negative	Number of times the server did not have enough disk space to satisfy a write request or a Trans2/SetEndOfFile request. These requests are not accelerated.
Free Disk Space Stale	Number of times write or Trans2/SetEndOfFile requests were not accelerated because the server's disk space information on the WX was out of date.
Free Disk Space Unavailable	Number of times write or Trans2/SetEndOfFile requests were not accelerated because the server's disk space information was unavailable.
Prefetch Reuse OK	When a file is closed, the read prefetch data for the file is kept for potential reuse. This counter is the number of times the prefetch data was found to be eligible for reuse.
Prefetch Reuse Not OK	Number of times the prefetch data was not eligible for reuse.
Whole File Prefetches	Number of times an entire file was prefetched for read acceleration.
Whole File Prefetch Alloc Failed	A file prefetch failed (TRUE or FALSE).
MID Table Full	The table of Multiplex IDs used to track CIFS requests became full at least once (TRUE or FALSE).
Write Update Error (Prefetch Buffer)	Indicates whether an update error occurred when an accelerated write overlaps with a read prefetch, and the write data is used to update the read prefetch data (TRUE or FALSE).
Close Accl Failure	Indicates whether an accelerated Close failed on the server (TRUE or FALSE).
Write Through on File Open	Indicates whether a file open operation writes the file directly to disk, rather than to a cache in memory (TRUE or FALSE).
Write Through On Write	Indicates whether file open operations were written directly to disk, rather than cached in memory (TRUE or FALSE). These write operations are not accelerated.
Send Buffer Alloc Failed	Indicates whether the allocation of a send buffer failed (TRUE or FALSE).
HTTP	
Accl Transactions	Number of accelerated transactions.
Partially Accl Transactions	Number of requests for partially prefetched objects.
Non Accl Transactions	Number of transactions not accelerated.
Not Accl: Header Not Parsed	Number of transactions not accelerated due to header parsing problem.
Not Accl: Cache Miss	Number of transactions not accelerated because the requested content was not in the cache.
Not Accl: In Cache But Cookie Mismatch	Number of cached transactions not accelerated due to a cookie mismatch.
Not Accl: Cant Cache MIME Type	Number of transactions not accelerated because the MIME type cannot be cached, such as when the response code is not 200 (OK) or 304.
Not Accl: In Cache But Other Problem	Number of cached transactions not accelerated due to other problems.
Not Accl: Un-Cacheable Request	Number of transactions not accelerated because the requests required that they not be cached..
Not Accl: Request Mismatch	Number of transactions not accelerated because the requests and cached content for the same object were different.
Not Accl: Collision	Number of transactions not accelerated because the cache location was already occupied.

Parameter	Description
Not Accl: Connection Max	Number of transactions not accelerated the number of requests exceeded the maximum number of connections.
Exchange	
Read PDUs	Number of read Protocol Data Units (PDUs) .
Write PDUs	Number of write PDUs.
Other PDUs	Number of other PDUs (not read or write).
Accl Reads	Number of accelerated read operations.
Accl Writes	Number of accelerated write operations.
Compression	
Using Default Decompressor	Indicates whether the default decompressor was used (TRUE or FALSE).
Decompressor IP	IP address of the remote WX that is decompressing the flow. "Not Applicable" indicates no remote WX available, or no data was sent, and "0.0.0.0" indicates that no WX lookup occurred.
Compression Percent	Percentage compression in traffic volume.
Network Sequence Caching	
Configured for Destination	Indicates whether NSC is enabled for the remote WX device (TRUE or FALSE).
Configured for App	Indicates whether NSC is enabled for the flow application (TRUE or FALSE).
Saw Reliable Packet	Indicates whether TCP acceleration is enabled, which is required for NSC (TRUE or FALSE).
Fast Compression	Indicates whether fast compression is enabled, which disables NSC (TRUE or FALSE).
NSC negotiated	Indicates whether NSC was negotiated with the remote device (TRUE or FALSE).
QoS	
Exclusion Filter Hit	Indicates whether the traffic is excluded from QoS (TRUE or FALSE).
Bytes Dropped	Number of bytes dropped by QoS.
Queue Depth	Number of flow packets waiting in the outbound queue.
Traffic Class	Name of the traffic class that the flow application belongs to.

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 33: 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

Demo Mode

The following topics describe how to configure and use a WX device in Demo Mode:

- “About Demo Mode” in the next section
- “Pre-Installation Tasks” on page 455
- “Installing a WX 15, WX 20, or WXC 250 in Demo Mode” on page 456
- “Installing Other Platforms in Demo Mode” on page 459
- “Configuring Demo Mode through the Web Console” on page 461
- “Viewing Performance Reports” on page 464
- “Exporting Performance Data” on page 466
- “Converting from Demo Mode to Active Mode” on page 466

About Demo Mode

Demo Mode is a passive mode of operation that lets you quickly baseline a WX device’s effectiveness on WAN traffic compression in your network. Demo Mode also provides an estimate of the average acceleration gains that are possible for your TCP applications. In Demo Mode, the WX device processes the observed data on a mirrored port, and has no effect on the actual network traffic. This lets you see the value of the WX device before you commit the device to your network. The WX device can be configured and installed in about five minutes.

Purpose and Benefits

Demo Mode enables network managers to identify the compression rates of their IP traffic by using a single WX device connected to a mirrored port on your network. The statistics from Demo Mode operation are presented in an easy to understand Web-based graphical format.

The benefits of Demo Mode include:


- A quick and simple method to evaluate the WX device (five-minute setup).
- A risk-free experience to the ease of administration and manageability of the WX device.

- A measurement of the effectiveness of data compression in your own network environment, including an ROI analysis.
- A confirmation the device's ability to learn and operate transparently in your network.

Sample Topology

The WX device is connected to a switch with a 10/100 Ethernet interface (WX 15, WX 20, and WXC 250) or 10/100/1000 Ethernet Interface (WX 60, WX 100, WXC 500, and WXC 590). The switch must be able to mirror traffic destined to the edge router, a common feature found in most switches. The WX device can be connected to any subnet that can see all traffic destined to/from the WAN at a specific edge/core location.

Figure 186: Sample Topology of a WX Device in Demo Mode



In Demo Mode, the WX device observes all traffic that passes through the device and generates real-time reports on the potential compression gains.

Security

Security is a top concern for all networking equipment within your network. WX devices in Demo Mode operate as follows:

- No packets are collected, only the potential data compression is measured.
- Secure CLI access using SSH, therefore no clear text passwords
- Secure Web access using SSL
- Secure ACL (Access Control Lists)
- MD5 Authentication

Return on Investment

Using your Demo Mode statistics along with your WAN cost structure, your WX sales team will work with you to generate an accurate, real, and defensible return on investment (ROI).

Pre-Installation Tasks

Before you install the WX device in Demo Mode, complete the following pre-installation tasks.

1. Identify interesting WAN links, which may include one or more of the following:
 - Heavily loaded links
 - Very expensive links
 - Links to locations targeted for growth
 - Links targeted for cost cutbacks, consolidation, or compression
2. Identify a suitable aggregation device (typically a switch) to connect the WX device in Demo Mode.
3. Reserve an IP address, and identify the subnet mask and default gateway for the WX device. The default gateway is the next hop on the WAN side of the WX device.
4. Set up the Ethernet mirror port (with Cisco switches use SPAN or PORT MONITOR) and check the port statistics to verify that traffic is being mirrored to this port.



NOTE: The WX device's Ethernet interfaces are auto-sensing.

After you have identified and set up a mirrored port, continue to one of the following sections depending on the type of WX device you have:

- “Installing a WX 15, WX 20, or WXC 250 in Demo Mode” in the next section
- “Installing Other Platforms in Demo Mode” on page 459.

Installing a WX 15, WX 20, or WXC 250 in Demo Mode



NOTE: The WX 15 is limited to WAN link speeds of 1 Mbps, while the WX 20 and WXC 250 support WAN speeds of 2 Mbps. If your WAN link speed exceeds 2 Mbps, you must use a WX 60, WX 100, WXC 500, or WXC 590.

Hardware Installation

After you have completed the pre-installation tasks, you are ready to install the WX device to a mirrored port in your network.

To install the WX device to a mirrored port in your network:

1. Set up the chassis.
 - To install the WX device in a 19-inch device rack, install the supplied brackets (front panel forward) to the sides of the device with the countersunk screws provided in the kit. Next, install the chassis in your network device rack.
 - To install the WX 15 on a desktop, 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. To install the WX 20 or WXC 250 on a desktop, you must first install the supplied rubber feet in the marked areas on the bottom of the chassis.
2. Connect the network cables to the WX device.




NOTE: Do not connect power to the device until Step 3.

The WX 15, WX 20, and WXC 250 have two 10/100 BaseT auto-sensing Ethernet interfaces. These RJ-45 ports are labeled REMOTE and LOCAL on the back of the chassis.

Using an Ethernet cable, connect a mirrored port on the aggregation device (such as a switch) to the LOCAL port of the WX device.

Figure 187: Connecting the WX 20 to a Mirrored Port



3. Connect the supplied power cord to the back of the chassis, and then connect the power cord to the local power source. Next, turn on the power switch
4. Proceed to the next section to configure the network settings.

Configuring Network Settings

After you have installed and powered on the WX device, the next step is to configure network settings for the device.

To configure the network settings for the WX 15, WX 20, or WXC 250, connect an ANSI-compatible terminal to the device's serial port and use a terminal emulation program, such as TeraTerm or HyperTerminal.




NOTE: The serial port is of type RS-232 (AT-compatible) with a male, DB-9 connector. You should use a female/female DB-9 crossover cable (null-modem cable) when connecting directly to a PC serial port.

To set IP parameters for the device using a terminal emulation program:

1. Connect an ANSI-compatible terminal to the serial port on the back of the WX device (Figure 188).

Figure 188: Connecting the WX 20 to an ANSI-compatible Terminal



2. Verify the serial port settings are as follows:
 - Baud rate: 9600 bps
 - Data bits: 8
 - Parity: none
 - Stop bits: 1
 - Flow control: none
 - Smooth-scroll: disabled
3. Start the terminal emulation program (such as HyperTerminal), and choose to connect via the serial port.

4. Type **admin** for the user name and **juniper** for the password (you may have to press Enter to see the first prompt).

You will now configure the IP parameters (IP address, subnet mask, and default gateway) to enable connectivity for this device. Once these parameters are configured, you can run the Quick Setup process through the Web console.

After running the Quick Setup process, additional management tasks can be performed via the Command Line Interface (CLI) or Web console.

5. To set the IP address, IP subnet mask, and default gateway:
 - a. Type an IP address, and then press Enter.
 - b. Type the subnet mask for the network, and then press Enter.
 - c. Type the default gateway for this device, and then press Enter.



NOTE: The default gateway is typically the next hop on the WAN side of the WX device.

6. The required parameters are now configured. You should now commit and save the configuration. To commit and save the configuration file with the default name and location, type the following commands:

```
commit
save-config
```

The configuration file is saved as “startup.cfg” and will be used if you reboot the device.

To save the configuration file with another name, type:

```
save-config <file name>
```

The name can be up to 8 characters. Do not include a file name extension (such as “.txt”).

7. On the back of the WX device, verify that the LINK LED for the LOCAL port is on. If not, toggle the MDI/MDI-X button (WX 20 and WXC 250 only).
8. On the front of the WX 15, verify that the “Operational” LED is on. On the WX 20 and WXC 250, verify that the “Bypass” LED is off.

You are now ready to log in to the Web console and run the Quick Setup program. Refer to “Configuring Demo Mode through the Web Console” on page 461.

Installing Other Platforms in Demo Mode

This section describes the hardware installation and the configuration steps for setting up a WX 60, WX 100, WXC 500, or WXC 590 in Demo Mode.

Hardware Installation

After you complete the pre-installation tasks, connect the WX device to a mirrored port in your network:


1. Set up the chassis.
 - If you plan to install the WX device in a 19-inch device rack, install the supplied brackets (front panel forward) to the sides of the device with the countersunk screws provided in the kit. Next, install the chassis in your network device rack.
 - If you plan to install the WX device on a desktop, place the chassis upside down on a smooth, flat surface. Next, install the supplied rubber feet in the marked areas on the bottom of the chassis. Finally, 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.
2. Connect an Ethernet cable from a mirrored port on the aggregation device (such as a switch) to the LOCAL port of the WX device.



NOTE: Do not connect power to the device until Step 3.

The WX 60, WX 100, and WXC 500 are configured with two 10/100/1000 BaseT auto-sensing Ethernet interfaces. These RJ-45 ports are labeled REMOTE and LOCAL on the back of the chassis. On the WXC 590, the ports are on the front panel.

Figure 189: Connecting the WX Device to a Mirrored Port

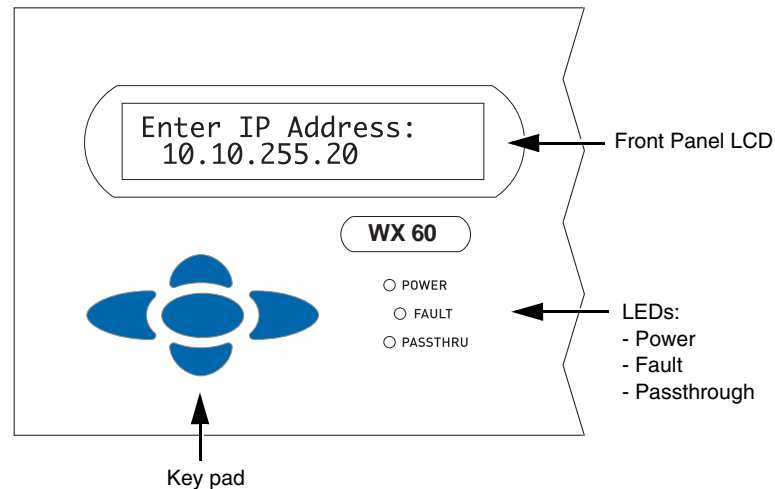


3. Connect the power cable to the back of the chassis, and then connect the other end of the power cable to your local power source.

Configuring Network Settings

After you install and start the WX device, the next step is to use the front-panel keypad and LCD to enter the network address information for the device. The LCD shown below is used for the WX 60 and WXC 500 (the WX 100 LCD is shown in “Configuring Network Settings” on page 49).

Figure 190: WX 60 Front Panel Keypad and LCD



When you start the device, “Juniper Networks” appears in the front panel LCD.

1. Press the Enter button (center button) to start.
 - a. At the “Select Setup Network_” prompt in the LCD, press Enter. You are prompted to enter network address information for the device.
 - b. Use the front-panel keypad to assign an IP address, subnet mask, and the default gateway for the WX device as follows:
 - 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.
 - Use the center button (**Enter**) to make a selection.



NOTE: The default gateway is typically the next hop on the Remote side of the WX device.

2. After entering network address and interface information, use the left arrow button to select Save and Commit, and press Enter to save the device configuration.
3. On the back of the device, verify that the LINK LED for the Local port is on.
4. On the front of the device, verify the “Passthru” LED is off.

Configuring Demo Mode through the Web Console

After assigning IP parameters to the WX device, you are ready to configure the device for Demo Mode operation, as described in the following sections:

- “Running Quick Setup” in the next section
- “Defining Virtual Devices in Demo Mode” on page 461
- “Excluding Traffic to the Local Subnet” on page 463

The WXOS Web console supports Microsoft Internet Explorer version 6.0 and later. Data is securely transmitted through HTTPS.

Running Quick Setup

After starting on the WX device and assigning IP parameters, you are ready to run Quick Setup and configure the device for Demo Mode operation.

To run Quick Setup:

1. From a local workstation, start your web browser and enter the following URL:

https://<IP address of the WX device>



NOTE: If the switch does not allow IP access to the mirror/SPAN port, connect the WX to another switch port, run Quick Setup, and then reconnect to the mirrored port. After collecting sufficient traffic, connect the WX to an accessible port to view the results. Alternatively, you can connect a terminal to the WX console port.

2. Depending on your browser settings, the Security Alert dialog box may appear, click **Yes** to proceed.
3. In the Login page, type **admin** for the user name and **juniper** for the password, and click **Login**.
4. Select Demo Mode, click **Next**, and then click **Finish**.

The WX is now configured for Demo Mode. The next time you log in to the Web console, the top banner will indicate Demo Mode. The front panel LCD (if any) also indicates Demo Mode.

Defining Virtual Devices in Demo Mode

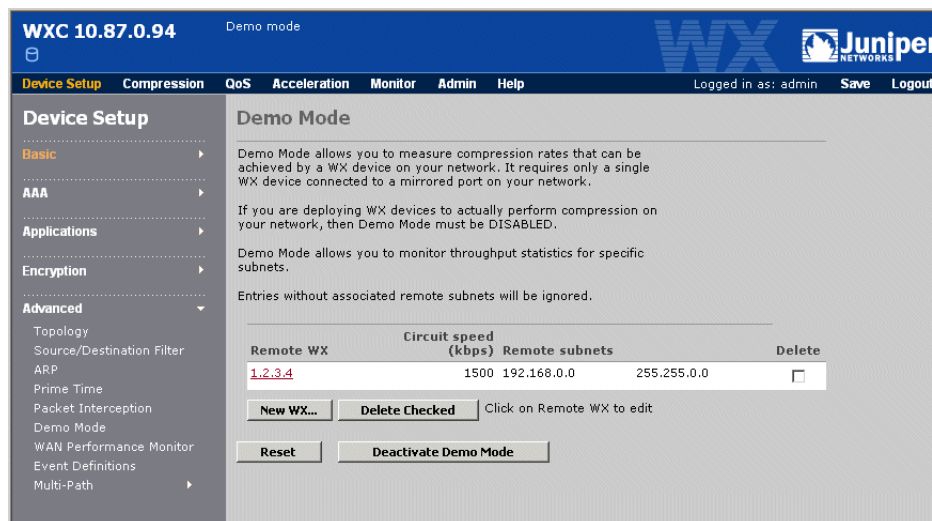
Demo Mode lets you see how a WX device performs in your network without affecting network traffic. In Demo Mode, the device passively calculates potential data compression statistics for all traffic and for individual applications.

To view the performance for specific remote subnets, you can define “virtual” WX devices and associate one or more subnets with each virtual device. On the monitoring reports, you can select a virtual device from the Destination menu to view the performance for the associated subnets (refer to “Monitoring and Reporting” on page 245).

To define remote subnets in Demo Mode:

1. Install the WX device as described in “Demo Mode” on page 453.
2. In the Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Demo Mode** (available only if the device is installed in Demo Mode).

Figure 191: Adding Virtual Devices in Demo Mode



On the Demo Mode page, you can:

- Add a virtual WX device, as described in Step 3.
- Change a virtual device. Click the virtual device address, change the remote subnets, and click **Submit**.
- Delete a virtual device. Click the check box next to the virtual device, and click **Delete Checked**.
- Switch from Demo Mode to Active Mode. Click **Deactivate Demo Mode** to reboot the device (the device restarts as a registration server). Verify that the Local and Remote interfaces are installed properly for live operation, as described in “Installation” on page 29.



NOTE: Network data cannot pass through the device while Demo Mode is enabled.

3. To add a virtual WX device and its remote subnets:
 - a. Click **New WX**.

Figure 192: Defining Remote Subnets in Demo Mode

The screenshot shows the Juniper WX Web Console interface. The top header displays 'WXC 10.87.0.94' and 'Demo mode'. The navigation menu on the left includes 'Device Setup', 'Compression', 'QoS', 'Acceleration', 'Monitor', 'Admin', and 'Help'. The 'Device Setup' menu is expanded, showing options like 'Basic', 'AAA', 'Applications', 'IPSec', and 'Advanced'. The 'Advanced' menu is further expanded, showing 'Topology', 'Source/Destination Filter', 'ARP', 'Prime Time', 'Packet Interception', 'Demo Mode', 'WAN Performance Monitor', 'Event Definitions', and 'Multi-Path'. The main configuration area is titled 'Demo Mode > New WX'. It contains a text box for 'Remote WX IP Address', a text box for 'Circuit speed (kbps)' with a value of '0', and a large text area for 'Remote subnets'. Below these fields are 'Submit', 'Reset', and 'Cancel' buttons.

b. Specify the following information:

- | | |
|----------------------|---|
| Remote SR IP Address | Enter any IP address for the virtual device (it need not be a real address). You can select this address from the Destination menu on compression reports to view the performance for the associated remote subnets.

The maximum number of virtual endpoints depends on the device type (refer to “WX Device Specifications” on page 421). |
| Remote subnets | Enter the remote subnets (one per line) associated with this virtual device. The subnet format is:

<IP address>/<subnet mask> |

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.

Excluding Traffic to the Local Subnet

To improve the accuracy of the compression statistics, exclude all traffic sent to the local subnet where the WX device is installed. This traffic would normally be decompressed by the device and will lower the average compression percentages if it is not excluded.

1. In the Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Source/Destination Filter**.
2. Select **DO NOT compress data between the following source/destination pairs**.
3. Enter an asterisk (*) in the Source field and the local subnet and mask in the **Destination** field. Do NOT select the **Bidirectional** check box.

You can define additional source/destination filters as needed (refer to “Using Source/Destination Filters” on page 112).


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

To further customize performance in Demo Mode, you can enable or disable data compression for specific applications (refer to “Compressing Traffic by Application” on page 151).

Viewing Performance Reports


After installing the WX device in Demo Mode, you can use the Monitor pages of the WXOS Web console to view potential compression and performance statistics:

- Compression statistics show the potential data compression for all traffic that traverses the WX device.



- Application Summary statistics include a pie chart of the nine monitored applications that have the highest percentage of the total traffic into the compression engine. The accompanying table shows the traffic statistics and the estimated percentage of data compression for each monitored application.

The **Remaining Applications** category in the chart shows the traffic for all other applications (both defined and undefined). The **Others** category in the table is for compressed applications that are undefined or unmonitored.



- Fast Connection Setup statistics show an estimate of the average reduction in session time that Fast Connection Setup may achieve for each application's "short" sessions. 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. No other acceleration statistics are available in Demo Mode.

Fast Connection Setup: Last 60 Minutes

Start Time: 02:41 PM 04/25/2005 Application: all Destination: all

Application	Total TCP Sessions	Short Sessions*		Average Short Session Time (msec)		Average Short Session Acceleration (percent)	
	(count)	(count)	(percent)	with Accel.	w/o Accel.		
CIFS	35	9	25.7%	24.00	39.78	39.7%	
LDAP	16	7	43.8%	39.43	58.14	32.2%	
Exchange	6	2	33.3%	10.00	20.00	50.0%	
HTTP	6	2	33.3%	20.00	34.28	41.7%	
Others	16	2	12.5%	42.00	63.00	33.3%	
AOL	0	0	0.0%	0.00	0.00	0.0%	
CVS	0	0	0.0%	0.00	0.00	0.0%	
Clearcase	0	0	0.0%	0.00	0.00	0.0%	
DNS	4	0	0.0%	0.00	0.00	0.0%	
FTP	0	0	0.0%	0.00	0.00	0.0%	
Filenet	0	0	0.0%	0.00	0.00	0.0%	
Hostname Resolution	0	0	0.0%	0.00	0.00	0.0%	

- Outbound QoS statistics can be viewed by configuring the WX device as its own "endpoint," which simulates an environment where all outbound traffic is sent to the same remote device. Note that the "Other traffic" endpoint is not used, and "virtual" devices cannot be used as endpoints for outbound QoS. For more information about outbound QoS, refer to "Configuring Outbound QoS Policies" on page 177.
- Inbound QoS is not applicable in Demo Mode.
- WAN Throughput and WAN Application Summary reports show only traffic sent to the WAN. Traffic received from the WAN is not monitored in Demo Mode.

For more information on viewing performance results, refer to “Monitoring and Reporting” on page 245.

Exporting Performance Data

While in Demo Mode, you can export performance data to a file in comma-separated variable (CSV) format. The exported data is similar to the data displayed in the Monitor page of the WXOS Web console. The CSV file can then be sent to your sales representative, or imported into a spreadsheet application (such as Microsoft Excel) or other data evaluation program.

To export data to CSV format:

1. In the WXOS Web console, click **Αδμιν** in the menu frame, click **Tools** in the left-hand navigation frame, and then click **Export Data**.
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)**.
3. Click **Submit**, and then click **Save** and specify a file name and location.

Converting from Demo Mode to Active Mode

To switch from Demo Mode to Active Mode (live operation):

1. Verify that the Local and Remote interfaces are installed properly for live operation, as described in “Installation” on page 29.
2. In the Device Setup page, click **Advanced** in the left-hand navigation frame, and then click **Demo Mode** (available only if the device is installed in Demo Mode).
3. Click **Deactivate Demo Mode**. The device restarts as a registration server and the Remote interface is activated.
4. To specify another WX device as the registration server:
 - a. Click **Registration Server** in the Device Setup page, click **Transfer registration server designation to another device**, specify the IP address of your current registration server, and click **Submit**.
 - b. Click **Registration Server** in the Device Setup page, enter the password of the registration server, and click **Submit**. For more information about the registration server, refer to “Configuring Registration Servers and Communities” on page 82.

Note that to switch back to Demo Mode from Active Mode, you must disconnect the Remote interface, reconnect the Local interface to a mirrored port, load the factory default settings, and then enter the network information and run Quick Setup.

Appendix F

Safety and EMC Certifications

The following table lists the safety and EMC certifications for each type of WX and WXC device.

Table 34: Safety and EMC Certifications for WX Devices


Description	WX 15	WX 20	WX 60	WX 100	WXC 250	WXC 500	WXC 590
Conformity for EMC							
EN 55022 Class A	X	X	X	X	X	X	X
EN 55024 Class A	X	X	X	X	X	X	X
FCC Part 15 Class A	X	X	X	X	X	X	X
EN 61000-3-2	X	X	X	X	X	X	X
EN 61000-3-3	X	X	X	X	X	X	X
Safety Standard							
CAN/CSA-C22.2 No 60950-1-03 - UL 60950-1	X	X	X	X	X	X	X
EN 60950-1	X	X	X	X	X	X	X
Gost	X	X	X	X	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 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 G

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

access control list	List of IP addresses from which an administrator can login to a WX device.
assembly	Process by which a WX device decompresses compressed traffic.
auto-negotiation	A protocol that enables Ethernet systems at the end of a twisted-pair or optical fiber segment to negotiate configuration parameters such as speed, half or full-duplex mode, and use of flow control.
bandwidth	The amount of data that can be sent through a network connection, measured in bits per second (bps).
bridge	A device that partitions a network into separate segments. The bridge allows a packet to be transmitted from one segment to the other only if it is addressed to a host on the other segment.
CLI	See <i>command line interface</i> .
command line interface (CLI)	A method of configuring the WX device by typing in commands via the local serial interface or remote SSH session.
community	Two or more WX devices that can compress and decompress 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.
endpoint	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.
filter	Operator defined IP addresses or TCP port numbers that determine valid addresses or applications for compression processing. A single filter or a list of filters can be defined for each system.
full-duplex	A mode of operation that enables a pair of systems connected by a link to transmit frames to one another at the same time.
gateway	A device that connects and forwards packets between computers or different networks. See also, <i>router</i> .
half-duplex	A mode of operation that allows only a single station to successfully transmit a frame at a given time.
hardware passthrough	Hardware-driven process by which all traffic is passed through the WX device at wire-speed. It is invoked automatically upon disruption.

HTTP	Hyper Text Transfer Protocol. The protocol most often used to transfer information from World Wide Web servers to browsers.
ICMP	Internet Control Message Protocol. An Internet Protocol used to communicate between devices on a network to manage errors and generate control messages.
Interior Gateway Protocol (IGP)	A group of protocols that provide routing information to the routers within an autonomous network.
Internet Protocol (IP)	The protocol that is used to route a data packet from its source to its destination over the Internet.
IP address	A numeric address, such as 10.10.187.22, assigned to every device on the network.
IP subnet mask	A numeric address, such as 255.255.0.0, used to define an IP subnet or to determine membership of an IP address in an IP subnet.
IP subnet	A group of IP addresses defined by the IP address and IP subnet mask pair, such as 10.10.0.0/255.255.0.0.
latency	The time necessary for a packet of data to travel from a source to a destination across a network.
local port	Ethernet port on the back of the WX device. Use to connect to a LAN aggregating switch. <i>See also, remote port.</i>
log	A record of device activity. Logs are recorded for system information, performance, backup, and recovery.
MIB	Management Information Base. A database containing ongoing configuration information and statistics of a device in a network. MIBs are used with SNMP.
MTU	Maximum Transmission Unit. The largest size packet that can be transmitted by a device on a network.
netmap	Compression 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.
operator interface	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.
OSPF	Open Shortest Path First. An interior gateway protocol that routes messages according to the least expensive path.
packet	A unit of data formatted for transmission on a network. Data is broken down into packets for sending over a packet switched network. Each packet has a header containing its source, destination, other control information, and a payload of data to be transmitted.
passthrough mode	A function of the WX device where data passes through at wire-speed upon device disruption or overflow.

ping	A program used to test whether a particular network destination is online, by sending an Internet control message protocol (ICMP) echo request and waiting for a response.
compression rate	The rate of data compression in percentage of a WX device.
compression subnets	The subnets for which a WX device can decompress compressed data. Each WX device advertises its compression subnets to the other devices in the community.
registration server	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.
remote port	Ethernet port on the back of the WX device. Used to connect to WAN router Ethernet port. <i>See also, local port.</i>
response time	The time it takes for a host to respond to a user command.
RIP	<i>See Routing Information Protocol.</i>
round-trip time (RTT)	The time it takes to send a packet to a remote host and receive a response; used to measure delay on a network at a given time.
router	Specialized computer that forwards data packets between networks. Routers can exchange information about their network connectivity (or accessibility) with neighboring network routers using standard routing protocols. This information is used by the router to determine an optimal path for a packet being forwarded.
Routing Information Protocol (RIP)	An interior gateway protocol used in IP networks.
Secure Shell	A program used for secure remote login to a WX device.
Simple Network Management Protocol (SNMP)	The Internet standard protocol for network management software.
Simple Network Time Protocol (SNTP)	A protocol that can synchronize clocks on local computers with time clocks on the Internet.
software passthrough	Software-driven process by which packets are passed through the WX device without any processing.
SSH	see Secure Shell.
static IP address	A permanent IP address for a client, server, or other network device.
Switch	A networking device that sends packets directly to a port associated with a given network address.
TCP	Transmission Control Protocol. The most common Internet transport layer protocol, defined in RFC 793. TCP is connection-oriented and stream-oriented, and provides for reliable communication over packet-switched networks.
tunneling	Encapsulating one type of packet inside the data field of another packet.

User Datagram Protocol (UDP)	User Datagram Protocol. UDP is connectionless and does not guarantee reliable communication; the application itself must process any errors and check for reliable delivery. Defined in RFC 768.
warm reboot	A reboot of the WX device without powering off the unit.
Web Console	A method for configuring and monitoring the statistics of the WX device using a Web browser.

Index

Numerics

- 3DES encryption
 - for IPSec231, 237, 348
 - for SSL optimization240
- 802.1q VLAN support.....67, 345

A

- AAA settings86, 324
- acceleration
 - CIFS traffic218, 326
 - Exchange traffic221, 326
 - Fast Connection Setup.....210, 213, 326
 - Forward Error Correction.....210, 326
 - HTTP traffic.....223, 326
 - reports266
 - resetting traffic flows314, 343
 - TCP Acceleration210, 212, 326
- access control lists388
- access control log file297, 400
- active FTP101
- advertising compression subnets148, 378
- AES encryption
 - for IPSec231, 237, 348
 - for SSL optimization240
- ageout time, device381
- aggregate local WAN speed170
- Application Flow Acceleration
 - about.....214
 - CIFS and Exchange reports.....271
 - CIFS traffic218, 326
 - Exchange traffic221, 326
 - HTTP reports.....272
 - HTTP traffic.....223, 326
- applications
 - about application definitions96
 - accelerating
 - CIFS traffic218, 326
 - Exchange traffic221, 326
 - Fast Connection Setup.....213, 329
 - HTTP traffic.....223, 326
 - TCP Acceleration212, 329
 - assigning to traffic classes.....104
 - common port numbers451
 - compressing.....151, 342
 - defining

- manually99, 332
- using the Traffic report.....275
- defining gateways for377
- defining the IPSec filter240
- defining whether encryption is required.....233
- monitoring105, 151, 351
- summary statistics
 - all traffic258
 - WAN traffic248
 - visibility in tunnels161, 374
- ARP, configuring.....114, 335
- asymmetric routing support for TCP Acceleration...330
- authentication methods, selecting.....87, 324
- automatic installation.....30

B

- backup devices, configuring336
- bandwidth detection181, 195, 367
 - about172
- bandwidth management
 - inbound199, 363
 - outbound, see "outbound QoS"
- baud rate
 - default.....34, 44, 50
 - setting.....339
- BGP routes, polling from a router387
- boot images
 - activating.....290, 319
 - loading.....287, 323
- browser support54
- buttons
 - Bypass/Disable.....35, 39
 - front panel.....38, 460
 - MDI/MDI-X35, 39
- bypass condition, Multi-Path134, 354
- Bypass/Disable button35, 39
- bypass/disable command45, 51, 314
- bytes graph.....257

C

- caching, for HTTP acceleration
 - about216
 - configuring330
- carving out unreachable addresses
 - and outbound QoS177

enable/disable.....	379	remote-routes.....	382
certificates, importing SSL.....	241, 391	rip.....	383
certifications	467	route	384
CIFS acceleration		route-poll.....	387
about.....	215	security.....	388
configuring.....	218, 326	snmp.....	389
reports.....	271	sntp.....	390
circuit speeds		ssl certificate.....	391
and router overhead.....	170	ssl optimization.....	392
configuring.....	180, 193, 367	stack-group	392
Citrix names, in application definitions.....	103	syslog.....	394
classes, traffic		system	394
inbound.....	199, 363	tacplus	395
outbound.....	182, 188, 366	top-talker.....	396
outbound QoS and Multi-Path	104	wan-performance-monitor.....	397
CLI commands		file management	
about		copy	314
basics of using.....	307	list	315
command modes.....	308	remove	320
command summary.....	309	show commands	
entering commands		show acceleration.....	399
from a file.....	323	show access-log.....	400
from a terminal or SSH program.....	305	show event.....	403
from the Web console.....	295	show flow-details	404
configuration		show flow-reset	405
aaa.....	324	show interface.....	405
acceleration	326	show ipsec	406
application	332	show log.....	407
arp	335	show path-mtu-discovery.....	410
backup.....	336	show reduction.....	412
clock	338	show reg-detail	413
console.....	339	top level.....	313
Demo Mode.....	362	activate more.....	313
dns.....	339	commit.....	313
event.....	340	embed	314
filter	342	flow-details.....	404
flow-reset	343	flow-reset	314
interface	344	import-route-table.....	315
ip.....	345	load-config	316
ipsec	346	packet-capture.....	317
license	349	ping.....	318
log.....	350	reboot	319
mon-apps.....	351	reset.....	320
multi-path	352	rollback.....	320
ospf.....	356	save-config.....	321
packet-interception.....	357	set	322
path-mtu-discovery.....	361	shutdown	322
prime-time	361	source	323
qos inbound.....	363	support	323
qos outbound	365	traceroute.....	324
radius.....	370	upgrade	323
reduction.....	371	client devices	
reduction-subnet	378	client-mode command	392
reg-server	380	connecting	52

- disconnecting.....54
- clusters for TCP Acceleration.....330
- CMS
 - about.....28
 - exclusive access to WX devices.....388
- command modes308
- communities
 - defining82, 380
 - deleting devices from382
- community topology.....108, 394
- Compressed traffic class.....199, 363
- compression statistics
 - bytes graph.....257
 - peak compression.....256, 279
 - viewing255
- compression subnets
 - configuring.....148, 378
 - filtering source/destination112
- compression tradeoff for speed.....377
- configuration file
 - displaying.....285, 321
 - loading.....286, 316
 - saving.....283, 321
 - setting to the factory default.....288, 316
- connect timeout, registration server381
- console port
 - baud rate, setting339
 - DB9 cable pin-outs.....426
 - default settings34, 44, 50, 306
- copying files314
- CSV, interpreting results.....437

D

- data packets, Forward Error Correction.....210, 328
- dead interval, OSPF.....357
- dead-time interval, RADIUS90, 370
- decompressors
 - default.....157, 374
 - preferred159, 374
- dedicated WANs.....171
- default decompressors.....157, 374
- default gateway, configuring
 - in CMS30
 - in front panel.....38, 460
 - in Web/CLI63, 345
- default IPsec policy346
- Default traffic class
 - inbound QoS.....199, 363
 - outbound QoS.....188, 366
 - outbound QoS and Multi-Path104
- default user name and password.....34, 44, 50
- Demo Mode
 - about.....453
 - converting to live operation.....466

- defining remote subnets.....362, 461
- excluding traffic to the local subnet.....463
- running Quick Setup461
- virtual devices.....362, 461
- deployment, examples.....22
- DER certificates.....242, 391
- device configuration
 - displaying.....285, 321
 - loading.....286, 316
 - saving.....283, 321
 - setting to the factory default.....288, 316
- device names.....64, 322
- diagnostic files, generating.....299, 323
- diagnostics, traffic flow
 - description of.....447
 - viewing and exporting.....300
- disk access policy, NSC373
- disk icons.....150
- diversion settings, Multi-Path136, 355
- DNS servers, configuring for the Traffic report65
- domain names
 - configuring.....65, 339
 - in flow diagnostics301
 - in the Traffic report.....276
- downgrading to a previous release.....287, 320
- DSCP values, see "ToS/DSCP values"
- dynamic resource allocation (DRA), configuring376
- dynamic routes
 - importing from a file.....78, 315
 - polling from a router.....77, 387
 - using OSPF.....76, 356
 - using RIP76, 383

E

- EMC certifications.....467
- encryption, see "IPSec"
- endpoints
 - compression.....145, 373
 - IPSec.....230, 348
 - Multi-Path135, 354
 - NSC150, 373
 - outbound QoS.....191, 367
 - Packet Flow Acceleration208, 328
 - summary report277
 - WAN performance monitoring.....138, 397
- erasing the disks289, 316
- Ethernet ports, connecting the cables.....37, 46
- events
 - configuring.....140, 340
 - report280
- Exchange acceleration
 - about.....215
 - configuring.....221, 326
 - reports271

Executive report	278	overriding.....	187
exporting data		H	
device performance statistics.....	298, 438	hardware passthrough	35, 39, 45, 51
interpreting performance results	437	hardware passthrough, disabling.....	314
packet capture.....	293, 317	heartbeat packets	
secondary registration server database	382	for all service tunnels.....	375
traffic statistics	274, 396	for high-loss tunnels.....	331
external routing for packet interception.....	118, 357	hello interval, OSPF.....	357
F		high-availability support	65
factory default configuration.....	288, 316	HMAC/SHA-1 authentication	
Fast Connection Setup		for IPSec	231, 237, 348
configuring.....	210, 213, 326	for SSL optimization	240
report.....	268	HTTP acceleration	
fast service tunnels.....	376	about.....	216
filters		configuring.....	223, 326
configuring application.....	151, 342	reports	272
source/destination.....	112, 342	Hub and Spoke topology	111, 394
firewall requirements	30	I	
flow details, viewing.....	404	IANA port map.....	276, 352
flow diagnostics		ICMP settings	384
description of.....	447	icons	
viewing and exporting.....	300	disk	150
Forward Error Correction		endpoint and tunnel	147
configuring.....	210, 326	IPSec status.....	235
report.....	269	Multi-Path status.....	136
fragments, compressing.....	343	on EndPoints Summary report.....	278
front panel		IDEA cipher	240
securing.....	95, 388	idle user timeout.....	93, 324
using the buttons	38, 460	importing routes	
FTP application type	101, 333	by polling a router.....	77, 387
FTP servers, using		from a file.....	78, 315
to copy a packet capture.....	318	inbound QoS	364
to copy system files	314	configuring.....	199, 363
to export diagnostic files.....	323	report.....	264
to import routes	78, 315	inbound speed	201
to load a boot image.....	287, 323	inline deployment	31
to load configuration files	286, 316	installation	
to pre-sync files with NSC	374	automatic	30
to pre-sync files with NSM	162	inline and off-path.....	31
to roll back a boot image	320	post-install tasks.....	59
to save configuration files.....	283, 321	pre-install tasks.....	29
G		WX 100	46
gateways, configuring		WX 15, WX 20, and WXC 250	32
application	377	WX 60 and WXC 500	36
default		WXC 590	40
in CMS.....	30	interface	
in front panel.....	38, 460	link failure propagation	67, 344
in Web/CLI	63, 345	manual mode test	67, 344
in Multi-Path configurations.....	132, 353	periodic mode test	345
general specifications	421	settings, configuring	
guaranteed bandwidths		in front panel.....	39
configuring.....	184, 190, 366	in Web/CLI	65, 344

- statistics.....344, 405
- Intranet traffic class.....199, 363
- IP address, configuring
 - in CMS30
 - in front panel.....38, 460
 - in Web/CLI63, 345
 - secondary address, Multi-Path.....131, 353
- IP compression, meta-packet
 - configuring.....160, 374
 - firewall requirements30
- IPSec
 - configuration procedure.....229, 346
 - defining endpoints.....230, 348
 - defining templates.....236, 347
 - using the Setup Wizard229

J

- JVM support.....54

K

- keep-alive packets
 - for all service tunnels.....375
 - for high-loss tunnels.....331
- key lifetimes.....237, 347
- keyboard shortcuts, CLI308
- keys
 - IPSec.....237
 - OSPF.....76, 356
 - RADIUS.....90, 92, 370, 395

L

- LAN-WAN routing check375
- latency threshold
 - Multi-Path.....136, 355
 - WAN performance monitoring.....139, 397
- Layer 2 multicast traffic.....261
- LEDs, checking
 - WX 10051
 - WX 15, WX 20, and WXC 250.....35
 - WX 60 and WXC 50039
 - WXC 590.....45
- license keys, entering.....69, 349
- lifetimes, IPSec key237, 347
- link failure propagation.....67, 344
- load balancing
 - across routers
 - route-based.....80, 385
 - router-based using ToS.....385
 - across WX 100 clients.....53
 - across WX devices.....155, 374
- loading a boot image287, 323
- local domain name.....65, 339
- local routes
 - about.....73

- adding static.....75, 384
- from OSPF.....76, 356
- from RIP76, 383
- importing from a file.....78, 315
- polling from a router.....77, 387
- local users, defining.....92, 324
- log files
 - access control297, 400
 - system296, 407
- logging in
 - CLI307
 - Web console61
- logging levels.....350
- login retries, SSH.....88, 325
- loss threshold
 - Multi-Path.....355
 - WAN performance monitoring.....397

M

- MAC addresses
 - in ARP entries.....114, 335
 - in decompressed packets.....377
- management traffic, encrypting.....348
- manual and automatic installation30
- marking methods, Multi-Path.....131, 355
- maximum bandwidths
 - inbound201, 364
 - outbound
 - configuring.....184, 190, 366
 - overriding.....187
- MD5
 - for IPSec.....231, 237, 348
 - for OSPF.....356
 - for SSL optimization.....240
- MDI/MDI-X buttons.....35, 39
- Mesh topology.....111, 394
- meta packets
 - configuring size and wait time377
 - disabling multi-packet.....377
 - IP compression.....160, 374
 - wait time377
- metrics, event.....141, 340
- minimum WAN speed.....181, 195, 367
- Molecular Sequence Reduction, see "MSR"
- monitor settings.....54
- monitoring
 - applications.....105, 151, 351
 - virtual endpoints246, 248
 - WAN performance
 - configuring.....138, 397
 - viewing reports.....249
- MSR
 - about.....19
 - symbol size376

MSS override for TCP Acceleration	331
MTU discovery	361, 410
multi-flow emulation	161, 374
Multi-Path configurations	
about	129
defining endpoints	135, 354
defining templates	133, 354
router configuration	137
viewing reports	249
multiple tunnels on the WX 100	144, 392

N

names, special characters in	63
NetFlow records	
generating	294, 396
packet contents	437
network	
cables, connecting	37, 46
interfaces, configuring	
in front panel	39
in Web/CLI	65, 344
settings, configuring	
in front panel	38
in Web/CLI	63, 345
Network Sequence Caching, see "NSC"	
non-WX endpoints	
in Demo Mode	461
outbound QoS	194, 368
NSC	
defining applications	151, 373
defining endpoints	150, 373
disk access policy	373
file pre-synchronization	161, 374
overflow mode	373
NTP, configuring	68, 390

O

off-path deployment	
configuring	116, 357
installing	31
operator access, securing	94, 388
OSPF, configuring	76, 356
outbound QoS	
about	168
and Packet Flow Acceleration	208
bandwidth detection	181, 195, 367
configuration procedure	178
dedicated and oversubscribed WANs	171
defining endpoints	180, 191, 367
defining settings by endpoint	186, 234
defining templates	179, 189, 366
defining traffic classes	104, 182, 188, 366
excluding LAN/WAN addresses	195
non-WX endpoints	194, 368

outbound speed	
about	170
defining	180, 191, 192, 366
report	262
starting and stopping	198
ToS/DSCP values	196, 199, 369
using the Setup Wizard	179
outbound speed	
about	170
defining	180, 192, 366
outbound speed defining	191
overflow mode, NSC	373
overflow, traffic volume	260, 371, 412
oversubscribed WANs	171
P	
packet age-out setting	366
packet capture	
enabling user privilege	93
using	293, 317
Packet Flow Acceleration	
Fast Connection Setup	210, 213, 326
Forward Error Correction	210, 326
reports	266
TCP Acceleration	210, 212, 326
packet fragments, compressing	343
packet interception	116, 357
packet size distribution statistics	261
pass-phrase, IPSec	231, 348
passthrough statistics	260
passwords	
default	34, 44, 50
defining	92, 324
OSPF	76, 356
registration server	83, 381
RIP	77, 118, 383
path MTU discovery	361, 410
peak compression	256, 279
PEM certificates	242, 391
performance data, exporting	298, 438
performance events	
configuring	140, 340
reports	280
performance monitoring, WAN	
configuring	138, 397
viewing reports	249
periodic interface mode test	345
permanent license keys	69, 349
ping utility	291, 318
PKCS12 certificates	242, 391
point-to-multipoint configuration	23
Point-to-Point topology	111, 394
policy routes, defining gateways by application	377
policy-based routing for packet interception	118, 357

- port numbers
 - common application 451
 - in application definitions 102, 333, 404
 - in flow diagnostics 301
 - RADIUS server 90, 92, 370, 395
 - required for TCP and UDP 30
 - viewing on Traffic by Port report 352
- post-installation tasks 59
- preferred decompressors 159, 374
- preferred path 134, 354
- pre-fetch, for HTTP acceleration
 - about 216
 - configuring 330
- pre-installation tasks 29
- pre-synchronization, file 161, 374
- primary boot image 319
- prime time
 - defining 115, 361
 - viewing on reports 246
- privilege level, user 93, 324
- protocols
 - in application definitions 102, 334
 - in flow diagnostics 301, 404
- Q**
 - QoS, inbound
 - configuring 199, 363
 - report 264
 - QoS, outbound, see "outbound QoS"
 - queue lengths
 - inbound QoS 364
 - outbound QoS 366
 - queue processing by ToS/DSCP values 199
- R**
 - rack-mount installation
 - WX 100 46
 - WX 15, WX 20, and WXC 250 32
 - WX 60 and WXC 500 36
 - RADIUS servers and server groups, defining 89, 370
 - RC2 and RC4 ciphers 240
 - rebooting the device 290, 319
 - recompression and tunnel switching
 - about 163
 - enabling 375
 - recovery image 319
 - recovery packets, Forward Error Correction 210, 328
 - redirect age-out setting, ICMP 384
 - registration servers
 - configuring 82, 380
 - configuring in Quick Setup 57
 - deleting devices from 382
 - remote circuit speeds
 - and router overhead 170
 - configuring 180, 193, 367
 - remote routes, viewing 154, 382
 - reports
 - about 245
 - acceleration 266
 - Application Summary
 - all traffic 258
 - WAN traffic 248
 - CIFS and Exchange acceleration 271
 - Compression 255
 - Endpoints Summary 277
 - Events 280
 - Executive 278
 - Fast Connection Setup 268
 - Forward Error Correction 269
 - HTTP acceleration 272
 - Inbound Bandwidth 264
 - Outbound Bandwidth 262
 - Packet Size Distribution 261
 - Passthrough Data 260
 - TCP Acceleration 266
 - throughput
 - all traffic 253
 - WAN traffic 246
 - Traffic 274
 - WAN/Multi-Path performance 249
 - requirements
 - browser and JVM versions 54
 - SSH version 306
 - resetting traffic flows 314, 343
 - retransmissions, RADIUS 90, 92, 370, 395
 - retries, SSH login 88, 325
 - RIP
 - for dynamic routes 76, 383
 - for packet interception 116, 357
 - rolling back to a previous release 287, 320
 - route injection 116, 357
 - router balancing
 - route-based 80, 385
 - router-based using ToS 385
 - router configuration
 - for packet interception 120
 - Multi-Path 137
 - routes
 - adding static 75, 384
 - configuring local 73
 - from OSPF 76, 356
 - from RIP 76, 383
 - importing from a file 78, 315
 - LAN-WAN check 375
 - polling from a router 77, 387
 - remote, viewing 154, 382
 - RSA for SSL key exchange 240
 - RTT

and meta-packet wait times	377
reported by ping	291, 318
reported by traceroute	292, 324
S	
Safe Mode.....	290, 319
safety and EMC certifications	467
sample topologies.....	22
secondary boot image	319
secondary IP address, Multi-Path.....	131, 353
secondary registration server.....	82, 380
secret key, RADIUS.....	90, 92, 370, 395
Secure Shell (SSH), supported version	306
secure wipe	289, 316
security associations	347, 406
security features	86
controlling operator access.....	94, 388
defining local users	92, 324
defining RADIUS servers and server groups.....	89
defining TACACS+ servers.....	91
securing front panel access	95, 388
selecting authentication methods	87, 324
serial port	
baud rate, setting	339
default settings	34, 44, 50
Server/Client Summary.....	303
servers	
DNS.....	65
NetFlow	294, 396
NTP	68, 390
RADIUS.....	89, 371
registration.....	82, 380
syslog.....	72, 394
TACACS+	91
WX 100	
connecting client devices.....	52
disconnecting client devices.....	54
installing.....	46
Server/Client Summary	303
service tunnels	
dynamic resource allocation.....	376
enabling endpoints	145, 373
fast	376
heartbeat packets.....	375
meta packets	377
MSR symbol size	376
source MAC addresses.....	377
statistics.....	371, 412
tunnel switching.....	375
Setup Wizard	
IPSec.....	229
outbound QoS	179
severity levels	
logging.....	350
performance events.....	142, 341
syslog.....	72
system events.....	427
SMB signing	
applying.....	220
disabling.....	221
SNMP	
configuring.....	71, 389
list of traps	427
SNTP, configuring	68, 390
software passthrough.....	35, 39, 45, 51
source address in RADIUS packets.....	371, 396
source MAC address, changing	377
source/destination subnets.....	112
special characters	63
specifications, device	421
Spoke topology	111, 394
SSH interface	
downloading SSH applications	306
enabling and disabling	388
SSL optimization	
enabling applications.....	243, 392
identifying eligible applications.....	101, 333
importing certificates.....	241, 391
overview.....	240
standalone mode, WX 100	393
static routes, adding	75, 384
statistics	
acceleration.....	266
application	
all traffic	258
WAN traffic	248
compression	255
events summary	280
executive summary	278
exporting	298, 438
Fast Connection Setup.....	268
Forward Error Correction.....	269
inbound bandwidth	264
interface	344, 405
interpreting CSV exports.....	437
outbound bandwidth	262
packet size distribution	261
passthrough traffic	260
TCP Acceleration	266
throughput	
all traffic	253
WAN traffic	246
traffic	274
WAN/Multi-Path performance	249
straight-through cable	37, 46
subnet mask, configuring	
in CMS	30
in front panel.....	38, 460

- in Web/CLI 63, 345
- subnets
 - advertising for compression 148, 378
 - defining whether encryption is required... 232, 238, 349
 - discovering 73
 - excluding from compression 112, 342
 - excluding from default decompressors 158, 374
 - excluding from outbound QoS 195, 369
 - filtering flow diagnostics 301
 - filtering the Traffic report 275
 - unadvertised subnets and outbound QoS 177
- support
 - browser and JVM 54
 - generating diagnostic files 299, 323
 - SSH version 306
 - technical 18
- switch-to-wire 21
- symbol size, MSR 376
- syslog
 - configuring 72, 394
 - list of messages 427
- system events
 - configuring 140, 340
 - reports 280
- system log file 296, 407
- system software, upgrading 287, 323
- system-level CLI commands 309

T

- TACACS + servers, defining 91, 395
- TCP
 - required ports 30
 - traffic class 199, 363
- TCP Acceleration
 - clusters 330
 - configuring 210, 212, 326
 - report 266
- technical support 18
- templates
 - IPSec, defining 236, 347
 - Multi-Path, defining 133, 354
 - outbound QoS
 - defining 189, 366
 - names of 188, 365
- terminal emulation program 33, 44, 49, 306
- thresholds
 - event 141, 340
 - Multi-Path 136, 355
 - WAN performance monitoring 139, 397
- throughput statistics
 - all traffic 253
 - WAN traffic 246
- time settings

- manual 68, 338
- NTP server 68, 390
- timeout
 - idle user 93, 324
 - RADIUS server 90, 92, 370, 395
 - registration server 381
- topology
 - sample 22
 - setting 108, 394
- ToS/DSCP values
 - defining by QoS traffic class 196, 369
 - in application definitions 103, 334
 - in Multi-Path configurations 132, 353
 - in UDP heartbeat packets 375
 - processing queues by 199
- traceroute utility 292, 324
- tradeoff, compression for speed 377
- traffic classes
 - inbound 199, 363
 - outbound 182, 188, 366
 - outbound QoS and Multi-Path 104
- traffic flows
 - exporting to CSV file 274, 396
 - resetting for acceleration 314, 343
 - sending to NetFlow server 294, 396
 - viewing details of one flow 404
 - viewing top flows 274
- Traffic report 274
- traps
 - configuring 71
 - list of 427
- TTL response, ICMP 384
- tunnel balancing across WX 100 clients 53
- tunnel mode 160, 374
- tunnel switching
 - about 163
 - enabling 375
- tunnels, service 145, 373
- types of applications 101, 332

U

- UDP
 - and heartbeat packets 331, 375
 - and meta packets 160, 374
 - required ports 30
- unadvertised subnets and outbound QoS 177
- undefined applications, defining
 - manually 99, 332
 - using the Traffic report 275
- upgrading the SRS software 287, 323
- URLs, in application definitions 103, 334
- user names and passwords
 - default 34, 44, 50
 - defining 92, 324

V

validating remote routes.....	155, 382
virtual endpoints	
for outbound QoS	194, 368
in Demo Mode.....	362
in Profile Demo	461
monitoring	246, 248
VLAN 802.1q support	67, 345

W

WAN circuit speeds	
and router overhead.....	170
bandwidth detection.....	181, 195, 367
configuring.....	180, 193, 367
WAN compression subnet	
CLI option	379
for off-path devices.....	149
WAN performance monitoring	
configuring.....	138, 397
viewing reports.....	249
WAN statistics	246
WCCP for packet interception.....	117, 357
Web console	
about.....	62
allowing access.....	94
allowing access by address.....	388
browser and JVM support	54
enabling and disabling	388
logging in	61
monitor settings	54
Weighted Fair Queuing	199, 365
Weighted Strict Priority.....	199, 365
wiping the disks.....	289, 316
Wizard	
IPSec.....	229
outbound QoS	179
WX 100 clients	
and multiple tunnels.....	144, 392
balancing tunnels across.....	53
client-mode command	392
connecting	52
disconnecting	54
Server/Client Summary	303
WX 100 installation.....	46
WX 100, standalone mode.....	393
WX 15, WX 20, and WXC 250 installation	32
WX 60 and WXC 500 installation.....	36
WXC 590 installation	40
WXC devices	
enabling NSC	150, 373
enabling NSC for applications	151, 373
file pre-synchronization	161, 374
rebooting.....	290, 319
wiping the disks	289, 316

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.
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)

All rights reserved.

This package is an SSL implementation written by Eric Young (ey@cryptsoft.com).

The implementation was written so as to conform with Netscape's SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered 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.