



WX Application Acceleration Platforms

WX Central Management System (CMS) Administrator's Guide

Release 5.5

Juniper Networks, Inc.

1194 North Mathilda Avenue

Sunnyvale, CA 94089

USA

408-745-2000

www.juniper.net

Part number 530-021230-01

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

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

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

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

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

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

The following are trademarks of Juniper Networks, Inc.: Active Flow Pipelining, AFP, Application Flow Acceleration, AppFlow, Central Management System, CMS, ERX, E-series, ESP, Fast Connection Setup, Instant Virtual Extranet, Internet Processor, J2300, J4300, J6300, J-Protect, J-series, J-Web, JUNOS, JUNOScope, JUNOScript, JUNOSe, M5, M7i, M10, M10i, M20, M40, M40e, M160, M320, M-series, MMD, Molecular Sequence Reduction, MSR, NetScreen-5GT, NetScreen-5XP, NetScreen-5XT, NetScreen-25, NetScreen-50, NetScreen-204, NetScreen-208, NetScreen-500, NetScreen-5200, NetScreen-5400, NetScreen-IDP 10, NetScreen-IDP 100, NetScreen-IDP 500, NetScreen-Remote Security Client, NetScreen-Remote VPN Client, NetScreen-SA 1000 Series, NetScreen-SA 3000 Series, NetScreen-SA 5000 Series, NetScreen-SA Central Manager, NetScreen Secure Access, NetScreen-SM 3000, NetScreen-Security Manager, Network Sequence Caching, NSC, NMC-RX, Packet Flow Acceleration, PFA, Policy-Based Multipath, PBM, SDX, Stateful Signature, T320, T640, T-series, 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 © 2001-2006, Juniper Networks, Inc. All rights reserved. Printed in USA.

WX Central Management System (CMS) Administrator's Guide, Release 5.4

Revision History
August 2007 —Rev 1

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

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

Year 2000 Notice

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

Software License

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

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

For complete product documentation, please see the Juniper Networks Web site at www.juniper.net/techpubs.

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 NONINFRINGEMENT. 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	13
Audience	13
Document Contents	13
Document Conventions.....	15
Typographical Conventions	15
Technical Support	15
Obtaining Additional Product Information	15
Chapter 1 Introduction	17
About CMS	17
What's New in Version 5.5.....	18
How CMS Works	19
Understanding CMS	20
CMS Support of Device Software Versions	20
Logging In to CMS	20
CMS Web Console Interface.....	21
Where to Go Next	22
Chapter 2 Installing CMS	23
System Requirements	23
Supported Browsers and Character Sets.....	24
Installation Procedure	25
Pre-Installation Tasks	25
Upgrading From a Previous Release	27
Installing WX CMS.....	28
Uninstalling WX CMS.....	31
Logging In for the First Time	31
Using CMS Servers with Multiple Network Interfaces	35
Recommended Configuration Tasks	35
Where to Go Next	36
Chapter 3 Managing Devices	37
Viewing Devices	38
Managing Devices	42
Scheduling Tasks for Multiple Devices	43
Viewing Device Events	44
Loading Device Boot Images.....	45
Rolling Back Device Boot Images	48
Rebooting Devices	50

Viewing Device Configuration Summaries	51
Analyzing Device Configurations	52
Loading Device Configurations	55
Rolling Back Device Configurations	58
Backing Up Device Configurations	60
Restoring Device Configurations	61
Retrieving Device Files	63
Applying a Registration Server Password	65
Putting Devices in Safe Mode	67
Accessing the WXOS Web Console from CMS	68
Exporting Community and Device Information	68
Managing CMS Schedules	69
Managing Scheduled Tasks	69
Exporting a Schedule Log	74

Chapter 4 Managing Device Configurations **75**

Overview of Device Configurations.....	75
Configuration Settings for WXOS 5.x Devices	76
Downloading Global and Partial Configurations	78
Merging Configurations	79
Creating and Editing Configurations	80
Consistency Checking	82
Tracking Configuration Versions	82
Tips for Managing Configurations	82
Viewing Configurations.....	83
Managing Configurations	85
Extracting Configurations	85
Duplicating Configurations	87
Creating New Configurations with Factory Defaults	88
Comparing Configurations	91
Displaying Configurations	92
Viewing Configuration History	93
Publishing Configurations	94
Importing Configurations	95
Exporting Configurations	96
Changing Referenced Configurations	97
Upgrading Configurations	98
Deleting Configurations	98
Defining Configuration Settings	98
Configuring Device Settings.....	103
Configuring Device Addresses.....	103
Defining Communities	105
Configuring Time Zone Settings	107
Configuring the ARP Table	108
Advertising Compression Subnets	109
Defining Outbound QoS Exclusions.....	111

Adding Static Routes	112
Configuring Router Polling	114
Configuring Multi-Path Addresses	116
Configuring the RADIUS Source Address	118
Configuring Basic Setup Parameters	118
Configuring the Interface Settings	119
Configuring NTP	122
Enabling SNMP	123
Defining Syslog Servers	124
Configuring Dynamic Local Routes	126
Enabling Route-Based Router Balancing	128
Designating a Registration Server	130
Generating NetFlow Records	131
Configuring AAA Settings	132
Selecting Authentication Methods	133
Enabling Authorization Checking	134
Defining RADIUS Servers	135
Defining Local Users	137
Securing Operator Access	138
Securing Front Panel Access	140
Configuring Application Settings	141
Default Application Definitions for WXOS 5.4 Devices	142
Viewing the Application Overview	144
Configuring Application Definitions	145
Testing New Application Definitions	149
Assigning Applications to Traffic Classes	149
Monitoring Applications	151
Configuring Compression Settings	152
Configuring Endpoints for Compression	152
Configuring Network Sequence Caching	154
Compressing Applications	156
Configuring Remote Routes	158
Configuring Tunnel Load Balancing Policies	160
Configuring Default Decompressors	162
Defining Preferred Decompressors	165
Configuring Tunnel Mode Settings	167
Configuring QoS Settings	168
Using Outbound QoS to Enhance Performance	169
Understanding Outbound QoS	170
Traffic Classes and Bandwidths	171
QoS Templates and Endpoints	172
WAN Circuit Speeds and Router Overhead	173
Dedicated, Oversubscribed, and Variable Rate WANs	174
Direct Setup Versus Wizard Configuration Results	176
Class Priorities and Excess Bandwidth Allocation	178
ToS/DSCP Values	180
Unadvertised Subnets	180

Procedure for Configuring Outbound QoS Policies	180
Using the Outbound QoS Setup Wizard.....	182
Defining Outbound QoS Settings by Endpoint.....	191
Defining Outbound QoS Templates.....	193
Defining Outbound QoS Endpoints	195
Changing Outbound ToS/DSCP Values	199
Starting and Stopping Outbound QoS.....	202
Configuring Inbound QoS Policies	203
Configuring Traffic Acceleration	206
Overview of Packet Flow Acceleration	206
TCP Acceleration	206
Forward Error Correction.....	208
Fast Connection Setup	208
Overview of Application Flow Acceleration	210
Microsoft CIFS and Microsoft Exchange Acceleration	210
HTTP Acceleration	212
Enabling Acceleration by Endpoint	214
Enabling Acceleration by Application	219
Enabling TCP Acceleration by Application	220
Enabling Fast Connection Setup by Application	221
Enabling Microsoft CIFS Acceleration	222
Enabling Microsoft Exchange Acceleration	225
Enabling HTTP Acceleration	227
Configuring Advanced Setup Parameters.....	229
Configuring Topology Settings.....	229
Selecting a Topology	230
Partial Mesh Example	230
Tiered Network Example	231
Selecting a Community Size	232
Configuring Source/Destination Filters	238
Defining the Prime Time	241
Configuring Packet Interception	242
Methods of Packet Interception	242
Configuring Packet Interception for Off-Path Devices	244
RIP Router/Switch Configuration Commands.....	247
WCCP Router Configuration Commands.....	250
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 133.	251
External Policy-Based Router Commands.....	255
Alternatives to Packet Interception	256
Configuring WAN Performance Monitoring.....	258
Configuring Multiple Tunnels Between WX 100 Servers.....	260
Adding CLI Commands to Configurations.....	262
Configuring Policy-Based Multi-Path	263
Procedure for Configuring Multi-Path	264
Enabling Policy-Based Multi-Path	265

Defining Multi-Path Templates	266
Defining Multi-Path Endpoints	268
Configuring Routers to Support Multi-Path	270
Configuring IPSec	272
Default IPSec Policy	273
IPSec Implementation Details	273
Procedure for Configuring IPSec Policies	274
Defining IPSec Settings by Endpoint	275
Defining IPSec Templates	277
Defining the Default IPSec Policy	279
Configuring Events	280
Chapter 5 Automatic Deployment and License Management	287
About Automatic Deployment	287
Configuring Auto-Deployment	288
Auto-Deployment Procedure	289
Defining Deployment Groups	290
Defining Deployment Records	292
Viewing the Auto-Deployment Status	294
Configuring License Management	296
Licensing Procedure	296
Importing and Validating Authorization Codes	297
Generating and Applying Licenses	299
Viewing the License Status	302
Chapter 6 Monitoring Performance	305
Viewing and Printing Reports	305
Configuring the My WAN Page	306
Viewing Reports on the Monitor Page	311
WAN Statistics	311
WAN Performance Statistics	312
WAN Throughput Statistics	317
WAN Application Summary	319
Compression Statistics	321
Data Compression Statistics	321
Application Summary Statistics	328
Passthrough Statistics	330
Packet Size Distribution Statistics	332
Monitoring Tunnel Status	333
QoS Statistics	336
Outbound QoS Statistics	336
Inbound QoS Statistics	340
Acceleration Statistics	344
Acceleration Summary	344
TCP Acceleration Statistics	346

Fast Connection Setup Statistics	348
CIFS and Exchange Acceleration Statistics	351
HTTP Acceleration Statistics	353
Top Traffic Statistics	355
Executive Summary	358
Trend Reports	360
Events Reports	362
Scheduling Reports	368
Managing Scheduled Reports	369
Chapter 7 Content Management	373
Defining and Distributing Content	373
Defining Distribution Groups	378
Viewing the Status of Content Distribution Tasks	380
Viewing the Schedule Log	382
Accessing Network Drives	383
Chapter 8 CMS Setup and Administration	385
Administering CMS Users	385
Partitioning Users by Customer	386
Viewing My Account and Changing Passwords	387
Defining CMS User Accounts	388
Defining User Groups	391
Defining User Roles	393
Viewing Logged In Users	396
Administering WX Devices	397
Importing and Managing Communities	397
Changing a Registration Server Address or Password	401
Managing Device Groups	402
Configuring Device Polling	405
Configuring Data Retention	407
Uploading a Boot Image	408
Viewing the Polling Catch-Up and Failure Logs	409
Managing Event Forwarding Filters and Email Distribution Lists	411
Defining Email Distribution Lists	412
Defining Event Forwarding Filters	413
Administering CMS	417
Setting WAN Performance Thresholds	418
Entering a Permanent License Key	419
Controlling Client Device Access to CMS	420
Configuring AAA Settings for Remote Authentication	421
Defining the Session Timeout	425
Defining an FTP Server	426
Defining an SMTP Server	427
Enabling Syslog Reporting	429

Stopping and Starting the Scheduler	430
Changing the CMS Server IP Address.....	430
Changing the Web Server Port	431
Viewing System Logs	432
Generating a Diagnostic File	433
Backing Up and Restoring the Database	434
Manual MySQL Database Backups.....	434
Restoring Manual MySQL Database Backups.....	434
Automatic MySQL Database Backups	435
Restoring Automatic MySQL Database Backups	436
CMS Configuration Backups	436
Moving CMS to Another Disk Drive.....	437
Purging Temporary Java Files	438
Changing the CMS Time Zone.....	438
Appendix A CMS Licenses	439
Appendix B System Events	441
Severity Levels	441
CMS Syslog Messages.....	442
WX System Events and SNMP Traps	442
WX Syslog Messages	447
Appendix C Understanding Exported Data Results	453
NetFlow Version 5 Export	453
Performance Statistics Export	455
General Device Information.....	455
Data Section Information.....	456
System Session Statistics	456
Compression Session Statistics	459
Application Session Statistics	459
WAN Statistics	460
Application Flow Acceleration Statistics	460
Bandwidth Management Statistics	461
WAN Performance Statistics	462
Inbound Traffic By Port Statistics.....	463
Top Traffic Export	463
Appendix D Common Application Port Numbers	465
Glossary	467
Index.....	471

About This Guide

Welcome to the Juniper® WX Central Management System (CMS) — a powerful management and configuration tool for the WX and WXC Application Acceleration Platforms. This section describes the audience, organization, and typographical conventions used in this manual.

Audience

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

Document Contents

- Chapter 1, “Introduction” on page 17

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

- Chapter 2, “Installing CMS” on page 23

This chapter describes how to install the CMS software.

- Chapter 3, “Managing Devices” on page 37

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

- Chapter 4, “Managing Device Configurations” on page 75

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

- Chapter 5, “Automatic Deployment and License Management” on page 287

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

- Chapter 6, “Monitoring Performance” on page 305

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

- Chapter 7, “CMS Setup and Administration” on page 385

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

- Appendix A, “CMS Licenses” on page 439

This appendix describes the evaluation and permanent licenses for CMS.

- Appendix B, “System Events” on page 441

This appendix describes the SNMP traps and Syslog messages for the system events generated by CMS and the WX devices.

- Appendix C, “Understanding Exported Data Results” on page 453

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

- Appendix D, “Common Application Port Numbers” on page 465

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

- “Glossary” on page 467

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

Document Conventions

This section describes conventions used throughout this manual.

Typographical Conventions

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

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

Technical Support

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

For technical support, use the following methods:

- Go to *<http://www.juniper.net/support>* and open a case online using the Case Manager.
- Send email to *support@juniper.net*.
- Call +1-888-314-JTAC (U.S, Canada, and Mexico) or +1-408-745-9500

Obtaining Additional Product Information

In addition to this manual, refer to the *CMS 5.5 Release Notes* document enclosed with the product. Also refer to the *WX/WXC Operator's Guide* and the *Quick Start* cards enclosed with each device.

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

Chapter 1

Introduction

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

- “About CMS” in the next section
- “What’s New in Version 5.5” on page 18
- “How CMS Works” on page 19
- “Understanding CMS” on page 20

About CMS

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

- **Cost effective** — CMS reduces the cost of ownership for WX devices by creating a single location from which to manage all devices and leverage configurations on devices already deployed in the network.
- **Eases configuration** — Using CMS, you can quickly and easily configure tens or hundreds of newly deployed devices, modify the configuration of already deployed devices, and view and manage the newly created WAN capacity generated by our Molecular Sequence Reduction (MSR)[™] and Network Sequence Caching (NSC)[™] technology.
- **Simplifies software deployment** — CMS dramatically simplifies the configuration and management of software upgrades. From a single location, and in a single operation, you can upgrade all devices in the same community to a new software version.

- **Creates global policies** — CMS allows network managers to centrally manage and modify global and device-specific configuration settings on all devices. Global settings include basic and advanced setup options, such as for NTP and SNMP, authentication settings, application definitions, outbound QoS settings, and the applications being compressed, monitored, and accelerated.
- **Manage events** — CMS can centrally configure WX system and performance events, monitor CMS and WX event occurrences, and selectively forward events to Syslog servers and/or an email server to generate email alerts.
- **Schedules all tasks** — Using CMS, you can schedule all device management tasks to be performed at the optimal time for the individual location.
- **Centrally views all application acceleration and WAN performance statistics** — CMS provides a single, clear window into the performance of devices around the globe. It presents historical per-tunnel and per-application compression statistics for each device.
- **Centrally views global device and tunnel status** — Using CMS, you can immediately view the status of all tunnels and of each deployed device.

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

What's New in Version 5.5

CMS 5.5 has the following new features:

- Supports devices running WXOS 5.2 through 5.4.
- Centralized content distribution that maximizes user response times by preloading WXC compression dictionaries for large files accessed over the Windows file system (CIFS). For more information, refer to “Content Management” on page 373.
- User roles now provide more detailed control over which features users can access (refer to “Defining User Roles” on page 393).

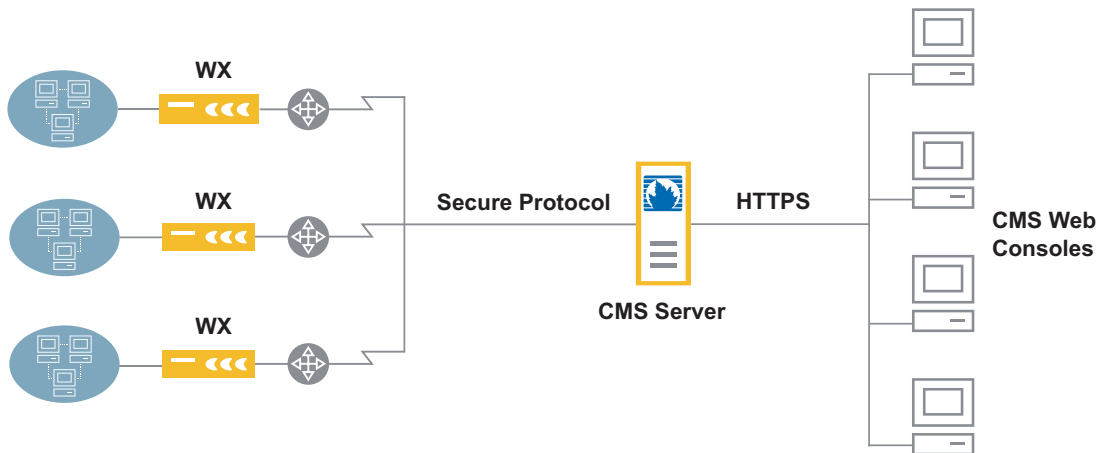
How CMS Works

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

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

Figure 1 shows a logical flow of the communication between the WX devices, a CMS server, and the CMS Web consoles. Configuration data between the devices and the CMS server is securely transmitted via a proprietary protocol (UDP is used to poll devices for some basic status information). Monitoring data is collected from the devices in clear text (compressed). Data between the CMS server and the Web consoles is securely transmitted via HTTPS.

Figure 1: CMS Communication



Understanding CMS

The following sections provide general information about CMS.

- “CMS Support of Device Software Versions” in the next section.
- “Logging In to CMS” on page 20.
- “CMS Web Console Interface” on page 21.

CMS Support of Device Software Versions

CMS 5.5 manages devices running WXOS version 5.2 and greater. Devices running WXOS versions prior to 5.2 are displayed in some Web console pages (such as the Devices page), but they cannot be managed through CMS.

Logging In to CMS

When you log in to the CMS Web console for the first time, you must specify the user name “root” and a default password. The root user account has the CMS administrator role, which provides access to all CMS functions. As a CMS administrator, you can create additional user accounts and specify the level of access for each user, as described in “Defining CMS User Accounts” on page 388.

If two or more users modify the same settings concurrently, the last set of saved changes is used.

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



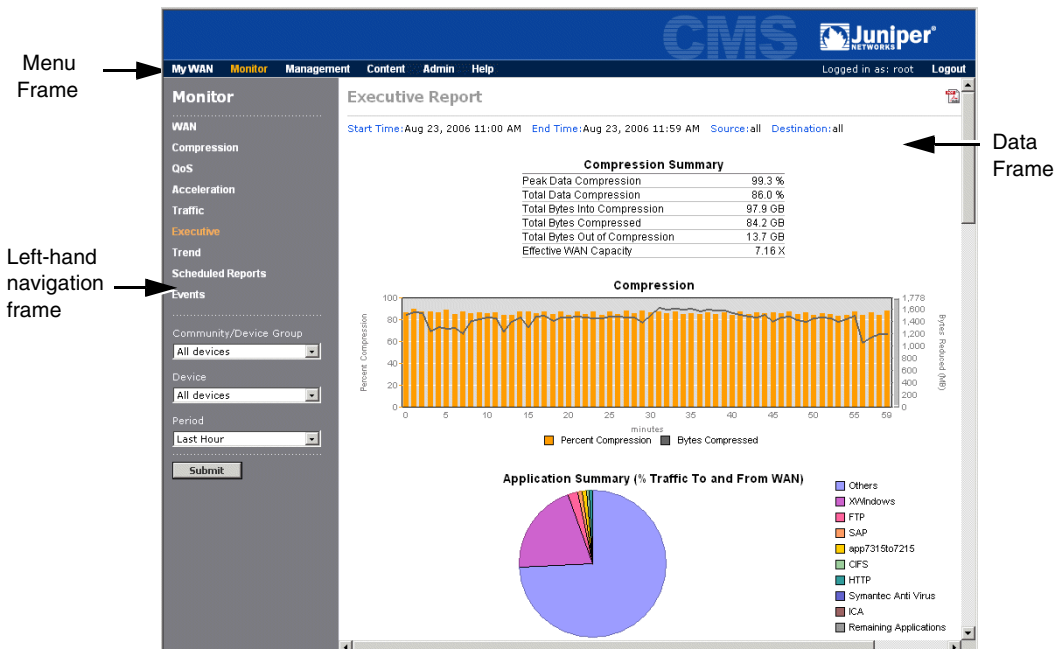
NOTE: If you close the Web browser without logging out, your session remains open until the session timeout time expires.

CMS Web Console Interface

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

- **My WAN** — Select and view a personalized set of performance charts specific to the user account.
- **Monitor** — Monitor tunnel status and performance statistics.
- **Management** — Manage devices, configurations, automatic deployment, and scheduled tasks.
- **Content** — Maximize response times by preloading compression dictionaries for large files and Web content.
- **Admin** — Administer CMS, such as add and delete user accounts, and import communities.
- **Help** — Open the About window, which includes the CMS software version and license information, or open a PDF version of this manual.
- **Logout** — Log out of the CMS Web console.

Figure 2: CMS Web Console Interface



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

Where to Go Next

You are ready to install the CMS software. Proceed to “Installing CMS” on page 23.

Chapter 2

Installing CMS

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

- “System Requirements” in the next section
- “Supported Browsers and Character Sets” on page 24
- “Installation Procedure” on page 25
- “Logging In for the First Time” on page 31
- “Using CMS Servers with Multiple Network Interfaces” on page 35
- “Recommended Configuration Tasks” on page 35

System Requirements

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



NOTE: CMS should be installed on a dedicated server that is not running any other applications, particularly database backup and anti-virus software, which may lock files and interfere with the CMS database. For online database backups, use the backup utility provided with CMS (refer to “Backing Up and Restoring the Database” on page 434).

- Microsoft Windows Server 2003 Standard Edition, Microsoft Windows Server 2003 Enterprise Edition with Service Pack 1, or Windows Server 2000 with Service Pack 4
- The following table shows the recommended and minimum CPU, memory, and disk space requirements for each range of devices being managed. These estimates assume a dedicated server with high speed drives, and a 60-minute polling interval.

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

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

Supported Browsers and Character Sets

CMS can be administered through the CMS Web console using Microsoft Internet Explorer version 6.0 or later. All other Web browsers are currently not supported. Only English characters are supported.

Verify that your browser accepts cookies (required to log in), and that the server is always checked for the latest configuration information:

1. Select Internet Options from the Tools menu.
2. Click Settings under Temporary Internet Files, select Every visit to the page, and click OK.
3. Click the Privacy tab and verify that the setting is Medium High or lower.
4. Click the Security tab, click Default Level, and verify that the setting is Medium or lower.

5. Click the Advanced tab, and verify that Play animations in web pages is selected (required for in-progress indicators on the Devices page).

On Windows Server 2003, the browser's default Security setting is High, which will prevent you from logging in to the Web console. Change the Security setting to Medium.



NOTE: The Executive report does not display correctly if the browser uses the Sun Java Virtual Machine. To display the Executive report correctly, use the Internet Explorer JVM.

Installation Procedure

The following sections describe how to install and uninstall WX CMS 5.5:

- “Pre-Installation Tasks” in the next section
- “Upgrading From a Previous Release” on page 27
- “Installing WX CMS” on page 28
- “Uninstalling WX CMS” on page 31

Pre-Installation Tasks

If you are installing WX CMS for the first time, complete all of the following pre-installation tasks:

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



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

- Verify that the system date, time, and time zone are accurate for your location. In addition to the time zone setting in the Windows Date/Time properties dialog box, check the time zone environment variable. Refer to your Microsoft Windows documentation for more information.

If you change the time zone after CMS is installed, you must restart the JuniperCMS service (refer to “Changing the CMS Time Zone” on page 438).

- Verify that the following ports are available and are not blocked by firewalls or other devices:

Ports	Description
443 or 8443	Port 443 is the default port used by the CMS Web server. If another server uses port 443, such as IIS, disable that server or specify port 8443 for the CMS Web server during installation. Port 443 or 8443 is required to support auto-deployment of WX devices.
3577 and 3578	<p>Ports 3577 and 3578 (both TCP and UDP) are used to communicate with the WX devices. WXOS uses 3577 as the destination port to send acknowledgements and statistics to CMS. If port 3577 is in use, you are prompted to enter another port number during CMS installation.</p> <p>If necessary, you can change the port number after installation as follows:</p> <ol style="list-style-type: none"> 1. Open a command window on the CMS server. 2. Enter the following commands: <pre>cd <CMS installation directory> /mysql/bin mysql -admin -peri use cmsdata; update ScheduleTable set StatsRecvPort = <port # >; quit</pre>



NOTE: On Windows Server 2003, ports 443 and 3577 are blocked by default, and must be added to the firewall exception list.

- Determine if the Sun™ Microsystems™ Java™ Runtime Environment (JRE™), which is a component of the Java 2 Platform, Standard Edition (J2SE™), is on your system. If JRE version 1.5.0 is not present, the CMS installation wizard will install it.
- Reserve a static IP address for the CMS server. If you later change the IP address of the CMS server, you must reboot the server and obtain a new license key before you can continue using CMS.
- The Microsoft FTP Server must be installed and running on the CMS server.

To install the FTP Server on Windows Server 2003:

1. Install the FTP service as described at <http://support.microsoft.com/?kbid=323384>.

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

To install the FTP Server on Windows Server 2000:

1. Click Start > Settings > Control Panel, and double-click Add/Remove Programs.
 2. Double-click Add/Remove Windows Components.
 3. Select Internet Information Services (IIS), and click Details.
 4. In the IIS window, select the check box for File Transfer Protocol Server and click OK.
 5. Click Next to install the service. When prompted, insert the Microsoft Windows 2000 Server CD into the CD drive.
- Verify that a Syslog server is not installed. CMS now includes a Syslog server to collect events from WX devices, but the server will not start if another Syslog server is already running.
 - Verify that numbers are formatted using a period as the decimal point and a comma as the separator (“1,234.56”). For example, to change the number formats on Windows XP, click Regional and Language Options in the Control Panel, click Customize, and select a period for the decimal symbol and a comma for the digit grouping symbol.

Upgrading From a Previous Release

To upgrade from a previous release of WX CMS:

1. To upgrade from CMS 5.1 or earlier, you must first upgrade to CMS 5.2.
2. To upgrade from CMS 5.2 through CMS 5.4, do the following (the upgrade will take several minutes and cannot be reversed):
 - a. Stop the CMS service on the server:
 - a. Click Start > Run, enter “services.msc” and click OK.
 - b. In the Services window, right-click on JuniperCMS and click Stop.
 - c. Back up the database, including all configuration files, as described in “Backing Up and Restoring the Database” on page 434.

- d. Install WX CMS 5.5, as described in “Installing WX CMS” in the next section).
3. If you upgrade from CMS 5.2, and you are using SSL to access an Active Directory server for remote authentication, you must reimport the SSL certificate after you install CMS 5.5 (click Admin > Setup > AAA > Import Certificate).

Installing WX CMS

To install the WX CMS software:

1. Log in to the Microsoft Windows 2000 or 2003 server as a user with administrator privileges. Next, close all windows and exit all programs, including any anti-virus programs running on the desktop.
2. Insert the CMS CD into the server's CD drive.

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



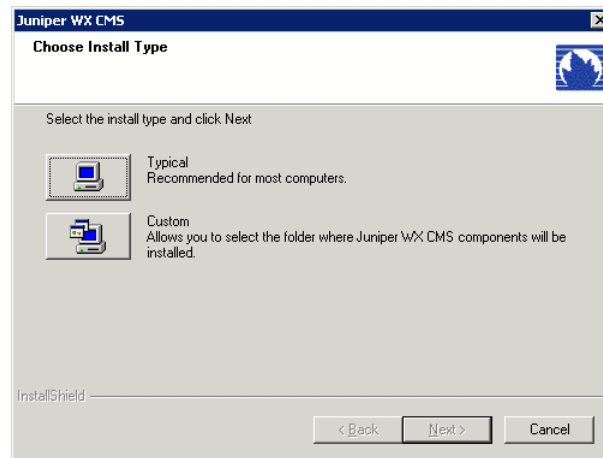
NOTE: Click OK for all security or AntiSpyware prompts encountered during the installation and initial setup of CMS.

3. Click Next. The CMS license agreement appears. Read the agreement carefully. To accept the terms of the agreement, click Yes. The Customer Information window opens (Figure 3).

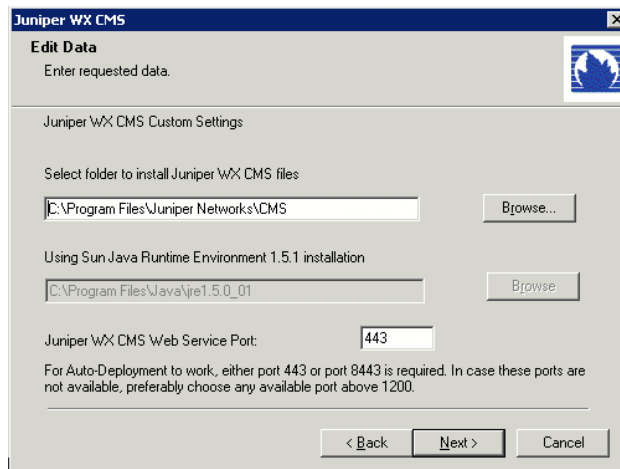
Figure 3: Entering Customer Information

4. Enter customer information:
 - a. Enter a user and company name.
 - b. If you have a permanent license key, enter it in the License Key field. If you do not have a permanent license key, leave “Evaluation” in the License Key field. For more information about licenses, refer to “CMS Licenses” on page 439.
 - c. Click Next. The Choose Install Type window opens (Figure 4).

Figure 4: Selecting the Installation Type



5. Select a Typical or Custom installation, as follows:
 - Click Typical to do the following:
 - Install the CMS files in *C:\Program Files\Juniper Networks\CMS*.
 - Install JRE version 1.5.0 in *C:\Program Files\Java\j2rel.5.0_06* if it is not already installed on your system.
 - Set the Web server port to 443 (the default HTTPS port). If port 443 is currently used by IIS or some other Web server, you are prompted to enter another port number (enter port 8443).
 - To change any of the default settings, click Custom: to open the Custom Settings window (Figure 5).

Figure 5: Customizing CMS Installation

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

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

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

Before restarting the system, remove any disks or CDs from the drives, and verify that the MySql and JuniperCMS services are installed (click Start > Run, enter "services.msc" and click OK). If the services are marked as disabled, reboot the system and repeat the installation procedure.

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



NOTE: Devices running WXOS versions prior to 5.2 are listed on some CMS pages, such as the Devices page, but they are not fully supported by CMS 5.5.

Uninstalling WX CMS

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

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

After you uninstall CMS, verify that the MySQL and JuniperCMS services have been removed. If they are marked as disabled, rebooting the system should remove them.

Logging In for the First Time

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

You can log into the CMS Web console from any workstation in your network. The Web console supports Microsoft Internet Explorer version 6.0 and greater. Data is securely transmitted from the CMS server to the Web browser via HTTPS.

To log in to the CMS Web console:

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

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

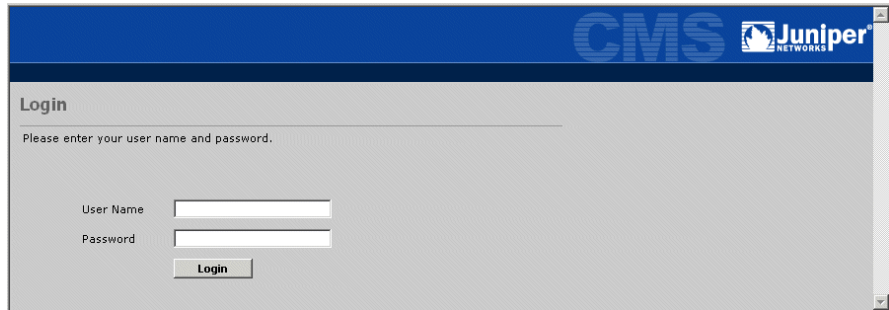
Be sure to use “https” instead of “http”. Also, if you have changed the CMS Web server port number from 443 to 8443, you must include “:8443” after the IP address. For example:

`https://10.10.0.1:8443`

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

2. Depending on your Web browser settings, the Security Alert dialog box may appear. Click Yes to open the Login page (Figure 6).

Figure 6: Logging In For The First Time

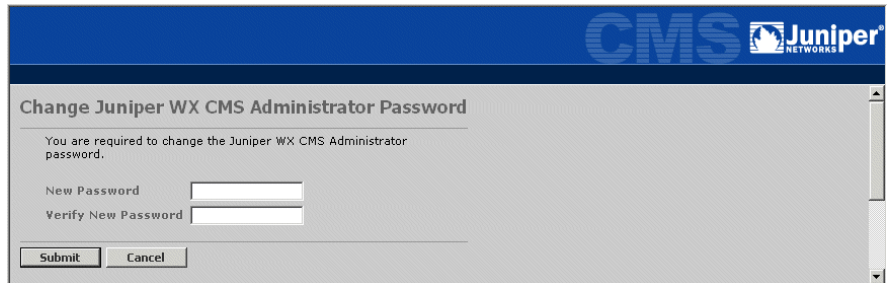


3. Enter the following user name and password, and click **Log In**.

- User name: root
- Password: juniper

The Change CMS Administrator Password page opens (Figure 7).

Figure 7: Changing the Default Password



4. Enter a new password for the root user in the New Password and Verify New Password fields, and click Submit. The password is case-sensitive.

5. A blank Communities page opens. To manage the WX devices in each community, you must import the communities defined on each device that acts as a registration server. For more information about registration servers, refer to “Designating a Registration Server” on page 130. If you do not yet have a registration server, you can import communities at a later time (refer to “Importing and Managing Communities” on page 397).

To import communities into CMS:

- a. Click Import to open the Communities > Import page (Figure 8).

Figure 8: Importing Communities to CMS

- b. Specify the following information, and click Next.

Public IP Address	Enter the IP address of a registration server.
Private IP Address	If the registration server is on the private side of a NAT device, enter the registration server's private IP address. All WX devices that use the registration server must also be on the private side of the same NAT device.
Password	Enter the registration server's password.

- c. Select the check box next to each community you want to import, and click Next.

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

- d. Select at least one user group that can access the imported devices, or click New User Group to define a new user group (refer to “Defining User Groups” on page 391), and then click Next.
- e. Click Finish. If the registration server is behind a NAT device, the Public IP Addresses page opens.

Figure 9: Entering Public IP Addresses

The devices listed below are located behind a NAT router. In order to access the devices, you are required to enter their public IP addresses.

If any public IP address fields are highlighted in yellow below, the device reports a conflicting private IP address or registration server. Check that the public/private address mapping is correct.

You may not change the public IP address for a registration server here. Use the Registration Server page accessible from the Communities page instead.

Showing 3 endpoints

Device Name	Private IP Address	Public IP Address
CMSQA-244	10.91.106.100	10.87.244.2
SR-192.168.101.100-NAT	192.168.101.100	10.87.241.2
CMSQA-243	192.168.243.2	10.87.243.2

Submit Cancel Clear Reset

- f. Enter the public IP address for each device listed, and click Submit. If you omit the public address, you can view the device in CMS and add it to device groups, but you cannot perform any other operations on the device.

The CMS quick setup is complete. If the server where you installed CMS has multiple network interface cards (NICs), refer to “Using CMS Servers with Multiple Network Interfaces” on page 35. To perform additional administrative tasks, refer to “Recommended Configuration Tasks” on page 35.

Using CMS Servers with Multiple Network Interfaces

If the CMS server has multiple network interface cards, you must specify which IP address is used by FTP on the remote WX devices to download WXOS images from the CMS server, even if only one interface is used (otherwise, the address is chosen randomly).

To specify the IP address:

1. Stop the JuniperCMS service:
 - a. Click Start > Run, enter “services.msc” and click OK.
 - b. In the Services window, right-click on JuniperCMS and click Stop.
2. In the following file, change the “CmsIP = ” line to specify the correct IP address:

`<Install>\CMS\data\preference\Installinfo.properties`

Where “< Install > ” is the location where CMS is installed. The default location is “C:\Program Files\Juniper Networks\CMS”.

3. Restart the JuniperCMS service.

Recommended Configuration Tasks

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

Task	Description
Synchronizing Clocks	<p>The clocks on all devices, including the CMS server, should be synchronized to the same Simple Network Time Protocol (SNTP) server or server hierarchy. To use CMS to configure an NTP server for your devices, refer to “Configuring NTP” on page 122.</p> <p>If you do not have an SNTP server in your network, you can use the address of the CMS server. During installation, CMS enables the Windows SNTP server. Be sure to verify that port 123 (UDP) is not blocked by firewalls or other devices.</p> <p>If you use some other SNTP server to synchronize your devices, the Windows SNTP agent on the CMS server should be pointed to the same SNTP server.</p>

Task	Description
Importing Communities	If you have multiple registration servers, import the communities from each server, as described in “Importing and Managing Communities” on page 397.
Defining User Groups	Create one or more user groups, and assign all users and imported communities to at least one user group, as described in “Defining User Groups” on page 391.
Using NAT	<p>If a registration server and its communities are on the private side of a NAT device, note the following:</p> <ul style="list-style-type: none"> ■ You can add or change a device's public address from the Communities page or the Devices page (refer to “Viewing Devices” on page 38). ■ If two registration servers have the same private address, their default communities have the same name. To distinguish the communities, change the private IP address of one of the registration servers, or, on one of the registration servers, move all devices to a non-default community and import that community. ■ If other WX devices have the same private address, change the device names to avoid confusion.
Uploading Boot Images	Upload WXOS boot images to the CMS server, as described in “Uploading a Boot Image” on page 408. An uploaded image can then be downloaded to selected devices.
Analyzing Configurations	<p>Use CMS to retrieve and analyze the differences between the configurations of selected devices. Extracted configurations can be used as a starting point in managing your device configurations. For more information about analyzing a configuration, see “Analyzing Device Configurations” on page 52. For more information about extracting a configuration, see “Extracting Configurations” on page 85.</p> <p>Analyzing configurations also helps you select a global configuration that can be modified and loaded on other devices. For more information about modifying and loading a configuration, see “Defining Configuration Settings” on page 98 and “Loading Device Configurations” on page 55.</p>

Where to Go Next

To view the devices that CMS discovers for the Community, refer to “Managing Devices” on page 37. To create user accounts or perform additional administrative functions, proceed to “CMS Setup and Administration” on page 385.

Chapter 3

Managing Devices

This chapter describes how to use CMS to centrally manage communities of WX devices. It covers the following topics:

- “Viewing Devices” in the next section
- “Managing Devices” on page 42
- “Accessing the WXOS Web Console from CMS” on page 68
- “Exporting Community and Device Information” on page 68
- “Managing CMS Schedules” on page 69

Viewing Devices

The Devices page lets you view the devices in each community or device group associated with your user groups, execute tasks for selected devices, and open the WXOS Web console for a specific device. You can also change the public IP address associated with a device's private NAT address.

To view the devices in each community or device group:

1. Click Management in the menu frame.
2. Select a community or device group from the Community/Device Group menu, and click Submit.

Figure 10: Viewing Devices



The screenshot displays the 'Devices' page in the Juniper CMS. The left sidebar shows the 'Management' menu with sub-items: Devices, Configurations, Auto-Deployment, License Management, Scheduled Tasks, and Schedule Log. The 'Devices' sub-item is selected. The main content area shows a table of devices. At the top of the table, there are links: 'Select All', 'Clear', 'Refresh', 'Export...', and 'Legend...'. Below these links is a dropdown menu for 'Apply to:' set to 'Checked devices only' and a 'GO' button. The table has 7 columns: Device Name, IP Address, Status, Duties, Model, SW, and License. There are 7 rows of device data. The first row has a checkbox, 'SR-10.87.240.2', '10.87.240.2', a green status icon, a 'Duties' icon, 'SR-50 -5.0', '5.4.3.6', and '1 Mbps'. The second row has a checkbox, 'WX-10.87.245.2', '10.87.245.2', a red status icon, a 'Duties' icon, 'SR-20 -2.0', '5.4.3.6', and '2 Mbps'. The third row has a checkbox, 'WX-10.87.246.2', '10.87.246.2', a green status icon, a 'Duties' icon, 'WX-100 -3.0', '5.4.3.6', and 'Eval'. The fourth row has a checkbox, 'WX-10.87.248.2', '10.87.248.2', a yellow status icon with a question mark, a 'Duties' icon, 'SR-20 -2.0', '5.4.1.0_pC*', and 'Eval'. The fifth row has a checkbox, 'WX-10.87.249.2', '10.87.249.2', a red status icon, a 'Duties' icon, 'WX-100 -2.0', '5.4.3.6', and 'Eval'. The sixth row has a checkbox, 'WX-10.87.247.2', '10.87.247.2', a green status icon, a 'Duties' icon, 'SM-250 -1.0', '5.4.3.6', and '2 Mbps'. The seventh row has a checkbox, 'WXC-500 10.87.242.2', '10.87.242.2', a yellow status icon, a 'Duties' icon, 'WXC-500 -1.5', '5.4.3.6', and 'Eval'. At the bottom of the sidebar, there is a 'Community/Device Group' dropdown set to 'default-10.87.240.2', a 'Show Exceptions Only' checkbox, and a 'Submit' button.

Device Name	IP Address	Status	Duties	Model	SW	License
<input type="checkbox"/> SR-10.87.240.2	10.87.240.2	↑ ●	⊕ ⬇ QoS	SR-50 -5.0	5.4.3.6	1 Mbps
<input type="checkbox"/> WX-10.87.245.2	10.87.245.2	↑ ●	⊕ QoS	SR-20 -2.0	5.4.3.6	2 Mbps
<input type="checkbox"/> WX-10.87.246.2	10.87.246.2	↑ ●	⊕ QoS	WX-100 -3.0	5.4.3.6	Eval
<input type="checkbox"/> WX-10.87.248.2	10.87.248.2	↑ ●?	⊕ QoS	SR-20 -2.0	5.4.1.0_pC*	Eval
<input type="checkbox"/> WX-10.87.249.2	10.87.249.2	↑ ●	⊕ QoS	WX-100 -2.0	5.4.3.6	Eval
<input type="checkbox"/> WX-10.87.247.2	10.87.247.2	↑ ●	⊕ QoS	SM-250 -1.0	5.4.3.6	2 Mbps
<input type="checkbox"/> WXC-500 10.87.242.2	10.87.242.2	↑ ●	⊕ ⬇ QoS	WXC-500 -1.5	5.4.3.6	Eval

3. If the device polling takes too long, you can click Stop in the upper-right corner of the page. Polling continues, but you can then execute tasks for one or more devices, such as loading a new boot image, as described in “Managing Devices” on page 42.

From the Devices page, you can also:

- Click a highlighted IP address to change the associated public address. A **NAT** in the Status column indicates the public address is missing or incorrect. To change a registration server's address, refer to “Changing a Registration Server Address or Password” on page 401.
- View the status, hardware model, software version, license, and functional properties of each device. An asterisk (*) next to the software version indicates that the device did not respond to the last query.

- Click Legend to view a brief description of the icons used on the Devices page (refer to Table 1 for more information).
- Click Refresh to view the latest device status information. Click the column headers to change the sort.
- Click Show Exceptions Only to view only devices that are not responding or that have:
 - Events that occurred in the last 60 minutes (WXOS 5.4 or later devices) or the last 24 hours (WXOS 5.3/5.2 devices).
 - Tasks that failed, an expired license, or a registration server password that does not match the password defined in CMS
- Click  or a colored icon (●●●●) in the Status column to view device events, as described in “Viewing Device Events” on page 44. The triangle icon is shown only for WXOS 5.3/5.2 devices. The colored icons indicate the highest severity event that has occurred on the device (WXOS 5.4 or later devices only).
- Open the WXOS Web console for a device by clicking the device name, as described in “Accessing the WXOS Web Console from CMS” on page 68.
- Click Export to export the device information to a CSV file, as described in “Exporting Community and Device Information” on page 68.
- Click  in the SW column to view the details of failed device tasks, as described in “Managing CMS Schedules” on page 69. Note that the status of content distribution tasks is reported separately (refer to “Viewing the Status of Content Distribution Tasks” on page 380).



NOTE: Devices with WXOS versions prior to 5.2 are displayed without a check box to indicate that they cannot be managed with CMS.

Table 1: Device Icons
































Icon	Table Column	Description
	Status	The device is up and running. Move the cursor over the icon to view the length of time since the device was started (or restarted).
	Status	The device is operating in safe mode.
	Status	Registration server password on the device does not match the password in CMS.
	Status	Indicates critical- or error-level events have occurred in the last 24 hours (WXOS 5.3/5.2 devices only). Move the cursor over the icon to view the number of events. Click the icon to view the event details.
	Status	Indicates the highest severity event that occurred on the device in the last 60 minutes (WXOS 5.4 or later devices only). Move the cursor over the icon to view the number of events. Click the icon to view the event details.
	Status	Storage space on the device is low. Move the cursor over the icon to view the number of bytes remaining.
	Status	The device is not reachable, has never responded. May occur when a private NAT address has an incorrect public address. An asterisk (*) after the software version indicates no response to the last query.
	Status	The private IP address has a missing or incorrect public address. Click the private address to enter or change the public address. To change a registration server's address, refer to "Changing a Registration Server Address or Password" on page 401.
	Status	The device is NOT using an NTP server to maintain an accurate device time. An NTP server is strongly recommended to ensure the accuracy of hourly reports (refer to "Configuring NTP" on page 122).
	Duties	The device is a hub in a Hub and Spoke topology.
	Duties	The device is a spoke in a Hub and Spoke topology. By default, a spoke compresses and decompresses data only for the hub devices.
	Duties	The device is part of a mesh topology.
	Duties	The device is the primary registration server.
	Duties	The device is the secondary registration server.
	Duties	The device is a backup for one or more devices. The icon flashes when the backup device is active.
	Duties	The device is part of a multi-node configuration.
	Duties	Outbound Quality of Service (bandwidth management) is enabled.

Table 1: Device Icons

Icon	Table Column	Description
	Duties	Packet Flow Acceleration is enabled.
	Duties	Network Sequence Caching (NSC) is enabled.
	Duties	Policy-Based Multi-Path is enabled.
	Duties	IPSec encryption is enabled.
	Duties	Packet interception using RIP is enabled.
	Duties	Packet interception using WCCP is enabled.
	Duties	Packet interception using external routing is enabled.
	Device	Indicates that you cannot change the device's boot image or configuration
	SW	One or more scheduled tasks is pending (does not include content distribution tasks). Click the icon to view more information about the scheduled tasks.
	SW	One or more scheduled tasks has failed (does not include content distribution tasks). Click the icon to view more information about the failed tasks.
	License	License key has an expiration date. Move the cursor over the icon to view the number of days remaining.
	License	License key expires in 3 to 10 days.
	License	License key expires in three days.
	License	Licensed throughput is exceeded.

Managing Devices

The following sections describe the device management tasks:

- “Scheduling Tasks for Multiple Devices” in the next section
- “Viewing Device Events” on page 44
- “Loading Device Boot Images” on page 45.
- “Rolling Back Device Boot Images” on page 48
- “Rebooting Devices” on page 50
- “Viewing Device Configuration Summaries” on page 51
- “Analyzing Device Configurations” on page 52
- “Loading Device Configurations” on page 55
- “Rolling Back Device Configurations” on page 58
- “Backing Up Device Configurations” on page 60
- “Restoring Device Configurations” on page 61
- “Retrieving Device Files” on page 63
- “Applying a Registration Server Password” on page 65
- “Putting Devices in Safe Mode” on page 67

Scheduling Tasks for Multiple Devices

Most device tasks, such as loading configurations and retrieving statistics, can be applied to multiple devices in a selected community or device group. On the Devices page, there are two ways to select the devices for a task:

- Select the check box next to the appropriate devices or click Select All at the top of the page to select all devices on the page. Only the selected devices are affected. If the task is scheduled, you can manually add devices to the scheduled task at a later time (refer to “Managing Scheduled Tasks” on page 69).
- Select Community/Device Group from the Apply To menu (Figure 11). The devices in the community or device group are determined dynamically at execution. This is particularly useful for recurring schedules, where the devices in a community or device group may change over time.

Figure 11: Devices Page

Task menu

Apply To menu

My WANMonitorManagementContentAdminHelp

Logged in as: rootLogout

Management

Devices

Configurations

Auto-Deployment

License Management

Scheduled Tasks

Schedule Log

Community/Device Group

default-10.87.240.2

Show Exceptions Only

Submit

Devices

Select AllClearRefreshExport...Legend...

Devices: 7

Device Name	IP Address	Status	Duties	Model	SW	License
<input type="checkbox"/> SR-10.87.240.2	10.87.240.2	↑	⊕ QoS	SR-50 -5.0	5.4.3.6	1 Mbps
<input type="checkbox"/> WX-10.87.245.2	10.87.245.2	↑	⊕ QoS	SR-20 -2.0	5.4.3.6	2 Mbps
<input type="checkbox"/> WX-10.87.246.2	10.87.246.2	↑	⊕ QoS	WX-100 -3.0	5.4.3.6	Eval
<input type="checkbox"/> WX-10.87.248.2	10.87.248.2	?	⊕ QoS	SR-20 -2.0	5.4.1.0_pC*	
<input type="checkbox"/> WX-10.87.249.2	10.87.249.2	↑	⊕ QoS	WX-100 -2.0	5.4.3.6	Eval
<input type="checkbox"/> WXC-10.87.247.2	10.87.247.2	↑	⊕ QoS	SM-250 -1.0	5.4.3.6	2 Mbps
<input type="checkbox"/> WXC-500 10.87.242.2	10.87.242.2	↑	⊕ QoS	WXC-500 -1.5	5.4.3.6	Eval

Viewing Device Events

To view device events from the Devices page:


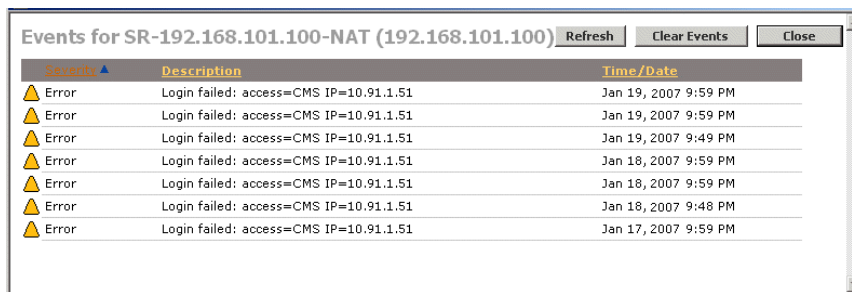
1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. For WXOS 5.3/5.2 devices, click  in the Status column to open the Events window (Figure 12). The event icon is displayed if critical- or error-level events occurred on the device in the last 24 hours.

Figure 12: Viewing Events for WXOS 5.3/5.2 Devices



Severity	Description	Time/Date
Error	Login failed: access=CMS IP=10.91.1.51	Jan 19, 2007 9:59 PM
Error	Login failed: access=CMS IP=10.91.1.51	Jan 19, 2007 9:59 PM
Error	Login failed: access=CMS IP=10.91.1.51	Jan 19, 2007 9:49 PM
Error	Login failed: access=CMS IP=10.91.1.51	Jan 18, 2007 9:59 PM
Error	Login failed: access=CMS IP=10.91.1.51	Jan 18, 2007 9:59 PM
Error	Login failed: access=CMS IP=10.91.1.51	Jan 18, 2007 9:48 PM
Error	Login failed: access=CMS IP=10.91.1.51	Jan 17, 2007 9:59 PM

The Events window displays the severity level, description, and date and time of the last 20 events. Event times are in the local CMS server time, not the device time. From the Events window, you can:





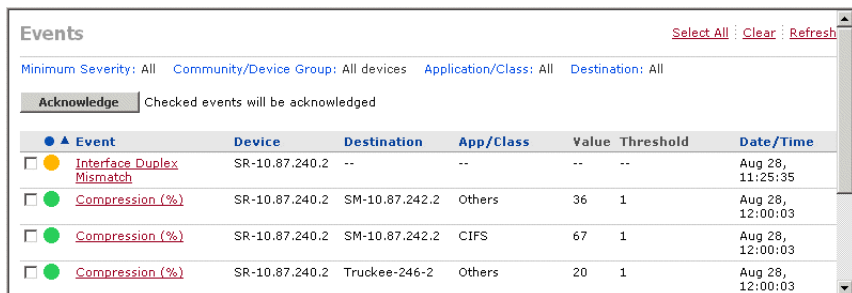




- Update the list by clicking Refresh.
 - Delete all events and remove the  from the Devices page by clicking Clear Events. This also clears the events on the device.
3. For WXOS 5.4 devices, click a colored icon (  ) in the Status column to open the Events window (Figure 13). The color indicates the highest severity CMS “Device Task Failed” event, WX system event, or WX performance event that occurred in the last 60 minutes.

Figure 13: Viewing Events for 5.4 Devices



Event	Device	Destination	App/Class	Value	Threshold	Date/Time
<input type="checkbox"/>  Interface Duplex Mismatch	SR-10.87.240.2	--	--	--	--	Aug 28, 11:25:35
<input type="checkbox"/>  Compression (%)	SR-10.87.240.2	SM-10.87.242.2	Others	36	1	Aug 28, 12:00:03
<input type="checkbox"/>  Compression (%)	SR-10.87.240.2	SM-10.87.242.2	CIFS	67	1	Aug 28, 12:00:03
<input type="checkbox"/>  Compression (%)	SR-10.87.240.2	Truckee-246-2	Others	20	1	Aug 28, 12:00:03

From the Events window, you can:

- Select the check box next to the appropriate events and click Acknowledge. Acknowledged events are removed from the list.
- Click an event name to view additional event details. Selecting a CMS “Device Task Failed” event opens the Schedule Details page where you can acknowledge or reschedule the failed CMS task (refer to “Managing Scheduled Tasks” on page 69).

For more information about viewing events, refer to “Events Reports” on page 362. For a description of the CMS and WX system events, refer to “System Events” on page 441.

Loading Device Boot Images

After you load a WXOS boot image on the CMS server, you can distribute the image to selected devices in a community or device group. To upload a boot image to CMS, refer to “Uploading a Boot Image” on page 408.



NOTE: If the CMS server has multiple network interface cards, you must specify which IP address the WX devices should use to download WXOS images from the CMS server (refer to “Using CMS Servers with Multiple Network Interfaces” on page 35).

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

Loading a boot image involves two tasks:

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

When loading a boot image to multiple devices, you should schedule the reboot separately after verifying that the boot image was loaded successfully on each device. This lets you activate the new boot image on all devices simultaneously. To verify a task was successful, refer to “Managing Scheduled Tasks” on page 69.

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



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

Monitor downgraded devices carefully. If problems occur, restore or roll back the configuration to the one used with the older boot image, if possible (refer to “Rolling Back Device Configurations” on page 58 and “Restoring Device Configurations” on page 61).


To load a boot image on selected devices:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices where you want to load the boot image. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.
3. From the Task menu, select Image > Load and click Go.

Figure 14: Loading a Boot Image on Devices

The screenshot shows the 'CMS Juniper NETWORKS' interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (highlighted), 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar under 'Management' lists 'Devices', 'Configurations', 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. The 'Devices' section is expanded, showing a 'Community/Device Group' dropdown set to 'default-10.87.72.10' and a 'Show Exceptions Only' checkbox. The main content area is titled 'Devices > Load image'. It contains the instruction 'Load the following boot image to the selected device(s). [Show selected devices](#)'. Below this is a 'Boot image' dropdown menu. The 'Schedule' section has two radio buttons: 'Load now' and 'Delay loading until:'. The 'Delay loading until' option is selected, showing 'Time' and 'Date' fields. The 'Reboot' section has a checkbox 'Reboot device(s) after loading boot image'. The 'Downgrade' section has a checkbox 'Allow image downgrade'. A detailed warning text is provided below the 'Downgrade' checkbox. At the bottom of the form are 'Submit' and 'Cancel' buttons.

4. Specify the following information:

Boot image	<p>Select the WXOS boot image you want to load. The default naming convention of boot images is:</p> <p>srs < rdmbb > . < zip or bin ></p> <p>where:</p> <p>< r > is the major release number.</p> <p>< d > is the minor release number.</p> <p>< m > is the maintenance release number.</p> <p>< bb > is the build number.</p> <p>The file extension must be “zip” or “bin”.</p>
Schedule	<p>Select Load now or select Delay loading until and enter a future time and date (in local CMS time):</p> <ul style="list-style-type: none">■ Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM.■ Enter a date in MM/DD/YYYY format or click  and select the month and date.
Reboot	<p>Click the check box to reboot the device after the image is loaded. The loaded image is not activated until the device is rebooted.</p> <p>To schedule the reboot as a separate task, which is recommended when updating multiple devices, refer to “Rebooting Devices” on page 50.</p> <p>NOTE: When you reboot a device, all unsaved configuration data is lost.</p>
Downgrade	<p>Click the check box if the selected boot image is older than the current version.</p> <p>CAUTION: Downgrading to a previous boot image may cause unpredictable behavior and should be avoided whenever possible.</p>

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

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

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

Rolling Back Device Boot Images

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



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

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

Rolling back a boot image involves two tasks:

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


When rolling back the boot image on multiple devices, you should schedule the reboot separately after verifying that the rollback was successful on each device. This lets you activate the boot image on all devices simultaneously. To verify a task was successful, refer to “Managing Scheduled Tasks” on page 69.

To roll back the boot image on selected devices:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices where you want to roll back the boot image. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.
3. From the Task menu, select Image > Rollback and click Go.

Figure 15: Rolling Back the Boot Image

4. Specify the following information:

- | | |
|----------|--|
| Schedule | <p>Select Roll back now or select Delay roll back until and enter a future time and date (in local CMS time):</p> <ul style="list-style-type: none"> ■ Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. ■ Enter a date in MM/DD/YYYY format or click  and select the month and date. |
| Reboot | <p>Click the check box to reboot the device after the rolled back image is loaded. The loaded image is not activated until the device is rebooted.</p> <p>To schedule the reboot as a separate task, which is recommended when updating multiple devices, refer to “Rebooting Devices” on page 50.</p> <p>NOTE: When you reboot a device, all unsaved configuration data is lost.</p> |

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

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

Rebooting Devices

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



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


To reboot selected devices:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices that you want to reboot. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.
3. From the Task menu, select Reboot and click Go.

Figure 16: Rebooting a Device

The screenshot shows the CMS Juniper Networks interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (selected), 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar shows 'Management' with sub-items: 'Devices' (selected), 'Configurations', 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. Under 'Devices', there is a 'Community/Device Group' dropdown set to 'default-10.87.72.10' and a 'Show Exceptions Only' checkbox. The main content area is titled 'Devices > Reboot'. It contains a warning: 'Reboot selected device(s) using the current boot image. Any unsaved changes will be lost. To delete the disk-based compression dictionary on WXC devices, check **Clean Reboot**. [Show selected devices](#)'. Below this, the 'Action' section has a checkbox for 'Clean Reboot (Requires WXOS 5.4 or later)'. The 'Schedule' section has two radio buttons: 'Reboot now' (selected) and 'Delay reboot until:'. The 'Delay reboot until' section includes 'Time' (HH:MM) and 'Date' (calendar icon) fields, with AM/PM selection. At the bottom are 'Submit' and 'Cancel' buttons.

4. For WXC devices running WXOS 5.4 or later, select Clean Reboot if you want to clear the disk-based compression dictionary used for Network Sequence Caching.
5. Select Reboot now or select Delay reboot and enter a time and date:
 - Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM.

- Enter a date in MM/DD/YYYY format or click  and select the month and date
6. To review the devices you selected, click Show selected devices.
 7. Click Submit to submit this task, or click Cancel.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

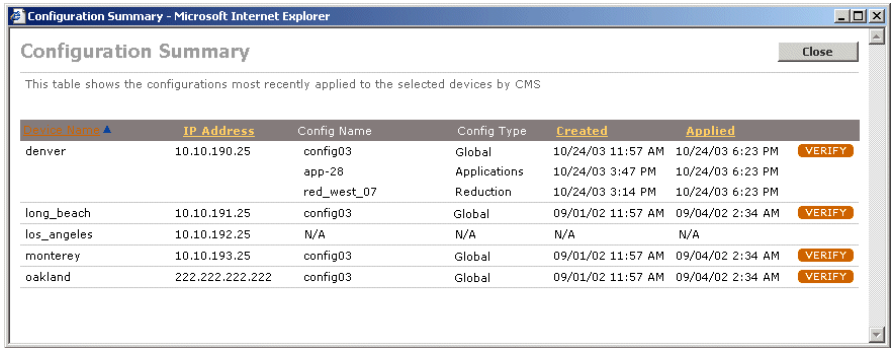
Viewing Device Configuration Summaries

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

To view configuration summaries:

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

Figure 17: Configuration Summary



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

The following information is displayed for each selected device:

- Name and type of the last global and partial configurations downloaded from CMS.

- Date and time the “load configuration” task was created and submitted to the scheduler.
- Date and time the configuration was applied to the device. The created and applied times are different if the load task was scheduled for a future time.

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

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

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

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

When you are done viewing the configuration, click Close.

To capture unsaved settings, you can extract the running configuration, as described in “Extracting Configurations” on page 85. Alternatively, you can incorporate the unsaved changes in an existing configuration (refer to “Defining Configuration Settings” on page 98).

Analyzing Device Configurations

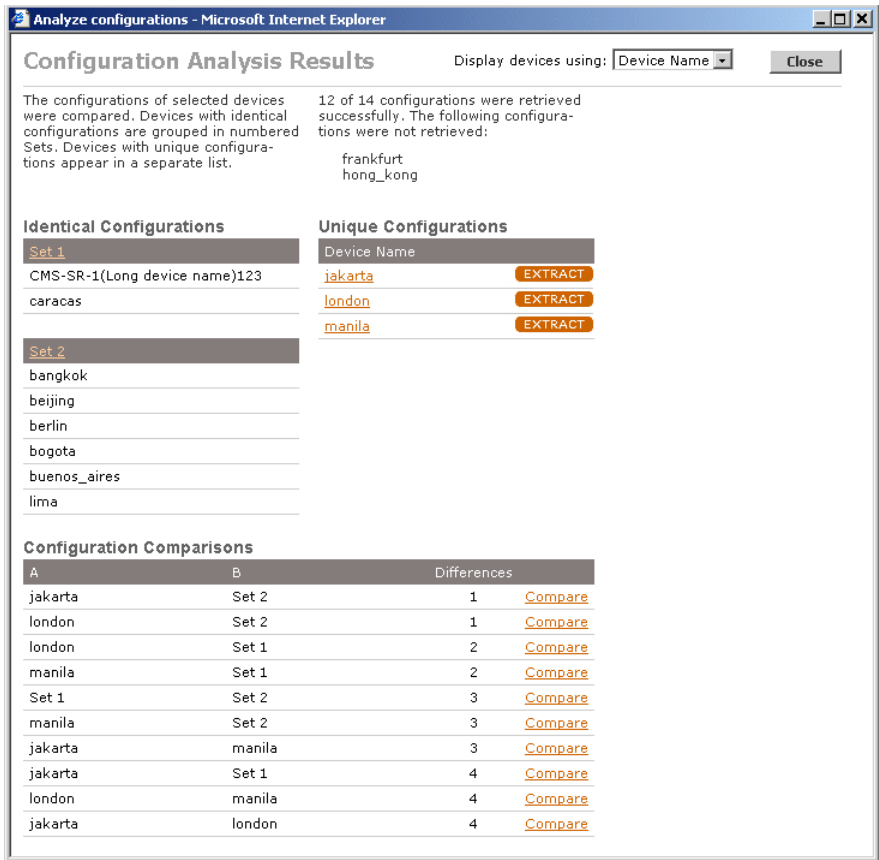
CMS lets you analyze the differences between the configurations of up to 30 devices in a community. This is particularly useful if the devices were installed and configured without CMS. Based on the analysis, you can eliminate unnecessary differences between devices, and extract global or partial configurations that you can load on other devices.

To analyze configurations:

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

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

Figure 18: Analyzing Configurations



The Configuration Analysis Results window includes the following:

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



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

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

When you are done viewing the configurations, click Close.

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

The extracted configuration is added to the Configurations page. You can then edit the configuration and load it on selected devices, as described in "Defining Configuration Settings" on page 98, and "Loading Device Configurations" on page 55.

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

Loading Device Configurations

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

Downloaded configurations are processed as follows:

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

By default, downloading a configuration replaces the corresponding settings on the device. If you elect to merge a downloaded configuration, list items such as applications, endpoints, and templates are added to the device configuration (refer to “Merging Configurations” on page 79). However, all individual items, such as the device IP address, always replace the corresponding setting on the device.

Any settings not defined in the downloaded configurations are left unchanged on each device, provided the settings were saved in the startup configuration file. Settings that can be specified only by CLI commands are retained on each device unless they are overridden by commands in the CLI section of a global or partial configuration.

You can preview the results for each device before submitting the task. For more information about global and partial configurations, refer to “Overview of Device Configurations” on page 75.

To load a configuration:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices where you want to load a configuration. To select all devices displayed on the page, click Select All. To select all devices in the community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.

- If you have previously loaded configurations from CMS, check each device for unsaved configuration settings (refer to “Viewing Device Configuration Summaries” on page 51). Unsaved settings are lost when a new configuration is loaded.



NOTE: Do not select a mix of WXOS 5.3/5.2 devices and WXOS 5.4 devices.

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

Figure 19: Loading a Configuration

Management

Devices > Load Configuration HELP

Load the following configuration to the selected device(s). [Show selected devices.](#)
Parameters in the selected configuration(s) will **replace** those on the devices. Verify that you have set all parameters correctly before loading configurations on the devices. For a detailed description of how this feature works, click on the **Help** button above.

Action

☒ Settings in selected configuration(s) **replace** those on the device.

☐ Settings in selected configuration(s) are **merged** with those on the device.

Global Configuration

☐ Do not load global configuration

☒ **factory-54** [History](#)

Partials

☒ No

☐ Yes

Schedule

☐ Load now

☒ Delay loading until:

Time HH:MM ☒ AM ☐ PM

Date

Recurrence No Recurrence

Reboot

☐ Reboot device(s) after loading configuration


Submit **Preview...** **Cancel**

- Specify the following information:

Global Configuration

To load a global configuration, select a global configuration from the menu. Click History to view the selected configuration and its history of changes. To create global configurations, refer to “Managing Configurations” on page 85.

To load only partial configurations, select Do not load global configuration.

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

6. To review the devices you selected, click Show selected devices.
7. Click Preview to view the resulting configuration for the first device. The blue text indicates the device-specific commands that can be set in the CMS Web console (CLI-only commands are not highlighted). Click Next to preview the configuration for each selected device.



NOTE: Always click Preview to verify there are no conflicts in the resulting device configurations. When you submit this page, CMS detects only the conflicts between the CMS configurations. Conflicts between the CMS configuration and the device configurations are detected only when the task is run, and then they appear as errors on the Schedule Details page (refer to “Managing Scheduled Tasks” on page 69).

To open the WXOS Web console, refer to “Accessing the WXOS Web Console from CMS” on page 68. To change a global or partial configuration, refer to “Defining Configuration Settings” on page 98.

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

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

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

Rolling Back Device Configurations

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



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


To roll back the configuration on selected devices:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices where you want to roll back a configuration. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.
3. From the Task menu, select Configuration > Rollback and click Go.

Figure 20: Rolling Back the Configuration

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

5. Specify the following information:

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

6. To review the devices you selected, click Show selected devices.
7. Click Preview to view the rollback configuration of the first device. Click Next and Back to page through all the configurations. When you're done, click Close.
8. Click Submit to submit this task, or click Cancel.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

Backing Up Device Configurations

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

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

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

The backup configuration files are archived as “Version 1.1”, “Version 1.2”, and so on. A new version is archived only if changes have occurred since the last backup. To restore an archived configuration file, refer to “Restoring Device Configurations” on page 61.


To back up device configurations:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices you want to back up. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.
3. From the Task menu, select Configuration > Backup and click Go.

Figure 21: Backing up Configuration Files

The screenshot shows the Juniper CMS web interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (highlighted), 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar shows the 'Management' menu with options like 'Configurations', 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. Under 'Devices', the 'Community/Device Group' is set to 'cms-private'. The main content area is titled 'Devices > Backup Startup File'. It contains a 'Backup' section with three radio button options: 'The saved startup.cfg file.', 'The running startup.cfg file.', and 'The running startup.cfg file after saving it first.' The 'Schedule' section has two radio button options: 'Backup now' and 'Delay backup until:'. The 'Delay backup until' section includes fields for 'Time' (HH:MM), 'Date', and 'Recurrence' (set to 'No Recurrence'). There are 'Submit' and 'Cancel' buttons at the bottom.

4. Specify the following information:

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

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

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

CMS reports on whether the task was submitted successfully. To view the status of the task and access the backed up configuration file, refer to “Managing Scheduled Tasks” on page 69.

Restoring Device Configurations


If you periodically back up the configuration file (*startup.cfg*) for each device, you can restore a backup configuration at any time. Note that each backup contains device-specific settings, so it can be restored only to its original device. To back up configuration files, refer to “Backing Up Device Configurations” on page 60.

To restore an archived configuration file:

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

Figure 22: Restoring a Configuration File

4. Specify the following information:

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

5. To verify the device you selected, click Show selected device.
6. Click Submit to submit this task, or click Cancel.

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

Retrieving Device Files

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

- **Diagnostic file.** Current configuration and the most recent system log and access control log.
- **System Log.** Critical, error, and information messages related to the operation of the device.
- **Access Control Log.** Log of each user who accessed the device in the last five days. Includes the workstation IP address for HTTPS and SSH access, and any configuration changes made by the user.
- **Monitor statistics.** All performance data from last month through the hour of the retrieval. The statistics are described in Appendix , “Understanding Exported Data Results.”
- **Flow statistics.** Top traffic flows collected on the device.

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

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

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

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



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

You can access the retrieved files from the Schedules page and download the files to the hard disk of your CMS Web console. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

To retrieve files from selected devices:

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

Figure 23: Retrieving Statistics/Diagnostic Files

Management

Devices

Configurations

Auto-Deployment

License Management

Scheduled Tasks

Schedule Log

Community/Device Group

default-10.87.240.2

☐ Show Exceptions Only

Submit

Devices > Retrieve Statistics/Diagnostic Files

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

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

Retrieve Files:

☐ Diagnostic file

☐ System Log

☐ Access Control Log

☐ Monitor statistics

☐ Flow statistics (Top talker data)

Schedule

☐ Retrieve now

☒ Delay retrieval until:

Time HH:MM ☐ AM ☐ PM


Date

Recurrence No Recurrence

The file downloaded for each device will be stored in the sub-directory identified by the device's ip address under the directory "C:\Program Files\Juniper Networks\CMS54\data\download". All files for each device will be bundled into one ZIP file.

Submit **Cancel**

4. Select the checkbox next to the files you want to retrieve. If you select the diagnostic file, the system and access control logs are included, so you do not need to select them separately.
5. Select Retrieve now or select Delay retrieval until and enter a time and date, and a recurring retrieval interval (optional):

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

CMS reports on whether the task was submitted successfully. To view the status of the task and access the retrieved files, refer to “Managing Scheduled Tasks” on page 69.

Applying a Registration Server Password

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

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

- Use the WXOS Web console to change the password on the registration server.
- Enter the new password in CMS (refer to “Importing and Managing Communities” on page 397).
- Apply the new password to the devices in each community, as described below.




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

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

1. On the Devices page, select a community from the Community/Device Group menu.
2. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu.
3. From the Task menu, select Apply password and click Go.

Figure 24: Applying a Registration Server Password

Show selected devices'. Under the 'Schedule' section, there are two radio buttons: 'Apply password now' and 'Delay until:'. The 'Delay until:' option is selected, with input fields for 'Time' (HH:MM) and 'Date' (MM/DD/YYYY), and radio buttons for 'AM' and 'PM'. At the bottom are 'Submit' and 'Cancel' buttons."/>

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

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

Putting Devices in Safe Mode

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

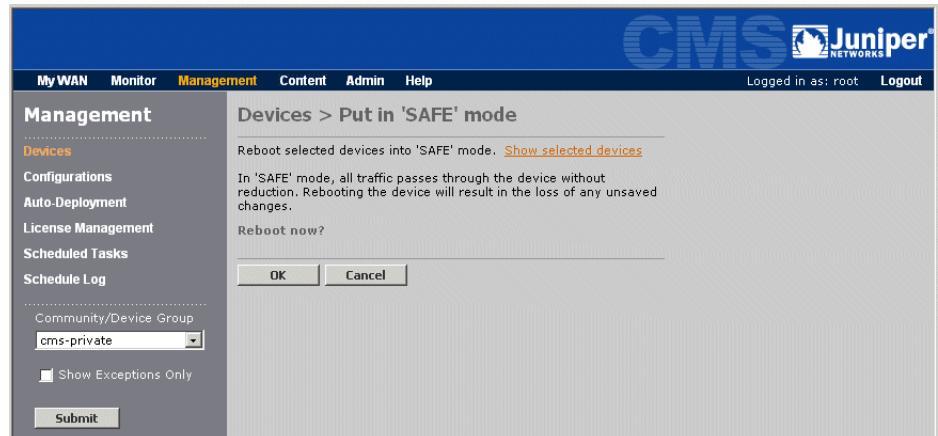


NOTE: Putting a device in safe mode reboots the device, so any unsaved configuration data is lost.

To put devices in safe mode:

1. On the Devices page, select a community or device group from the Community/Device Group menu.
2. Select the devices that you want to put in safe mode. To select all devices displayed on the page, click Select All. To select all devices in the selected community or device group, select Community/Device Group from the Apply To menu. Note that if you select Community/Device Group, the devices are determined dynamically when the task is run.
3. From the Task menu, select Put in 'SAFE' mode and click Go.

Figure 25: Putting Devices in Safe Mode



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

CMS reports on whether the task was submitted successfully. To view the status of the task, refer to “Managing Scheduled Tasks” on page 69.

6. Perform the appropriate troubleshooting and diagnostics. When you are done, reboot the devices to enable data compression (refer to “Rebooting Devices” on page 50).

Accessing the WXOS Web Console from CMS

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

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



NOTE: For WX devices that reside on the private side of a NAT device, the device link on the Devices page (and on the My WAN and Monitor pages) uses either the public or private IP address or domain name, depending on the Device Access URL setting in the user's account (refer to “Defining CMS User Accounts” on page 388).

Exporting Community and Device Information

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

The exported file contains the following information:

- Community or device group name
- Registration server IP address
- Date and time of the export
- The following information for each device:
 - Device name.
 - IP address.

- Model number.
- Serial number.
- Local interface MAC address.
- Remote interface MAC address.
- License speed.
- Software version.

To export device information:

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

Managing CMS Schedules

The following sections describe how to use the CMS scheduler:

- “Managing Scheduled Tasks” on page 69
- “Exporting a Schedule Log” on page 74

Managing Scheduled Tasks

The Schedules page lets you view all device tasks scheduled in the last 15 days, including tasks that are pending, in-process, successful, failed, or cancelled. To view tasks older than 15 days, refer to “Exporting a Schedule Log” on page 74. To view scheduled reports, refer to “Managing Scheduled Reports” on page 369.

Scheduled tasks can be viewed or changed by the user who scheduled the task, a CMS administrator, or the user group administrator. For tasks that retrieve files from a device or back up a configuration, you can open the retrieved files or download them to a local disk.

The following scheduling actions are available for tasks, depending on how the devices are specified:

- For tasks scheduled for selected devices, you can:
 - **Acknowledge.** Failed tasks can be acknowledged, which removes the failed task icon from the Devices page.
 - **Add Devices.** Pending and recurring tasks can be updated to include more devices.
 - **Cancel.** Pending tasks can be cancelled for all or selected devices.
 - **Reschedule.** Failed and pending tasks can be rescheduled.
- For tasks scheduled for an entire community or device group, you can cancel, reschedule, or acknowledge a failed task for all devices, but not for individual devices.

If your network is having problems, you can stop and restart the scheduler, as described in “Stopping and Starting the Scheduler” on page 430.

To manage the scheduled tasks:


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

Figure 26: Schedules Page

Management

Devices

Configurations

Auto-Deployment

License Management

Scheduled Tasks

Schedule Log

Community/Device Group

winnie's device group

Device

All devices

Scheduled Task Status

☒ Pending

☒ Successful

☒ Failed

☒ Cancelled



☒ In Progress

☐ Show Recurring Only


Submit

Scheduled Tasks

Refresh

Action	Reboot	Creation Time	User	Scheduled Time	Status	Devices
 Backup startup configuration		Aug 21, 2007 3:18 PM	root	Aug 29, 2007 3:20 PM	Failed	40
Retrieve files		Aug 21, 2007 10:53 PM	root	Aug 21, 2007 10:53 PM	Failed	40
Retrieve files		Aug 21, 2007 10:50 PM	root	Aug 21, 2007 10:50 PM	Success	3
Load image	<input checked="" type="checkbox"/>	Aug 21, 2007 6:37 PM	devadmin	Aug 21, 2007 6:37 PM	Success	6
Retrieve files		Aug 21, 2007 3:47 PM	root	Aug 21, 2007 3:47 PM	Failed	40
 Retrieve files		Aug 21, 2007 3:17 PM	root	Aug 21, 2007 3:20 PM	Cancelled	40

The Schedules page provides the following information for each task:

- **Action**—Task name. The  icon indicates a recurring task. Move the cursor over the icon to view the frequency and the next run time.
- **Reboot**—A check mark indicates that the task includes a reboot after the task is performed.
- **Creation Time**—Date and time the task was submitted to the scheduler.
- **User**—ID of the user who submitted the task.
- **Scheduled Time**—Date and time that the task is scheduled. For a recurring schedule that has run once, this is the scheduled time of the last run.
- **Status**—The status of the task. For a recurring schedule that has run once, this is the status of the last run.
- **Devices**—Number of devices for which the task is scheduled.

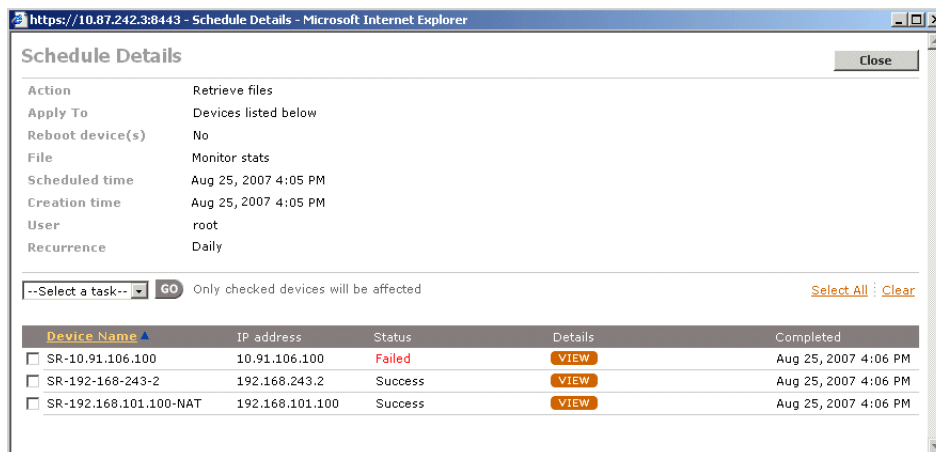


NOTE: A “Failed” status indicates the task has failed for at least one device.

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

To cancel a pending task, view the results of a task, or perform other scheduling functions, click the task name to open the Schedule Details page. For tasks scheduled for an entire community or device group, you can cancel, reschedule, or acknowledge a failed task for all devices, but not for individual devices.


Figure 27: Viewing Task Details



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

- The Apply To field indicates whether the task is scheduled for specific devices (“Devices listed below”) or a named community or device group. The devices listed for a named community or device group cannot be selected or changed (they are extracted from the database).
- The Details column contains one of the following:
 - Details about a failed task
 - A link to a retrieved file (for one-time retrievals). Click the link to open or save the file to a local disk.
 - A VIEW button for a recurring schedule. Click the button to view the details of each run for a specific device. A recurring backup or file retrieval includes a link to each retrieved file.

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

Add devices to pending or recurring tasks	<p>For tasks scheduled for specific devices, you can add more devices from the same community or device group:</p> <ol style="list-style-type: none"> 1. Select Add Devices on the task menu, and click Go. 2. Select the devices to be added and click Submit.
Reschedule failed or pending tasks	<ol style="list-style-type: none"> 1. For tasks scheduled for specific devices, click the check box next to the devices that you want to reschedule. For tasks scheduled for a community or device group, you must reschedule all devices. 2. Select Reschedule on the task menu, and click Go. 3. Reschedule the task and click Submit. <p>For each device, the status is changed to “Failed: Rescheduled” for failed tasks or “Cancelled: Rescheduled” for pending tasks, and a new “Pending” task is added to the Schedules page.</p> <p>If you reschedule a pending task for all its devices, the original task is changed to “Cancelled” on the Schedules page.</p>
Acknowledge failed tasks	<ol style="list-style-type: none"> 1. For tasks scheduled for specific devices, click the check box next to the devices for which you want to acknowledge the failed task. For tasks scheduled for a community or device group, you must acknowledge all devices. 2. Select Acknowledge ‘Failed’ Status on the task menu, and click Go. <p>The  icon is removed from the devices on the Devices page, and the status is changed to “Failed: Acknowledged” on the Schedule Details page.</p>
Cancel pending tasks	<ol style="list-style-type: none"> 1. For tasks scheduled for specific devices, click the check box next to the devices you want to cancel. For tasks scheduled for a community or device group, you must cancel all devices. 2. Select Cancel on the task menu, and click Go. <p>For each device, the status is changed to “Cancelled”. If the task is still pending for some devices, the task remains on the Schedules page as a “Pending” task.</p>

Exporting a Schedule Log

You can export a schedule log containing information about tasks submitted to the scheduler for a particular community or device group. You can save the file in CSV format on a local disk, and then import its contents into a spreadsheet program (such as Microsoft Excel).

The file contains the following information:

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

To export the schedule log:

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

Chapter 4

Managing Device Configurations

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

- “Overview of Device Configurations” on page 75.
- “Viewing Configurations” on page 83.
- “Managing Configurations” on page 85.
- “Defining Configuration Settings” on page 98.

Overview of Device Configurations

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

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

The Device Settings partial configuration can be used to configure settings for one device, such as the IP address. Alternatively, you can define device-specific settings through the device Web console (refer to “Accessing the WXOS Web Console from CMS” on page 68).

Configuration Settings for WXOS 5.x Devices

Table 2 lists the configuration settings that can be defined in each type of CMS partial configuration for WXOS 5.4 devices. Except for Device Settings, global configurations include all of the partial configuration settings. Most partial configurations correspond to parameter groups in the device Web console. Note that CMS provides separate procedures to deploy licenses and change the registration server password.

Table 2: Partial Configurations for WXOS 5.x Devices

Type	CMS Settings	Notes
Device Settings	<ul style="list-style-type: none"> ■ Addresses ■ Communities ■ Time zone ■ ARP ■ Compression subnets ■ Outbound QoS exclusions ■ Static local routes ■ Dynamic local routes (router polling) ■ Multi-Path (secondary IP address) ■ RADIUS source IP address 	
Basic Setup	<ul style="list-style-type: none"> ■ Interfaces ■ Time (NTP servers) ■ SNMP ■ Syslog server ■ Dynamic local routes (OSPF/RIP) ■ Router balancing ■ Registration server ■ NetFlow 	<ul style="list-style-type: none"> ■ To apply license keys from CMS, refer to “Automatic Deployment and License Management” on page 287. ■ To apply a new registration server password to multiple devices, refer to “Applying a Registration Server Password” on page 65.
AAA	<ul style="list-style-type: none"> ■ Authentication ■ Authorization ■ RADIUS ■ Local users ■ Operator access ■ Front panel access 	<ul style="list-style-type: none"> ■ Packet capture passwords must be defined on the device.
Applications	<ul style="list-style-type: none"> ■ Overview ■ Application definitions ■ Traffic classes ■ Monitoring 	

Table 2: Partial Configurations for WXOS 5.x Devices

Type	CMS Settings	Notes
Compression	<ul style="list-style-type: none"> ■ Endpoints ■ Network Sequencing Caching ■ Application filter ■ Remote routes ■ Load balancing ■ Default decompressors ■ Preferred decompressors ■ Tunnel mode 	
QoS	<ul style="list-style-type: none"> ■ Setup Wizard ■ Overview ■ Templates ■ Endpoints ■ ToS/DSCP ■ Start/stop ■ Inbound Qos 	
Acceleration	<ul style="list-style-type: none"> ■ Overview ■ Fast Connection Setup ■ TCP Acceleration ■ CIFS applications ■ HTTP applications ■ Exchange applications 	
Advanced Setup	<ul style="list-style-type: none"> ■ Topology ■ Source/destination filter ■ Prime time ■ Packet interception ■ WAN performance monitor ■ WX 100 Multi-Tunnel ■ CLI commands 	
Multi-Path	<ul style="list-style-type: none"> ■ Start/stop ■ Templates ■ Endpoints 	
IPSec	<ul style="list-style-type: none"> ■ Overview ■ Templates ■ Default policy 	<ul style="list-style-type: none"> ■ The IPSec Setup Wizard can be run only from the device Web console.
Events	<ul style="list-style-type: none"> ■ Event Definitions 	<ul style="list-style-type: none"> ■ WXOS 5.4 (or later) devices only

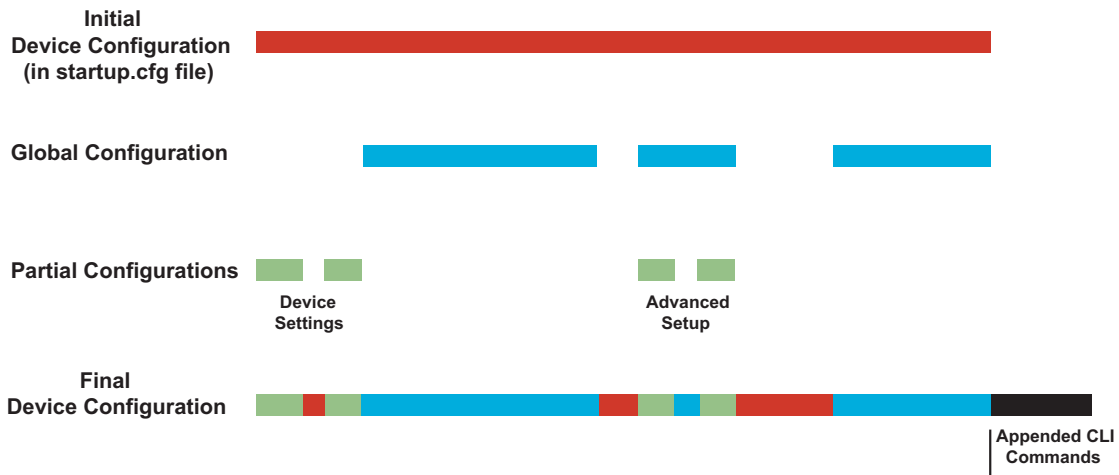
Downloading Global and Partial Configurations

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

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

Figure 28 shows how downloaded global and partial configurations are added to a device's configuration in the *startup.cfg* file (any unsaved device settings will be lost). The gaps indicate unspecified settings in the global and partial configurations. Any CLI commands specified in the CLI section of the downloaded configurations are appended to the *startup.cfg* file on the device.

Figure 28: Downloading Global and Partial Configurations on a Device



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

To configure a new device as a hub, simply download the two configurations. Table 3 shows the relevant topology CLI commands in each configuration.

Table 3: Combining Global and Partial Configurations

Global Configuration	Advanced Setup Configuration	Resulting Device Configuration
config reduction set topology-type spoke . . .	config reduction set topology-type hub config reduction set topology-size 1	config reduction set topology-type hub config reduction set topology-size 1



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

For descriptions of each CLI command, refer to the *WX/WXC Operator's Guide*. To download global configuration settings from CMS, refer to "Loading Device Configurations" on page 55.

Merging Configurations

By default, downloading a configuration replaces the corresponding settings on the device. If you elect to merge a downloaded configuration, list items such as applications, endpoints, and templates are added to the device configuration. However, all individual items, such as the device IP address, always replace the corresponding setting on the device.

For example, Table 4 shows a device with the default application definitions. If a CMS configuration with the WebApp and Mail applications is merged with the device configuration, the WebApp application is added to the device, and the Mail application overwrites the existing Mail definition. If the "replace" mode is used, all application definitions on the device are replaced with the two application definitions from CMS.

Table 4: Merging Configurations

Device Configuration	Downloaded Configuration	Resulting Device Configuration
Application definitions:	Application definitions:	Application definitions:
FTP	WebApp	WebApp
Telnet	Mail	FTP
Mail		Telnet
.		Mail
.		.
.		.
.		.



NOTE: Before downloading a configuration from CMS, always preview the device configuration to check for configuration conflicts or unexpected results (refer to “Loading Device Configurations” on page 55).

When merging configurations, note the following:

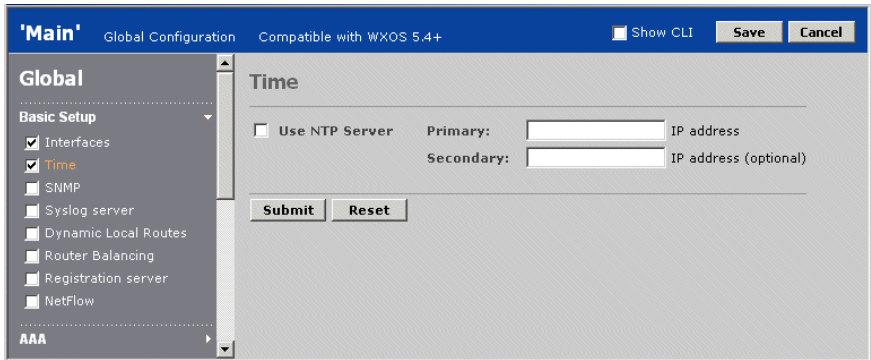
- Only the top-level lists are merged. For example, the list of application definitions are merged, but not the rules for a specific application. Similarly, the lists of traffic classes and the applications assigned to each class are merged, but not the traffic classes within a QoS template.
- Lists of applications enabled for compression (Application Filters) or TCP Acceleration can be merged. To disable applications for compression or TCP Acceleration, the device configuration settings must be replaced.
- Source/destination filters are replaced, not merged, if the mode of the configuration (off, include, or exclude) does not match the device.
- An error occurs if the merged configuration would exceed any maximum limits (such as 5 Syslog servers and 10 SNMP destinations).

Creating and Editing Configurations

Configurations can be created manually in CMS or by extracting a configuration from a device (recommended). New configurations can be accessed only by their creator (owner), a CMS administrator, or a user group administrator. To allow access by other users, a configuration must be published (refer to “Publishing Configurations” on page 94).

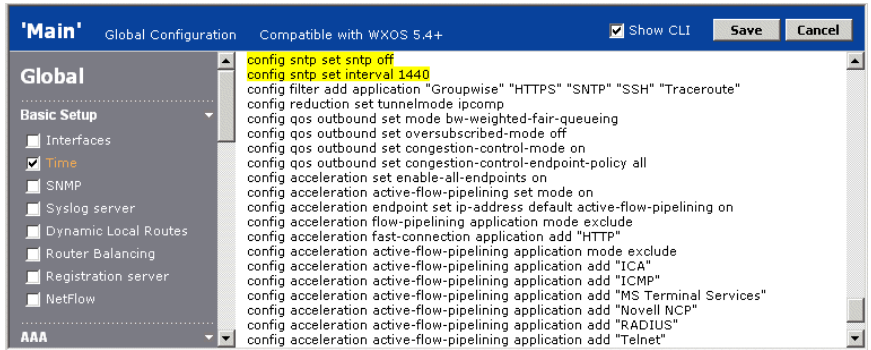
Configuration settings are selected by setting check boxes. Figure 29 shows a global configuration with the **Time** settings selected. Selecting a check box enables the default settings. You can then click the link next to the check box to change the defaults. If you clear the check box, the associated CLI commands are removed from the configuration.

Figure 29: Selecting Configuration Settings



Click **Show CLI** in the header to view the CLI commands in the current configuration. The commands associated with the currently selected link are highlighted. Click any link that has an enabled check box to highlight the associated commands (some default settings have no explicit commands). Clear the **Show CLI** check box to restore the configuration page.

Figure 30: Viewing CLI Commands



Consistency Checking

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

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

Tracking Configuration Versions

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

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

- **Version 2.2.** Compatible with WXOS 5.4 devices
- **Version 2.1.** Compatible with WXOS 5.2 and 5.3 devices

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

Tips for Managing Configurations

Review the following tips for managing global configuration settings:

- Analyze your existing device configurations to determine which configurations you want to extract and maintain on CMS. For more information, refer to “Analyzing Device Configurations” on page 52 and “Extracting Configurations” on page 85.
- Maintain a small number of unique global configurations, and define partial configurations to customize specific settings, such as topology settings, for selected devices and communities.

- Assign configuration names that reflect the contents of the configuration. Examples of configuration names include “hub-config,” “bandwidth-policy,” and “community-east.”
- Use the version history to track configuration changes and compare previous versions with the current version (refer to “Viewing Configuration History” on page 93).
- Publish your configurations to allow other users in your user groups to view or edit them. By default, configurations can be accessed only by their creator, a CMS administrator, or user group administrator (refer to “Publishing Configurations” on page 94).

Viewing Configurations

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

To view the configurations defined in CMS:

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

Figure 31: Configurations Page

Task menu

Task menu

Configurations

--Select a task-- GO Only checked configurations will be affected

Config Name	Config Type	Description	WXOS	References	Published	Modified
AAA-1	AAA		5.4+			05-15-2007
AAA-2	AAA		5.2/5.3			05-15-2007
Accel-1	Acceleration		5.4+	global-1		05-15-2007
Accel-2	Acceleration		5.2/5.3	global-2		05-15-2007
Adv-1	Advanced Setup		5.4+			05-15-2007
Adv-2	Advanced Setup		5.2/5.3			05-15-2007
App-1	Applications		5.4+			05-15-2007
App-2	Applications		5.2/5.3			05-15-2007
Basic-1	Basic Setup		5.4+			05-15-2007
Basic-2	Basic Setup		5.2/5.3			05-15-2007

From the Configurations page, you can:

- View the type, description, WXOS compatibility, and last-modified date for each global and partial configuration.

Some partial configurations (Acceleration, Compression, Event Definitions, Multi-Path, and QoS) must reference the application definitions in a global or Applications partial configuration. The References column indicates the referenced configuration. When a referenced configuration is changed, the referencing configurations are adjusted as needed.

- Move the cursor over the published icon to view the user groups that can access the configuration. If the icon is not shown, only the configuration owner (the creator), CMS administrator, or user group administrator can access the configuration.
- Click the column headers to change the sort.
- Execute a configuration task, such as generating new configurations, as described in “Managing Configurations” in the next section.
- Click the configuration name to change the settings, as described in “Defining Configuration Settings” on page 98.

Managing Configurations

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

- “Extracting Configurations” in the next section
- “Duplicating Configurations” on page 87
- “Creating New Configurations with Factory Defaults” on page 88
- “Comparing Configurations” on page 91.
- “Displaying Configurations” on page 92
- “Viewing Configuration History” on page 93
- “Publishing Configurations” on page 94
- “Importing Configurations” on page 95
- “Exporting Configurations” on page 96
- “Changing Referenced Configurations” on page 97
- “Upgrading Configurations” on page 98
- “Deleting Configurations” on page 98

Extracting Configurations

You can define new global configurations by extracting the saved configuration (the *startup.cfg* file) from a selected device. To determine which device configuration to extract, you can analyze the configurations of your current devices, as described in “Analyzing Device Configurations” on page 52.

You can also define new partial configurations by extracting the related configuration settings, such as application definitions, from a device or a CMS global configuration. After you extract and modify configurations, you can load them on selected devices, as described in “Loading Device Configurations” on page 55.

If you extract an Acceleration, Compression, Event Definitions, Multi-Path, or QoS partial configuration from a device, the applications and traffic classes will be undefined until you reference another configuration (refer to “Changing Referenced Configurations” on page 97).



NOTE: In a configuration extracted from a device, all settings are enabled (all check boxes are selected), but in a configuration extracted from a CMS global configuration, only the settings enabled in the global configuration are enabled in the extracted configuration.

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

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

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

Figure 32: Extracting Configurations

The screenshot shows the CMS Juniper Networks web console. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (selected), 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar under 'Management' lists 'Devices', 'Configurations' (selected), 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. The main content area is titled 'Configurations > Extract'. It contains a description: 'This page allows you to extract and save a configuration. You can extract a global configuration from a WX device or you can extract a partial configuration from a global configuration.' Below this are three radio button options: 'Extract global configuration from WX device' (selected), 'Extract partial configuration from WX device', and 'Extract partial configuration from global configuration'. There are two dropdown menus: 'Community/Device Group' set to 'All devices' and 'Device' set to '10.87.231.2'. Below these is a section 'Enter a name and description for the extracted configuration' with two text input fields for 'Configuration name' and 'Description'. At the bottom are 'Submit' and 'Cancel' buttons.

2. Do one of the following:
 - a. To extract a global configuration from a device, select a community from the Community menu, and select a device name from the Device menu.
 - b. To extract a partial configuration from a device, click Extract partial configuration from WX device, select a community from the Community menu, select a device name from the Device menu, and select the configuration type from the Partial Config Type menu.
 - c. To extract a partial configuration from a global configuration, click Extract partial configuration from global configuration, select a global configuration from the Global Configuration menu, and select the configuration type from the Partial Config Type menu.
3. Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_).
4. Enter a description of the configuration. The text “(Extracted from <source>)” is appended to the description, where <source> is the name of the device or the global configuration.
5. Click Submit to add the new configuration to the Configurations page. Extracting a global configuration from a device may take some time.
6. Click the configuration name to change its settings, as described in “Defining Configuration Settings” on page 98.

Duplicating Configurations

You can define new configurations by copying and modifying an existing global or partial configuration. After you copy and modify configurations, you can load them on selected devices, as described in “Loading Device Configurations” on page 55.

To copy an existing global or partial configuration:

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

Figure 33: Duplicating a Configuration

The screenshot shows the CMS web interface. The top navigation bar includes 'Dashboard', 'Monitor', 'Management' (highlighted), 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar under 'Management' lists 'Devices', 'Configurations' (highlighted), 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. Below these is a 'Show Configurations' dropdown set to 'All Configurations' and a 'Submit' button. The main content area is titled 'Configuration > Duplicate'. It contains the text 'Duplicate the device configuration extr-244.108'. Below this are two input fields: 'Duplicate configuration name' and 'Description'. At the bottom of the form are 'Submit' and 'Cancel' buttons.

3. Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_).
4. Enter a description of the configuration. The text “(Duplicated from < source >)” is appended to the description, where < source > is the name of the global or partial configuration.
5. Click Submit to add the new configuration to the Configurations page.
6. Click the configuration name to change its settings, as described in “Defining Configuration Settings” on page 98.

Creating New Configurations with Factory Defaults

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



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

To ensure that the password and registration server are correct, create new global configurations by extracting them from a working device (refer to “Extracting Configurations” on page 85).

After you create new configurations, you can load them on selected devices, as described in “Loading Device Configurations” on page 55.

To create a new configuration with factory defaults:

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

Figure 34: Creating a New Configuration with Factory Defaults

The screenshot shows the Juniper CMS interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (highlighted), 'Content', 'Admin', and 'Help'. The right side of the bar shows 'Logged in as: root' and a 'Logout' link. The left sidebar under 'Management' lists 'Devices', 'Configurations' (highlighted), 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. Below this is a 'Show Configurations' section with a dropdown menu set to 'All Configurations' and a 'Submit' button. The main content area is titled 'Configuration > New' and contains a 'Create new configuration' section. This section includes a text area for 'Configuration name', a text area for 'Description', a 'Compatibility' section with radio buttons for '5.4+' (selected) and '5.2/5.3', and a 'Config Type' section with radio buttons for 'Global' (selected) and 'Partial'. Next to the 'Partial' radio button is a dropdown menu currently showing 'AAA'. At the bottom of the form are 'Submit' and 'Cancel' buttons.

Specify the following information:

- | | |
|--------------------|---|
| Configuration name | Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_). |
| Description | Enter a configuration description (up to 100 characters). |
| Compatibility | Select 5.4 + or 5.2/5.3 to indicate the version of WXOS on the devices where the new configuration will be loaded. |
| Config Type | Select the new configuration type: <ul style="list-style-type: none">■ Global. Contains all settings that can be defined in CMS (except device-specific settings).■ Partial. Contains one group of settings.<p>For Compression, QoS, Acceleration, or Multi-Path partial configurations, you must also select a global configuration or an Applications partial configuration that contains the application definitions. For a Multi-Path partial configuration, the selected configuration must also specify QoS traffic classes.</p> |

2. Click Submit to add the new configuration to the Configurations page.
3. Click the configuration name to change its settings, as described in “Defining Configuration Settings” on page 98. For a new global or AAA partial configuration, you must change the default password for the admin account (refer to “Defining Local Users” on page 137).

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

- “Configuring Application Definitions” on page 145
- “Monitoring Applications” on page 151
- “Compressing Applications” on page 156
- “Configuring Tunnel Mode Settings” on page 167
- “Enabling TCP Acceleration by Application” on page 220
- “Enabling Fast Connection Setup by Application” on page 221



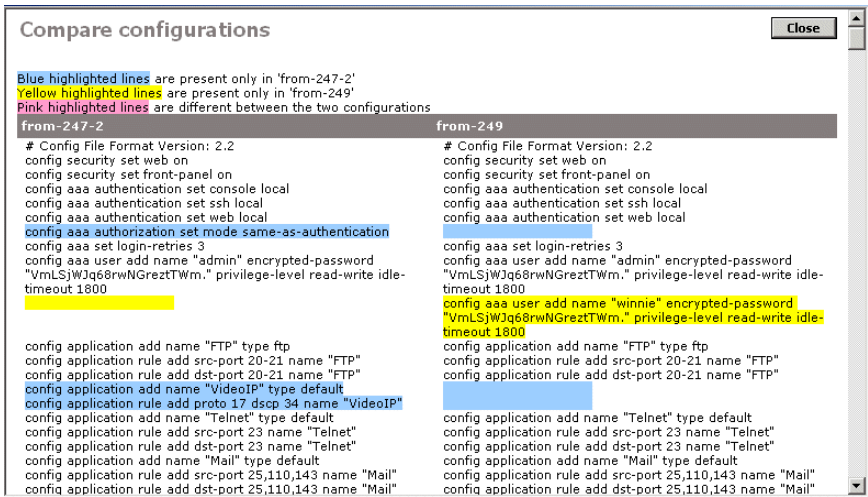
NOTE: If you do not specify the Topology settings, they default to a mesh topology with the lowest range of devices (refer to “Configuring Topology Settings” on page 229).

Comparing Configurations

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

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

Figure 35: Comparing Configurations



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

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

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

For descriptions of each CLI command, refer to the *WX/WXC Operator's Guide*.

Displaying Configurations

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

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

Figure 36: Displaying a Configuration



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

For descriptions of each CLI command, refer to the *WX/WXC Operator's Guide*.

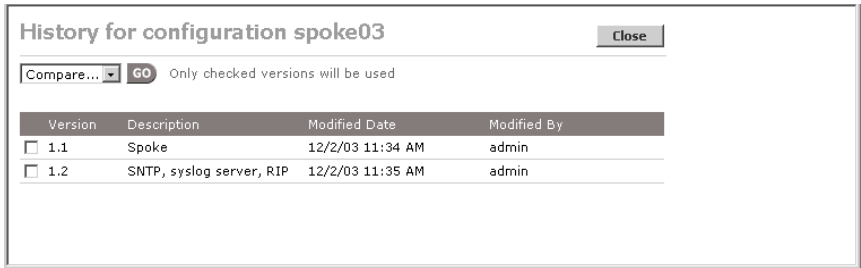
Viewing Configuration History

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

To view a configuration's history:

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

Figure 37: Viewing Configuration History



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

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

Publishing Configurations

New configurations can be accessed only by their creator (owner), a CMS administrator, or user group administrator. To allow access by other users, the owner (or administrator) can publish configurations to one or more user groups. The owner or user group administrator can publish only to the user groups that they belong to (a CMS administrator can publish to all user groups).



NOTE: User group administrators cannot view unpublished configurations owned by users who are authenticated remotely.

To publish a configuration:

1. On the Configurations page, select the checkbox next to the configurations that you want to publish.
2. From the Task menu, select Publish on the task menu, and click Go.

Figure 38: Publishing Configurations

3. To publish the selected configurations, select Publish to selected user groups below, and select the appropriate user groups.
4. To “unpublish” the selected configurations, clear the check box for one or more user groups or select Do not publish (affects all user groups).

- Click Submit to enter the changes, or click Cancel to discard them.



NOTE: Users who belong to only one of the published user groups can unpublish the configuration for the other user groups. Also, if the owner of a configuration is removed from a user group, the configuration remains published until the owner removes it.

Importing Configurations

Partial or global configurations can be imported to CMS from a text file. The file must reside on an accessible disk or an FTP server.

To import a configuration:

- On the Configurations page, select Import from the Task menu and click Go.

Figure 39: Importing Configurations

- Click Browse to import a local file, or click FTP Server and specify the server address, path name, user name, and password.
- If the file defines a partial configuration, select Partial, and select the type of partial configuration.
- Enter a name that reflects the contents of the configuration (up to 30 characters). Use only letters, numbers, hyphens (-), and underscores (_).

5. Enter a description of the configuration. The text “(Imported)” is used if the description is omitted.
6. Click Submit to enter the changes, or click Cancel to discard them.

An error occurs if the configuration name is not unique or the file contains invalid configuration commands.

Exporting Configurations

Partial or global configurations can be exported from CMS to a text file. If you export a partial configuration, you should indicate the type of partial in the file name. To import a partial configuration, you must specify the partial type (refer to “Importing Configurations” on page 95).

To export a configuration:

1. On the Configurations page, select the checkbox next to the configuration that you want to export.
2. From the Task menu, select Export on the task menu, and click Go.
3. Click Open to view the configuration, or click Save to save the configuration to a text file.

Changing Referenced Configurations

The partial configurations for Acceleration, Compression, Event Definitions, Multi-Path, and QoS must reference the application definitions (and traffic classes for QoS and Multi-Path) in a global or Applications partial configuration. If necessary, you can change the referenced configuration at any time.

To change a referenced configuration:

1. On the Configurations page, select the checkbox next to the partial configuration that you want to change.
2. From the Task menu, select Set Reference Config, and click Go.

Figure 40: Changing a Referenced Configuration

The screenshot shows the Juniper CMS interface for setting a reference configuration. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (highlighted), 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar shows the 'Management' menu with options like 'Devices', 'Configurations' (highlighted), 'Auto-Deployment', 'License Management', 'Scheduled Tasks', and 'Schedule Log'. The main content area is titled 'Configurations > Set Reference Configuration'. It contains a text box with the following configurations listed: 'Main', 'config-2.2', 'extr-246', 'factory-54', 'from-247-2', 'from-249', 'import-246', and 'upgrade-244'. Below the list are 'Submit' and 'Cancel' buttons. A 'Show Configurations' dropdown menu is set to 'All Configurations'.

3. Select a new partial or global configuration to be referenced.
4. Click Submit to enter the changes, or click Cancel to discard them.

The referencing configuration is updated to reflect the applications and traffic classes in the new referenced configuration. For example, application definitions are removed from the referencing configuration if they are not defined in the new referenced configuration.

Upgrading Configurations

A CMS 5.2/5.3 configuration can be loaded on a 5.4 device if it is first upgraded to the 5.4 format. To upgrade a configuration:

1. On the Configurations page, select the check box next to the 5.2/5.3 global or partial configuration you want to upgrade.
2. From the Task menu, select Upgrade and click Go.
3. Enter a name and description for the upgraded configuration, and click Submit.

Deleting Configurations

To delete configurations on CMS:

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



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

Defining Configuration Settings

After you generate a new configuration, you can define or change its settings. Remember that if you create a configuration as described in “Creating New Configurations with Factory Defaults” on page 88, all parameters have the same default settings as a new WX 50 device before Quick Setup is run.

All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *WX/WXC Operator's Guide*. To load a configuration on selected devices, refer to “Loading Device Configurations” on page 55.



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

To define configuration parameters:

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

The Configuration window opens for the selected global or partial configuration. Figure 41 shows the Basic Setup section of a global configuration.

Figure 41: Editing a Global Configuration

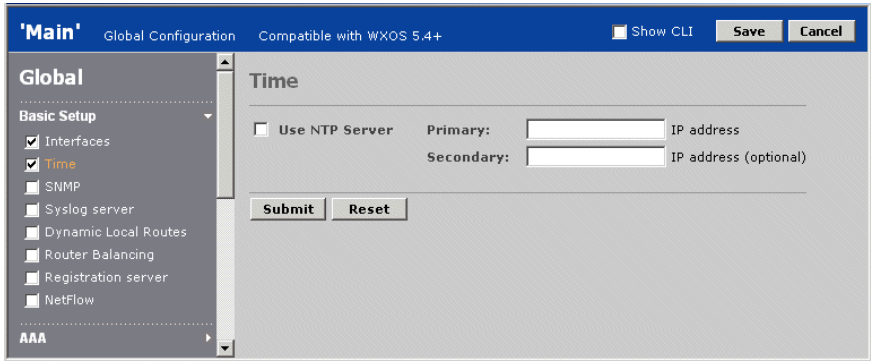
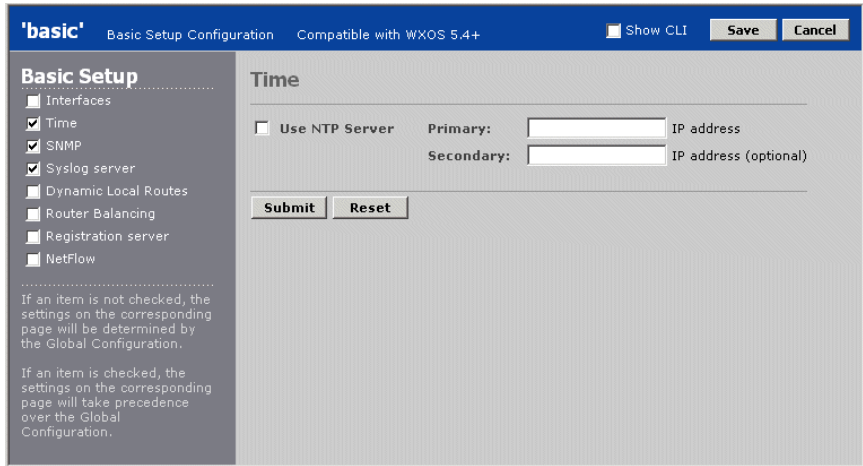


Figure 42 shows a Basic Setup partial configuration.

Figure 42: Editing a Partial Configuration



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

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

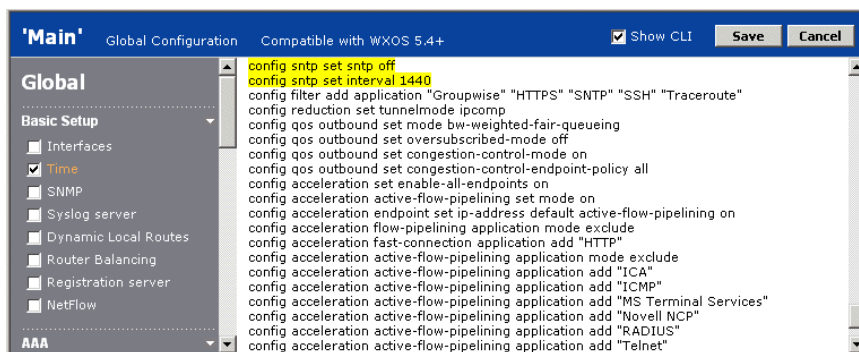


NOTE: Clearing a check box deletes the associated settings. When you load a configuration, only the checked settings affect the device. If a check box is not selected, the associated default settings are NOT applied.

Refer to the sections listed in Table 5 on page 101 for instructions on configuring each parameter.

- Click **Show CLI** in the header to view the CLI commands in the current configuration. The commands associated with the currently selected link are highlighted. Click any link that has an enabled check box to highlight the associated commands (some default settings have no explicit commands). Clear the **Show CLI** check box to restore the configuration page.

Figure 43: Viewing CLI Commands



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

If there are no errors, CMS creates an updated configuration with a new version number. The version number and change description can be viewed in the configuration history (refer to “Viewing Configuration History” on page 93).

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

Table 5 lists the sections that describe each group of configuration parameters for WXOS 5.4 devices. The Device Settings must be defined in a partial configuration. Each of the other parameter groups can be defined in a global or partial configuration.

Table 5: Directory of Configuration Parameters

Parameter Group	Sections
Device Settings	<ul style="list-style-type: none"> “Configuring Device Addresses” on page 103 “Defining Communities” on page 105 “Configuring Time Zone Settings” on page 107 “Configuring the ARP Table” on page 108 “Advertising Compression Subnets” on page 109 “Defining Outbound QoS Exclusions” on page 111 “Adding Static Routes” on page 112 “Configuring Router Polling” on page 114 “Configuring Multi-Path Addresses” on page 116 “Configuring the RADIUS Source Address” on page 118
Basic Setup	<ul style="list-style-type: none"> “Configuring the Interface Settings” on page 119 “Configuring NTP” on page 122 “Enabling SNMP” on page 123 “Defining Syslog Servers” on page 124 “Configuring Dynamic Local Routes” on page 126 “Enabling Route-Based Router Balancing” on page 128 “Designating a Registration Server” on page 130 “Generating NetFlow Records” on page 131
AAA	<ul style="list-style-type: none"> “Selecting Authentication Methods” on page 133 “Enabling Authorization Checking” on page 134 “Defining RADIUS Servers” on page 135 “Defining Local Users” on page 137 “Securing Operator Access” on page 138 “Securing Front Panel Access” on page 140
Applications	<ul style="list-style-type: none"> “Configuring Application Settings” on page 141 “Viewing the Application Overview” on page 144 “Configuring Application Definitions” on page 145 “Assigning Applications to Traffic Classes” on page 149 “Monitoring Applications” on page 151

Table 5: Directory of Configuration Parameters

Parameter Group	Sections
Compression	<ul style="list-style-type: none"> “Configuring Endpoints for Compression” on page 152 “Configuring Network Sequence Caching” on page 154 “Compressing Applications” on page 156 “Configuring Remote Routes” on page 158 “Configuring Tunnel Load Balancing Policies” on page 160 “Configuring Default Decompressors” on page 162 “Defining Preferred Decompressors” on page 165 “Configuring Tunnel Mode Settings” on page 167
QoS	<ul style="list-style-type: none"> “Understanding Outbound QoS” on page 170 “Using the Outbound QoS Setup Wizard” on page 182 “Defining Outbound QoS Settings by Endpoint” on page 191 “Defining Outbound QoS Templates” on page 193 “Defining Outbound QoS Endpoints” on page 195 “Changing Outbound ToS/DSCP Values” on page 199 “Starting and Stopping Outbound QoS” on page 202 “Configuring Inbound QoS Policies” on page 203
Acceleration	<ul style="list-style-type: none"> “Configuring Traffic Acceleration” on page 206 “Enabling Acceleration by Endpoint” on page 214 “Enabling Acceleration by Application” on page 219
Advanced Setup	<ul style="list-style-type: none"> “Configuring Topology Settings” on page 229 “Configuring Source/Destination Filters” on page 238 “Defining the Prime Time” on page 241 “Configuring Packet Interception” on page 242 “Configuring WAN Performance Monitoring” on page 258 “Adding CLI Commands to Configurations” on page 262
Multi-Path	<ul style="list-style-type: none"> “Enabling Policy-Based Multi-Path” on page 265 “Defining Multi-Path Templates” on page 266 “Defining Multi-Path Endpoints” on page 268 “Configuring Routers to Support Multi-Path” on page 270
IPSec	<ul style="list-style-type: none"> “Defining IPSec Settings by Endpoint” on page 275 “Defining IPSec Templates” on page 277 “Defining the Default IPSec Policy” on page 279
Events	<ul style="list-style-type: none"> “Configuring Events” on page 280

Configuring Device Settings

The Device Settings partial configuration lets you define device-specific configuration settings for a single WXOS 5.x device. Alternatively, you can define these settings in the WXOS Web console (refer to “Accessing the WXOS Web Console from CMS” on page 68).

If you use automatic deployment, a Device Settings partial configuration is generated automatically for each auto-deployed device (refer to “Automatic Deployment and License Management” on page 287).

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

- “Configuring Device Addresses” on page 103
- “Defining Communities” on page 105
- “Configuring Time Zone Settings” on page 107
- “Configuring the ARP Table” on page 108
- “Advertising Compression Subnets” on page 109
- “Defining Outbound QoS Exclusions” on page 111
- “Adding Static Routes” on page 112
- “Configuring Router Polling” on page 114
- “Configuring Multi-Path Addresses” on page 116
- “Configuring the RADIUS Source Address” on page 118

Configuring Device Addresses

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

To specify the network address and contact information:

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

Figure 44: Configuring Network Address and Contact Information

'Device 1' Device Settings Configuration Compatible with WXOS 5.4+ ☐ Show CLI **Save** **Cancel**

Device Settings

- ☒ Addresses
- ☐ Communities
- ☒ Time Zone
- ☐ ARP
- ☐ Compression Subnets
- ☐ Outbound QoS Exclusions
- ☐ Static Local Routes
- ☐ Dynamic Local Routes
- ☐ Multi-Path
- ☐ RADIUS

If an item is not checked, the settings on the corresponding page will be determined by the Global Configuration.

If an item is checked, the settings on the corresponding page will take precedence over the Global Configuration.

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*

Subnet Mask*

Default Gateway*

Device Name

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. To change the address of a registration server, you must first transfer the registration server to another device (refer to the <i>WX/WXC Operator's Guide</i>). |
| Subnet mask | Specify the network portion of the IP address. For example, "255.255.255.0" indicates that the first 24 bits of the IP address are used for the network portion of the address. |
| Default gateway | Enter the IP address of the default router (must be on the same subnet as the device). |
| Device name | 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.

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

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

Defining Communities

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

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

To configure the communities for a registration server:

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



NOTE: The Addresses section is selected automatically because the device address is needed to form the full name of the default community (“default- < IP address >”). Verify that the IP address, subnet mask, and gateway are defined in the Addresses section (refer to “Configuring Device Addresses” on page 103).

Figure 45: Defining Communities for a Registration Server

'Device 1' Device Settings Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Device Settings

- ☒ Addresses
- ☒ Communities
- ☒ Time Zone
- ☐ ARP
- ☐ Compression Subnets
- ☐ Outbound QoS Exclusions
- ☐ Static Local Routes
- ☐ Dynamic Local Routes
- ☐ Multi-Path
- ☐ RADIUS

If an item is not checked, the settings on the corresponding page will be determined by the Global Configuration.

If an item is checked, the settings on the corresponding page will take precedence over the Global Configuration.

Communities

This page allows you to manage communities. To create a new community, enter the new community name in the text field and click **Add**. To delete a community, click the corresponding **Delete** button. To view a list of devices belonging to a community, click the community name.

Community Name	Devices	
<u>default</u>	0	
<u>east</u>	0	DELETE

On the Communities page you can:

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

- e. Click Submit to enter the changes, or click Cancel to discard them, and then click Done to return to the Communities page.



NOTE: As new communities are added to a registration server, they must be imported into CMS (refer to “Importing and Managing Communities” on page 397). CMS queries the registration server(s) each day and automatically incorporates any changes to the imported communities.

Configuring Time Zone Settings

The Device Settings partial configuration lets you specify a device’s time zone and whether the device uses Daylight Savings Time. To specify a Network Time Protocol (NTP) server, refer to “Configuring NTP” on page 122.

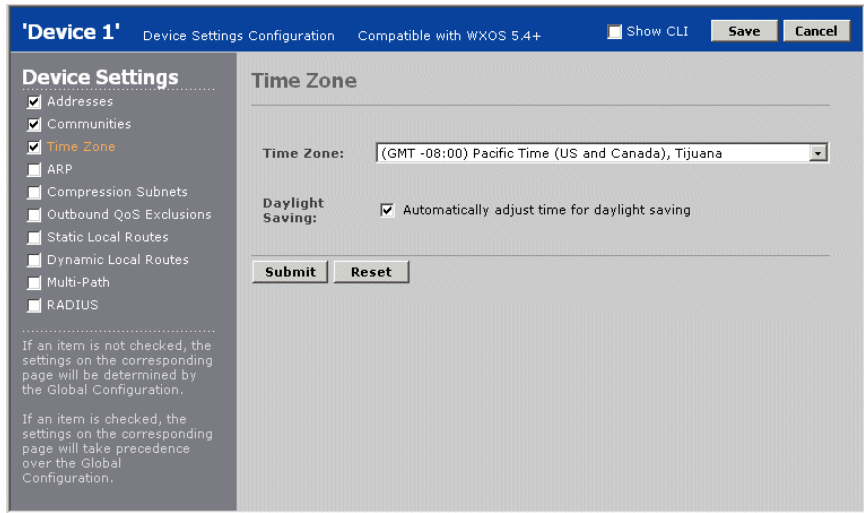


NOTE: When you view reports in the device’s time, the reported device times will be correct only if the time zone is set correctly.

To configure the time zone settings:

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

Figure 46: Configuring the Time Settings for a Device



2. Select the time zone of the device.
3. Select **Automatically adjust time for daylight savings**, if applicable.

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

Configuring the ARP Table

The Address Resolution Protocol (ARP) is used to:

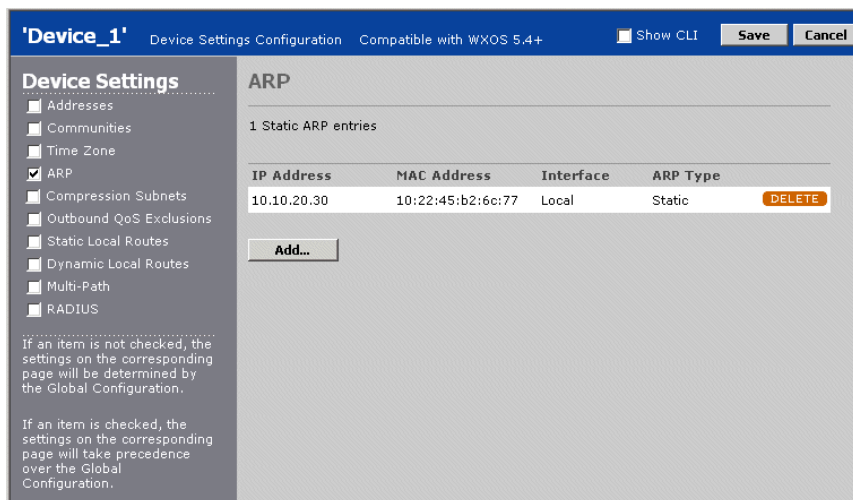
- Determine whether the gateway for a route is on the Local or Remote interface
- Discover the hardware (MAC) addresses of devices that are directly addressable on the Local and Remote interfaces

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

To add static entries to the ARP table:

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

Figure 47: Viewing the ARP Table



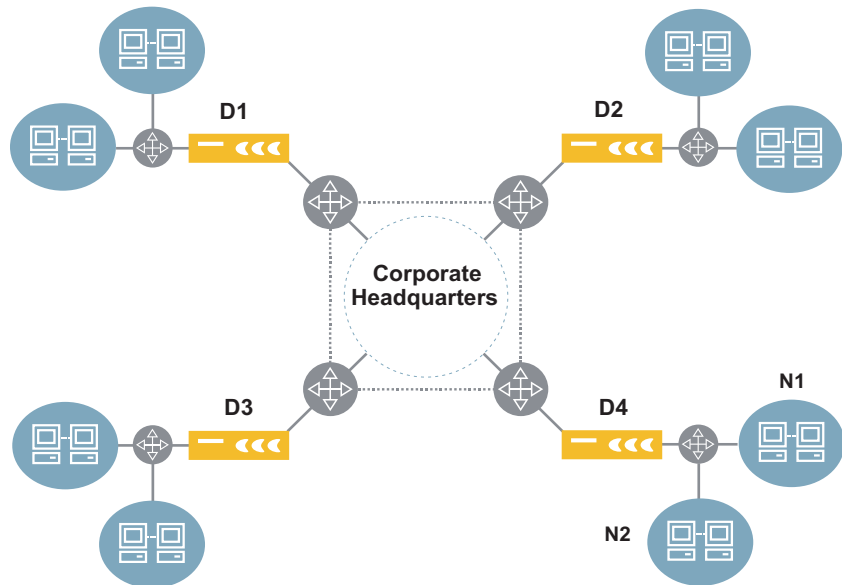
- To add one or more static ARP entries, click Add, enter the IP address and its associated MAC address, and select the Local or Remote interface. You can add up to five entries at one time.
- Click Submit to enter the new entries, or click Cancel to discard them.
- To delete an ARP entry, click DELETE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

Advertising Compression Subnets

Compression subnets are the subnets on the LAN side of the device that you can selectively advertise to the other devices in the community. The other devices can then compress and accelerate traffic sent to the advertised subnets. Initially, the only compression subnet is the subnet where the device is installed. To enable dynamic discovery of LAN-side subnets, refer to “Configuring Dynamic Local Routes” on page 126.

The set of subnets advertised by each device is called a “netmap.” By default, only the subnets you select are advertised. You can enable the advertisement of all subnets or just selected subnets. In Figure 48, each device has two subnets on its LAN side.

Figure 48: Selecting Specific Subnets for Data Compression



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

You can also control compression by application, as described in “Compressing Applications” on page 156 and by source/destination address, as described in “Configuring Source/Destination Filters” on page 238.



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

To advertise compression subnets:

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

The screenshot shows the 'Device 1' configuration window. The left-hand navigation pane is titled 'Device Settings' and lists various configuration options: Addresses, Communities, Time Zone, ARP, Compression Subnets (checked), Outbound QoS Exclusions, Static Local Routes, Dynamic Local Routes, Multi-Path, and RADIUS. Below this list, there is a note: 'If an item is not checked, the settings on the corresponding page will be determined by the Global Configuration. If an item is checked, the settings on the corresponding page will take precedence over the Global Configuration.'

The main content area is titled 'Compression Subnets'. It contains the following text: 'Subnets discovered by a Juniper WX device can be advertised to other Juniper WX devices so that data destined for them will be targeted for compression. (Note: Actual compression will also depend on filter settings.)' Below this, it says 'Select a method for advertising subnets for compression.' There are three radio button options: 'Advertise ALL discovered subnets', 'Advertise ONLY subnets listed below' (which is selected), and 'Advertise all discovered subnets EXCEPT those listed below'. Below these options, there is a text box with instructions: 'To add a subnet to the list, enter the IP address and subnet mask and click **Add**. Click **Delete** to remove a subnet from the list. When you are done, click **Submit**.' Below this text, there are two input fields labeled 'IP Address' and 'Subnet Mask', followed by an 'Add' button. At the bottom of the section, there is a 'Submit' button.

Figure 49: Configuring Compression Subnets

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

- **Advertise all discovered subnets EXCEPT those listed below.** All discovered subnets are advertised, except the ones you specify.



NOTE: Be careful to advertise only the LAN-side subnets that the device can access. Do not use the ALL option if the device is installed off-path (refer to “Configuring Packet Interception” on page 242) or if the WAN compression subnet option is enabled manually, such as in some VLAN environments. In these cases, all discovered LAN- and WAN-side subnets are eligible for advertisement.

3. Click Submit to enter the changes.

Defining Outbound QoS Exclusions

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



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

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

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

Figure 50: Excluding Subnets or Hosts from Bandwidth Management

The screenshot shows the 'Device_1' configuration page with the 'Outbound QoS Exclusions' section active. The left sidebar lists various settings, with 'Outbound QoS Exclusions' checked. The main area contains a checkbox for 'DO NOT impose Outbound QoS on traffic between the following network pairs.' which is checked. Below this is a text box with instructions: 'Enter IP address or address/subnet mask. Examples: 123.123.123.123 or 123.123.123.0/255.255.255.0. Enter asterisk (*) to indicate that the endpoint can be ANY address. Click Add to add a network pair to the list. Click Delete to remove a network pair from the list. When you are done, click Submit.' A table with two columns, 'Between LAN side network' and 'And WAN side network', contains one row with an asterisk (*) in the first column and '10.10.20.0/255.255.255.0' in the second column. There are 'Add' and 'Delete' buttons next to the table row, and a 'Submit' button at the bottom.

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

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

3. Click Submit to enter the changes.

Adding Static Routes

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

- Add static routes manually, as described here
- Add dynamic routes by enabling OSPF and/or RIP (v1 or v2), or by periodically polling the routing table of a Cisco router (refer to “Configuring Dynamic Local Routes” on page 126)
- Import a file of routes from an FTP server (refer to the *WX/WXC Operator’s Guide*)

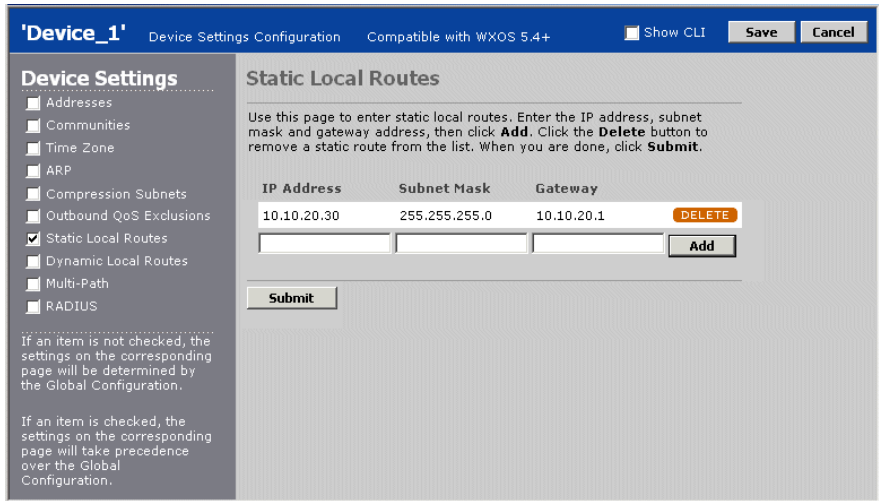
Each device can have a total of 8192 routes (static and dynamic).

If a subnet’s gateway is on the LAN side of the device (as determined by ARP), the subnet is added to the list of 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 109).

To manually add static network routes:

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

Figure 51: Adding a New Local Static Route



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

When you load the configuration, the static routes defined here replace the static routes defined on the device (if any). Also, LAN-side static routes are added to the compression subnets and advertised automatically to other devices, except when the WAN compression subnets option is enabled (refer to “Advertising Compression Subnets” on page 109).

Configuring Router Polling

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



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

Configuring Route Polling

To enable route polling:

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

Figure 52: Enabling Router Polling

The screenshot shows the 'Device_1' configuration window. The left sidebar lists various settings, with 'Dynamic Local Routes' checked. The main panel is titled 'Dynamic Local Routes' and contains three radio button options: 'No Dynamic Local Routing' (selected), 'Use OSPF/RIP', and 'Obtain routing table from router'. Under 'Use OSPF/RIP', there are sub-options for 'OSPF' and 'RIP', each with 'Start' and 'Stop' radio buttons and a corresponding configuration button ('OSPF...' and 'RIP...'). Below these, a note states: 'OSPF and RIP are configured from Basic Setup Partial and Global configurations only.' The 'Obtain routing table from router' option has a 'Poll router:' label and a 'Router...' button. At the bottom of the main panel are 'Submit' and 'Reset' buttons.

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

3. Specify the following information:

Poll router	Enter the IP address of a Cisco router and the port number used for <i>rsh</i> (the standard port is 514). NOTE: The IP address must be on the same subnet as the device.
Secondary router	Enter the IP address and port of a secondary Cisco router to be used when the primary router is unavailable.
Local user name	Enter a local user name that matches the <i>remote</i> user name specified on the Cisco router.
Remote user name	Enter a remote user name that matches the <i>local</i> user name specified on the Cisco router.
Polling 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 Dynamic Local Routes page.

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

If a subnet's gateway is on the LAN side of the device (as determined by ARP), the subnet is added to the list of 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 109).

Configuring a Cisco Router for Route Polling

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

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

Configuring Multi-Path Addresses

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

To use Multi-Path, you configure both devices so that outgoing packets intended for the secondary path are marked with a secondary source IP address and, optionally, with a specific gateway address or ToS/DSCP value. For more information about Policy-Based Multi-Path, refer to “Configuring Policy-Based Multi-Path” on page 263.

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

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

Figure 53: Multi-Path Secondary IP Address

'Device_1' Device Settings Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Device Settings

- ☐ Addresses
- ☐ Communities
- ☐ Time Zone
- ☐ ARP
- ☐ Compression Subnets
- ☐ Outbound QoS Exclusions
- ☐ Static Local Routes
- ☐ Dynamic Local Routes
- ☒ **Multi-Path**
- ☐ RADIUS

If an item is not checked, the settings on the corresponding page will be determined by the Global Configuration.

If an item is checked, the settings on the corresponding page will take precedence over the Global Configuration.

Multi-Path

Secondary IP Address

Supplemental Marking Method

Gateway IP Primary Secondary

Submit Reset

When Multi-Path is enabled, in the event of path degradation or failure, traffic is automatically diverted between the Primary and Secondary paths.

A link is considered to be degraded when excessive latency or packet loss is observed for a period of time.

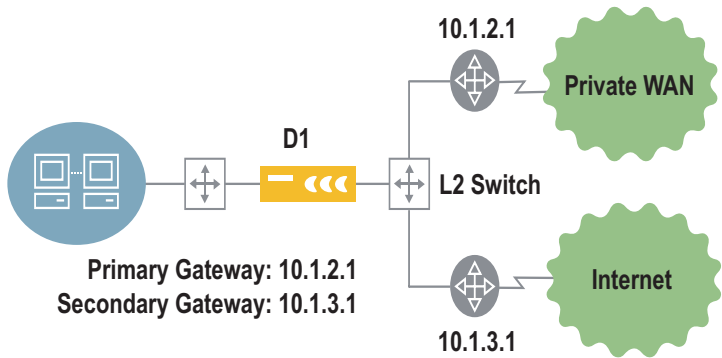
Supplemental marking methods can be used to mark traffic for diversion to the Primary or Secondary paths. If you intend to use Gateway IP as a marking method, fill in the Primary and Secondary gateway IP addresses.

2. Specify the following information:

Secondary IP Address	<p>Enter an IP address to be used as the source address on packets to be sent on the secondary path (packets sent on the primary path have the device address). The secondary IP address must be unique, and must be on the same subnet as the device address.</p> <p>Unless the WAN routers for the primary and secondary paths are also on this subnet (see Gateway IP below), the default gateway must be configured to route traffic with this source address to the appropriate WAN link (refer to “Configuring Routers to Support Multi-Path” on page 270).</p> <p>NOTE: If you enter an address assigned to another device, the path will remain inactive.</p>
Gateway IP	<p>If the WAN routers for the primary and secondary paths are on the same subnet as the WX, and the WX is connected to a Layer 2 switch (see Figure 54), enter the gateway IP addresses here.</p> <p>ARP is used to obtain the MAC addresses for the two gateways, and then traffic for the primary and secondary paths is marked with the MAC address of the appropriate gateway. In this case, no additional router configuration is needed.</p>

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

Figure 54: Multi-Path with Primary and Secondary Gateways



Configuring the RADIUS Source Address

The Device Settings partial configuration lets you specify an alternate source IP address for the RADIUS client. If you are using RADIUS servers to authenticate users (refer to “Defining RADIUS Servers” on page 135), replies from the RADIUS servers are sent to the specified source address. By default, all replies are sent to the device’s IP address.

To specify the RADIUS source address:

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

Configuring Basic Setup Parameters

The following sections describe the basic setup configuration settings:

- “Configuring the Interface Settings” on page 119
- “Configuring NTP” on page 122
- “Enabling SNMP” on page 123
- “Defining Syslog Servers” on page 124
- “Configuring Dynamic Local Routes” on page 126
- “Enabling Route-Based Router Balancing” on page 128
- “Designating a Registration Server” on page 130
- “Generating NetFlow Records” on page 131

Configuring the Interface Settings

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



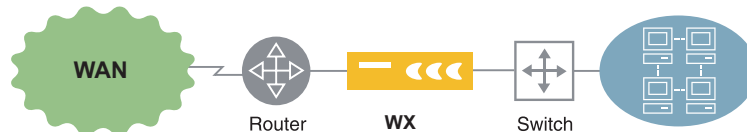
NOTE: The WX 15, WX 20, WX 50, and WXC 250 have two 10/100 NICs. The WX 60, WX 100, WXC 500, and WXC 590 have two 10/100/1000 NICs. The fiber WX 100 supports only 1 Gigabit speeds at full-duplex.

The interface settings let you do the following:

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

If you enable high-availability support, a failure detected on one interface causes the other interface to be turned off for 15 seconds. This allows the switch or router to detect the failure, and ensures that the routing mechanisms work as expected (Figure 55).

Figure 55: 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.

Note that you can also disable hardware passthrough so that the router detects the loss of traffic if the WX fails. On the WX 15, WX 60, WXC 500, and older, non-fiber WX 100s, you can press the Bypass Disable button on the back panel. On newer, non-fiber WX 100s (connectors on the front panel) and on all WXOS 5.4 devices, you can use the following CLI command to enable or disable hardware passthrough (enabled by default):

```
embed bypass-capability <on | off>]
```

To configure the interface settings:

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

Figure 56: Configuring Interface Speed and Duplex Mode Settings

The screenshot shows the 'Main' configuration window for WXOS 5.4+. The left navigation pane has 'Global' selected, with 'Basic Setup' expanded and 'Interfaces' checked. The main content area is titled 'Interfaces' and contains the following settings:

- LOCAL Speed:** Radio buttons for 'Auto' (selected) and 'Manual'. A dropdown menu is next to 'Manual'.
- REMOTE Speed:** Radio buttons for 'Auto' (selected) and 'Manual'. A dropdown menu is next to 'Manual'.
- Local link failure propagation:** An unchecked checkbox with the text: 'If LINK goes down on local interface, the remote interface will be shut down temporarily.'
- Remote link failure propagation:** An unchecked checkbox with the text: 'If LINK goes down on remote interface, the local interface will be shut down temporarily.'

At the bottom of the main area are three buttons: 'Submit', 'Reset', and '802.1q...'.

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

- **VLAN ID.** Enter a VLAN ID (1 through 4095) for the port where the Local interface of the device is connected. On ports that have multiple VLANs, specify the VLAN that has the largest number of hosts. Note that the device resides on one VLAN, but can 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.

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

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

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

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

Since both LAN and WAN-side subnets will be eligible for compression, be sure to advertise only the true LAN-side subnets (refer to “Advertising Compression Subnets” on page 109).

Configuring NTP

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



NOTE: Using an NTP server is highly recommended if you poll the devices hourly for performance statistics. If a device is more than three minutes slow, its hourly data may not be counted in the correct hour, making the hourly reports inaccurate (reports for longer periods will be correct). To change the polling interval, refer to “Configuring Device Polling” on page 405.

To configure NTP servers:

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

Figure 57: Configuring NTP

The screenshot shows the 'Main' Global Configuration window. The top bar includes 'Global Configuration' and 'Compatible with WXOS 5.4+'. On the right of the top bar are 'Show CLI', 'Save', and 'Cancel' buttons. The left sidebar is titled 'Global' and contains a 'Basic Setup' section with a list of items: 'Interfaces' (checked), 'Time' (checked), 'SNMP', 'Syslog server', 'Dynamic Local Routes', 'Router Balancing', 'Registration server', and 'NetFlow'. Below this is an 'AAA' section. The main content area is titled 'Time' and contains a checkbox 'Use NTP Server' which is checked. To the right of this checkbox are two input fields: 'Primary:' and 'Secondary:'. The 'Primary:' field is labeled 'IP address' and the 'Secondary:' field is labeled 'IP address (optional)'. At the bottom of the main area are 'Submit' and 'Reset' buttons.

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

Enabling SNMP

The following SNMP support is provided:

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



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

The Enterprise MIB can be used to view device performance statistics from a Network Management System (NMS). In addition, the devices can send SNMP traps to the NMS and other network devices. For a description of the SNMP traps, refer to “WX System Events and SNMP Traps” on page 442.

To enable SNMP:

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

Figure 58: Enabling SNMP

2. Select the SNMP Enabled check box to enable SNMP, and then enter the Read and Write Community Strings used by the NMS to access SNMP data. The defaults are “public” and “private”.
3. Select the Trap Enabled check box to generate SNMP traps (version 2 traps only).

4. To add one or more SNMP trap destinations:
 - a. For WXOS 5.4 configurations, enter an IP address and an associated community string (up to 30 characters), and click Add. To delete a trap destination, click Delete next to the destination.
 - b. For WXOS 5.3/5.2 configurations, enter one trap community string, and then enter the destination IP addresses (one per line). The default community string is “trap community”.
5. Select the Authentication Trap Enabled check box to generate traps for incorrect logins and unauthorized user access attempts.
6. Click Submit to enter the changes, or click Reset to discard them.

Defining Syslog Servers

Syslog messages can be sent to up to five Syslog servers. Syslog servers let you centrally log and analyze configuration events and system error messages, such as interface status, security alerts, and environmental conditions. To view events from WXOS 5.4 devices in CMS, each device must specify the CMS server as a Syslog server.

For a description of Syslog messages generated by CMS and the WX devices, refer to “System Events” on page 441.

To enable Syslog reporting:

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

Figure 59: Enabling Device Syslog Reporting

The screenshot shows the 'Main' Global Configuration interface for WXOS 5.4+. The left sidebar lists configuration categories: Global, Basic Setup, AAA, Applications, Compression, and QoS. Under 'Basic Setup', 'Syslog server' is selected. The main panel is titled 'Syslog Server' and contains the following options:

- Enable syslog reporting:** A checkbox labeled 'Yes' is checked.
- Syslog servers:** A text input field is empty. A tooltip indicates: 'Enter IP addresses, one per line with a maximum of 5 servers.'
- Syslog message severity:** Four checkboxes are present: 'Critical' (checked), 'Error' (checked), 'Information' (unchecked), and 'Notice' (unchecked). A tooltip indicates: 'Check message severity levels you want reported to the syslog server.'

At the bottom of the panel are 'Submit' and 'Reset' buttons. The top right of the interface includes a 'Show CLI' button and 'Save' and 'Cancel' buttons.

2. Select the Yes check box to enable Syslog reporting, and then enter the IP addresses of up to five Syslog servers (one per line).
3. Select the severity levels of the messages sent to the Syslog server:
 - **Critical.** Critical error messages about software or hardware malfunctions.
 - **Error.** Error messages, such as License expired.
 - **Informational.** Informational messages, such as reload requests and low-process stack messages.
 - **Notice:** Informational messages about unusual events that are not errors.
4. Click Submit to enter the changes, or click Reset to discard them.

Configuring Dynamic Local Routes

If your network uses OSPF or RIP, you can enable these protocols to discover routes dynamically on the local and remote sides of each device. Alternatively, you can configure a device to periodically poll a Cisco router on the same subnet (refer to “Configuring Router Polling” on page 114).

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



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

If a subnet's gateway is on the LAN side of the device (as determined by ARP), the subnet is added to the list of 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 109).

To configure RIP and/or OSPF:

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

Figure 60: Configuring RIP and OSPF

The screenshot shows the 'Dynamic Local Routes' configuration window. The left-hand navigation pane is expanded to 'Dynamic Local Routes'. The main area has the following settings:

- Dynamic Local Routes:**
 - ☒ No Dynamic Local Routing
 - ☐ Use OSPF/RIP
 - OSPF:** ☐ Start (Not configured) ☒ Stop [OSPF...](#)
 - RIP:** ☐ Start (Default) ☒ Stop [RIP...](#)
 - ☐ Obtain routing table from router
 - Poll router:** ☐ Start (Not configured) ☒ Stop [Router...](#)
 - Router polling is configured from Device Settings partial configurations only. Router polling should be used with Cisco routers only. For non-Cisco routers, OSPF or RIP are recommended.
- ☐ Static routes take precedence over Dynamic routes

Buttons at the bottom: [Submit](#) [Reset](#)

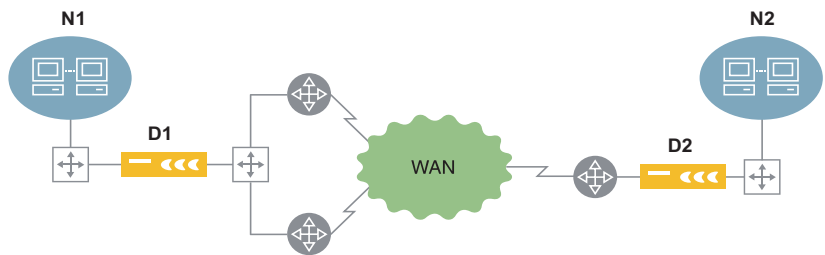
2. To enable OSPF:
 - a. Click OSPF... and enter the Area ID for OSPF.
 - b. If your network uses OSPF authentication, select Password and enter the password (up to 8 characters), or select MD5 and enter the key ID (0 to 255) and the MD5 key (up to 16 characters).
 - c. Click Submit.
 - d. Click Use OSPF/RIP, and select Start next to OSPF.
 - e. Click Submit to enter the changes, or click Reset to discard them.
3. To enable RIP:
 - a. Click RIP... and select the version of RIP used in your network (1 or 2).
 - b. If your network uses RIP authentication, select Password and enter the password (up to 15 characters).
 - c. Click Submit.
 - d. Click Use OSPF/RIP, and select Start next to RIP.
 - e. Click Submit to enter the changes, or click Reset to discard them.
4. By default, dynamic routes take precedence over static routes to the same destination. To give precedence to static routes, click Static routes take precedence over Dynamic routes, and click Submit.

Enabling Route-Based Router Balancing

You can configure devices to balance the compressed traffic load across multiple routers that have equal-cost paths to the same destination (route-based balancing). To configure a router to distribute traffic based on ToS values set by a WX (ToS marking for router-based balancing), refer to the “configure route” CLI command in the *WX/WXC Operator's Guide*.

Using route-based balancing, compressed traffic can be distributed across up to four different gateways. In Figure 61, 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 61: Configuring Router Balancing Policies



To identify gateways (up to four) that have equal cost paths to the same IP address, open the device Web console for a device and click Local Routes. Equal cost paths are grouped together in the Local Routes page (Figure 62).

Figure 62: Common Routes with Equal Cost Paths

53/22-Carson
Device Setup Compression QoS Acceleration Monitor Admin Help
Logged in as: admin Save Logout

Setup

Basic

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

AAA

Applications

IPSec

Advanced

Local Routes

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
		192.168.0.2	Dynamic

OSPF: Stopped
RIP: Stopped
Router Polling: None

☐ Static routes take precedence over dynamic routes

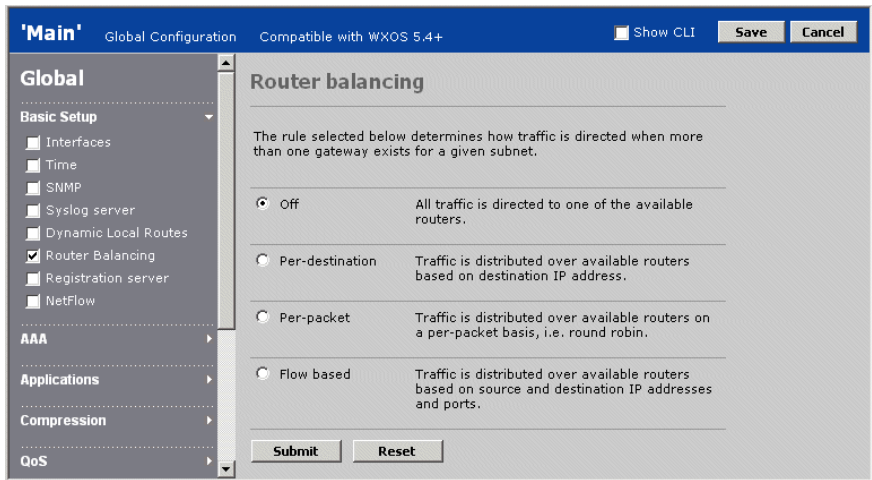
Submit Static... Dynamic... Import... Router Balancing...

Equal cost paths to the same destination

To enable route-based router balancing:

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

Figure 63: Configuring Route-Based Router Balancing



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



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

- **Flow based.** Traffic is distributed over available routers based on source and destination IP addresses and ports.
3. Click Submit to enter the changes, or click Reset to discard them.

Designating a Registration Server

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

In global configurations and Basic Setup partial configurations, a registration server address and password must be specified if the Registration Server check box is selected. If you change the password defined on a registration server, you can update the password in CMS, and download the new password to all devices (refer to “Changing a Registration Server Address or Password” on page 401).



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

To specify the registration server:

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

Figure 64: Designating a Registration Server

The screenshot shows the 'Main' configuration window with the 'Registration Server' tab selected. The left-hand navigation pane shows 'Basic Setup' expanded, with 'Registration server' checked. The main configuration area for 'Registration Server' includes the following fields and options:

- Registration server:** A text input field for the IP address.
- Password:** A password input field with masked characters (dots).
- SECONDARY registration server:** A section with two radio button options:
 - ☒ None
 - ☐ Use IP address (with an associated text input field)
- Buttons:** 'Submit' and 'Reset' buttons at the bottom.

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

When you save or download a configuration, an error occurs if the password does not match the current or previous password stored in CMS for the specified registration server (refer to “Changing a Registration Server Address or Password” on page 401).

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

Generating NetFlow Records

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

NetFlow data is sent in Version 5 format, as described in “NetFlow Version 5 Export” on page 453.

To generate NetFlow records:

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

Figure 65: Generating NetFlow Records

The screenshot shows a web-based configuration interface for a Cisco device. The top navigation bar includes 'Main', 'Global Configuration', and 'Compatible with WXOS 5.4+'. On the left, a 'Global' navigation pane lists various configuration options: Interfaces, Time, SNMP, Syslog server, Dynamic Local Routes, Router Balancing, Registration server, and NetFlow (which is checked). The main content area is titled 'Top Traffic > NetFlow'. It contains a message: 'In order to use the NetFlow Export feature the IP Address and Port of the NetFlow server must be entered below.' Below this message are two input fields: 'IP Address' and 'Port'. There is a checkbox labeled 'Enable NetFlow' which is checked. At the bottom of the form are 'Submit' and 'Reset' buttons. A footer note at the bottom right states: '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 enter the changes, or click Reset to discard them.

Configuring AAA Settings

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

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

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

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

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

The following sections describe the AAA configuration settings:

- "Selecting Authentication Methods" in the next section
- "Enabling Authorization Checking" on page 134
- "Defining RADIUS Servers" on page 135
- "Defining Local Users" on page 137
- "Securing Operator Access" on page 138
- "Securing Front Panel Access" on page 140

Selecting Authentication Methods

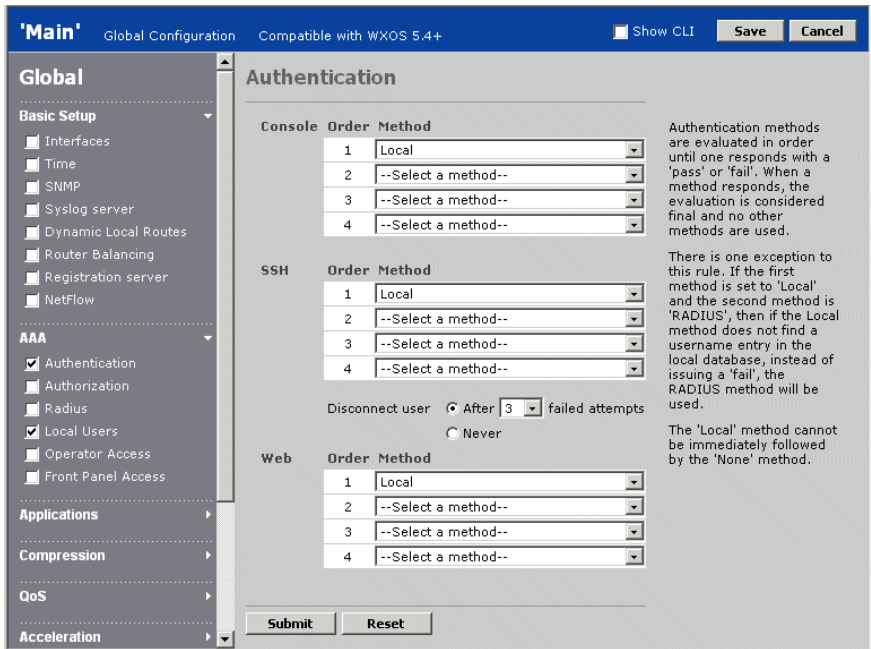
You can specify the order in which a device’s local database and RADIUS server groups are accessed to authenticate users on the Web, the SSH (CLI), and the console port. You can also specify the number of SSH login attempts allowed before a user is locked out. By default, all users are authenticated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

To define RADIUS servers and server groups, refer to “Defining RADIUS Servers” on page 135. To define user accounts locally, refer to “Defining Local Users” on page 137.

To select the authentication methods for each user interface:

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

Figure 66: Selecting Authentication Methods



2. Specify the following information:

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

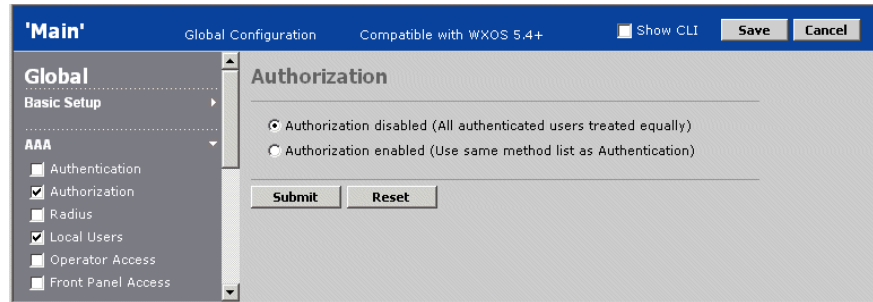
3. Click Submit to enter the changes, or click Reset to discard them.

Enabling Authorization Checking

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

To enable or disable authorization checking:

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

Figure 67: Enabling Authorization Checking

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

Defining RADIUS Servers

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

To specify the server groups used for authentication, refer to “Selecting Authentication Methods” on page 133.

To define RADIUS servers and server groups:

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

Figure 68: Defining RADIUS Servers and Server Groups

The screenshot shows the 'RADIUS' configuration page. On the left is a sidebar with a tree view containing 'Global', 'Basic Setup', 'AAA' (with sub-items: Authentication, Authorization, Radius, Local Users, Operator Access, Front Panel Access), 'Applications', and 'Compression'. The main content area is titled 'RADIUS'. It features two tables. The first table, 'RADIUS Servers', has columns: Name, IP Address, Auth. Port, Time-out (sec), Re-transmit, Dead Time (min), and Delete. It contains one entry: 'Main' with IP 10.20.30.40, Auth. Port 1812, Time-out 3, Re-transmit 3, and Dead Time 0. Below this table is a 'New Server...' button. The second table, 'RADIUS Server Groups', has columns: Name and Delete. It contains one entry: 'Central'. Below this table is a 'New Group...' button. At the bottom of the main area are 'Submit' and 'Reset' buttons. The top of the page has a blue header with 'Main', 'Global Configuration', 'Compatible with WXOS 5.4+', 'Show CLI', 'Save', and 'Cancel' buttons.

From the RADIUS page, you can:

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

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

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

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

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

Defining Local Users

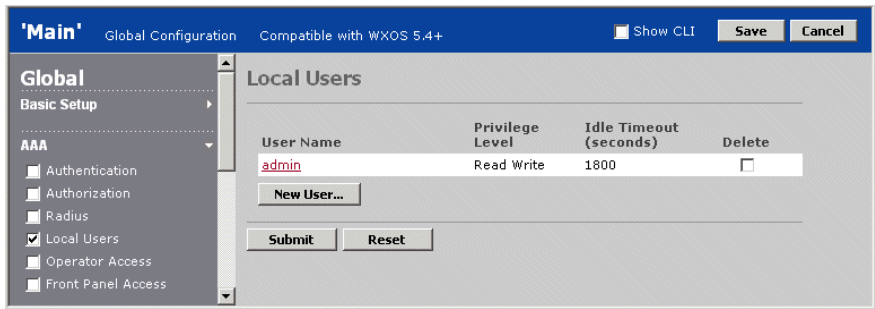
You can define up to 25 users that can be authenticated locally by each device. Each user can have full (admin) or read-only access privileges. The default password (juniper) of the predefined admin account must be changed for all new global configurations and for new AAA partial configurations where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

To specify how users are authenticated (locally and/or through RADIUS), refer to “Selecting Authentication Methods” on page 133.

To define local user accounts:

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

Figure 69: Defining Local Users



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

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



NOTE: Authorization checking is disabled by default, so that all authenticated users have read-write access and a 30-minute idle timeout. If you create read-only user accounts or change the default idle timeout, you must enable authorization checking (refer to “Enabling Authorization Checking” on page 134).

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

Securing Operator Access

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

To restrict operator access:

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

Figure 70: Configuring Device Operator Access

'Main' Global Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Global

Basic Setup

AAA

- Authentication
- Authorization
- Radius
- ☒ Local Users
- ☒ Operator Access
- Front Panel Access

Applications

Compression

QoS

Acceleration

Advanced Setup

Multi-path

IPSec

Events

Operator Access

The following lists are used to restrict operator access to this Juniper WX device from designated valid client addresses only. If both lists are empty, then operator access is unrestricted. If an address/subnet is entered in the Include list, then all other addresses/subnets are denied access.

Enter addresses/subnets, one per line. For an individual client, enter the IP address only. For a subnet, enter the IP address and subnet mask separated by a slash (/).

Example:
123.123.123.123
123.123.123/255.255.255.0

Also, if you want to preserve the changes, you must save the configuration to flash memory using the 'Save Configuration' page under the 'Maintenance' tab after submit.

Include list

Enter addresses/subnets which should have access to this Juniper WX device, one per line.

Exclude list

Enter addresses/subnets which should be denied access to this Juniper WX device, one per line.

- To allow access to a device only from specific IP addresses or subnets, enter the addresses or subnets in the Include list (one per line). The subnet format is:

<IP address>/<subnet mask>

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

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



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

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

Securing Front Panel Access

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



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

To lock the front panel keypad:

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

Figure 71: Securing Front Panel Access



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

Configuring Application Settings

Application definitions allow each device to identify the traffic of up to 256 applications (the 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”.

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

- Enable or disable compression, as described in “Compressing Applications” on page 156.
- Enable or disable monitoring of applications for reports, as described in “Monitoring Applications” on page 151.
- Accelerate an application’s traffic (if compression is enabled), as described in “Configuring Traffic Acceleration” on page 206.
- Assign the application to a traffic class to manage its outbound bandwidth allocation, as described in “Assigning Applications to Traffic Classes” on page 149. Traffic classes are also used for path optimization, as described in “Configuring Policy-Based Multi-Path” on page 263.
- Trigger events based on the application’s performance, as described in “Configuring Events” on page 280.
- View compression and acceleration statistics for monitored applications, as described in “Monitoring Performance” on page 305.

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

Default Application Definitions for WXOS 5.4 Devices

Table 6 lists the default application definitions for WXOS 5.4 devices. Each definition has rules to match any traffic that has the specified source or destination port number(s).

Table 6: Default Application Definitions

Application	Order	Port Numbers
AOL	47	5190-5193
CIFS	16	139, 445
Clearcase	34	371
CVS	44	2401
DNS	25	53
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

Application	Order	Port Numbers
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
SNTP	24	123
Microsoft SQL Server	18	No ports specified
SSH	23	22
Sybase SQL Anywhere	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

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

To view the application overview:

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

Figure 72: Application Overview Page

Application Name	Traffic Class	Compression NSC	TCP Acceleration	Fast Connection Setup	CIFS Acceleration	Exchange Acceleration	HTTP Acceleration	Monitoring
AOL		✓	✓	✓	✓	✓	✓	✓
CIFS		✓	✓	✓	✓	✓	✓	✓
Clearcase		✓	✓	✓	✓	✓	✓	✓
CVS		✓	✓	✓	✓	✓	✓	✓
DNS		✓	✓	✓	✓	✓	✓	✓
Exchange		✓	✓	✓	✓	✓	✓	✓
Filenet		✓	✓	✓	✓	✓	✓	✓
FTP		✓	✓	✓	✓	✓	✓	✓
Groupwise		✓	✓	✓	✓	✓	✓	✓
H.248		✓	✓	✓	✓	✓	✓	✓
H.323		✓	✓	✓	✓	✓	✓	✓
Hostname Resolution		✓	✓	✓	✓	✓	✓	✓
HTTP		✓	✓	✓	✓	✓	✓	✓

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

Traffic Class	Traffic class assigned to the application. To change the traffic class, refer to “Assigning Applications to Traffic Classes” on page 149.
Compression	Indicates whether the application’s traffic is compressed (refer to “Compressing Applications” on page 156.).
NSC	Indicates whether Network Sequence Caching is used for compression (refer to “Compressing Applications” on page 156). 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 220).

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 221).
CIFS Acceleration	Indicates whether CIFS traffic for the application is accelerated (refer to “Enabling Microsoft CIFS Acceleration” on page 222).
Exchange Acceleration	Indicates whether Exchange traffic for the application is accelerated (refer to “Enabling Microsoft Exchange Acceleration” on page 225).
HTTP Acceleration	Indicates whether HTTP traffic for the application is accelerated (refer to “Enabling HTTP Acceleration” on page 227).
Monitoring	Indicates whether you can view statistics for the application, as described in “Monitoring Applications” on page 151.

Configuring Application Definitions

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

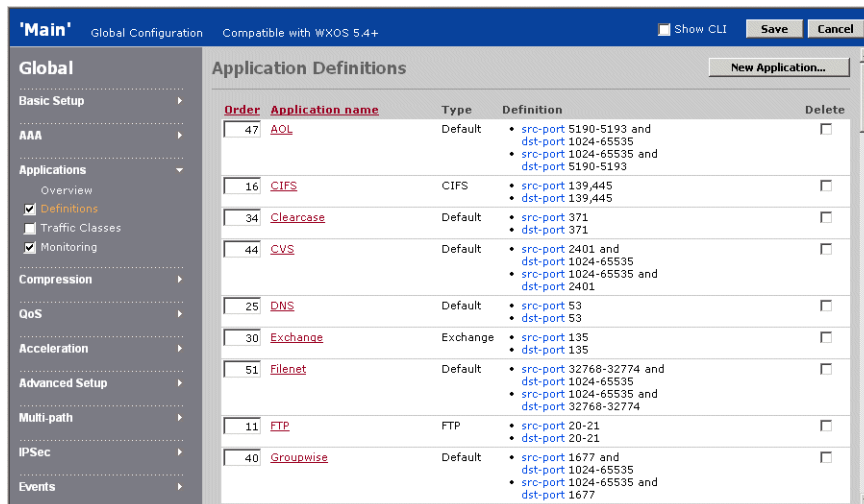
A packet matches an application definition if a match occurs on any of its rules. All the values defined in the same rule must be true for a match to occur on that rule. A packet is classified under the first application for which a rule match is found. Packets are compared against the definitions according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



NOTE: In the device Web console, you can add new definitions by selecting undefined applications from the Top Traffic report, as described in the *WX/WXC Operator’s Guide*. You can then extract the configuration settings from the device (refer to “Extracting Configurations” on page 85).

To add or change application definitions:

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

Figure 73: Application Management Page

From the Application Management page, you can:

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

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

Figure 74: Defining New Applications

Global ConfigurationCompatible with WXOS 5.4+Show CLISaveCancel

Global

Basic Setup

AAA

Applications

Overview

Definitions

Traffic Classes

Monitoring

Compression

QoS

Acceleration

Advanced Setup

Multi-path

IPSec

Events

Application definition > New

Application Name:

Application Type:Default

Application traffic will be identified using the following rules

Source Address	Source Port	Destination Address	Destination Port	Protocol	Advanced
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR
				Any	AdvancedCLEAR

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
(*).

SubmitCancel

3. Specify the following information:

Application name Enter a name for the application (up to 63 characters). Use only letters, numbers, blanks, and the following special characters:

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

- Application type Select one of the following application types:
- **Default.** No special processing.
 - **CIFS.** Apply to CIFS application definitions whose traffic you want to accelerate (refer to “Enabling Microsoft CIFS Acceleration” on page 222). 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 225). 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 227). Also allows a URL to be specified for pattern matching.

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

Source Address	<p>Enter a source IP address or subnet. The general format is:</p> <p>address/subnet_mask</p> <p>A blank indicates any source IP address.</p>
Source Port	<p>Enter a source port number, a series of comma-separated port numbers, or a range of port numbers separated by a hyphen (-). A blank indicates any port. For a list of common application port numbers, refer to “Common Application Port Numbers” on page 465.</p>
Destination Address	<p>Enter a destination IP address or subnet (same format as the source address). A blank or asterisk (*) indicates any destination IP address. Typically, source and destination addresses are specified in separate rules so that a match occurs on either one. A rule that specifies both source and destination addresses will match only the traffic between those addresses.</p>
Destination Port	<p>Enter one or more destination port numbers (same format as the source port). A blank indicates any port. Typically, source and destination ports are specified in separate rules so that a match occurs on either one. A rule that specifies source and destination ports will match only the traffic between those ports.</p>
Protocol	<p>Select an application protocol or select Any to indicate TCP or UDP. You can also type in a protocol number (0 to 134). By default, a match can occur on any TCP or UDP packet.</p> <p>NOTE: Any protocol defined by number is added to the Any list of defaults that applies to each rule that does not specify a protocol. To use application pattern matching (described below), select TCP.</p>

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

ToS Bits	<p>Select the check box, and then select one of the following:</p> <ul style="list-style-type: none"> ■ ToS. Select an IP precedence value (0 through 7). ■ DSCP. Enter a DSCP value (0 through 255). <p>For more information about ToS and DSCP, refer to “Changing Outbound ToS/DSCP Values” on page 199.</p>
----------	---

Application pattern matching If the application type is HTTP or Citrix, you can enter a URL or a Citrix client and/or application name.

A URL can be up to 127 characters. The general format is:

`<host> / <uri>`

Where:

<host> is up to eight strings separated by periods. You can use an asterisk (*) by itself to indicate any string. For example:

`www.juniper*.net/`

The slash is required even when only the host is specified. Consecutive periods, such as "...", are interpreted as ".*.*.*", and will match any host name.

<uri> is up to eight strings separated by slashes. You can use an asterisk (*) by itself to indicate any string.

For example:

`www.juniper*.net/* /index.htm`

When an asterisk is part of a string, it is treated as a single character (not a wildcard), such as "www.juniper*.net".

Click Continue to return to the Application Definition page.

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

Testing New Application Definitions

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

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

Assigning Applications to Traffic Classes

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

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

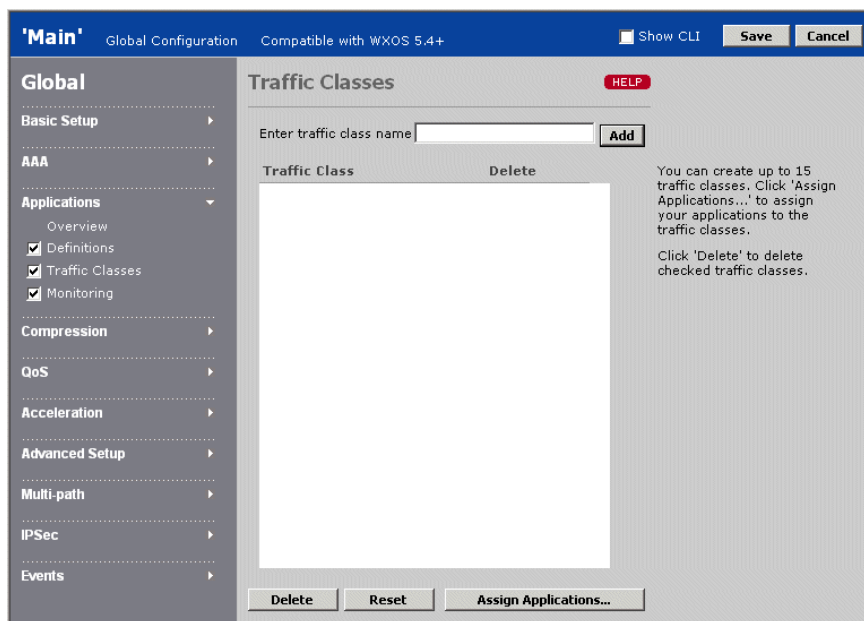
For more information about outbound QoS and Multi-Path, refer to:

- “Configuring QoS Settings” on page 168
- “Configuring Policy-Based Multi-Path” on page 263

To define traffic classes and assign applications to each class:

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

Figure 75: Assigning Applications to Traffic Classes



From the Traffic Classes page, you can:

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



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

- Change a class name. Click the class name, enter the new name, and click Submit.

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

Monitoring Applications

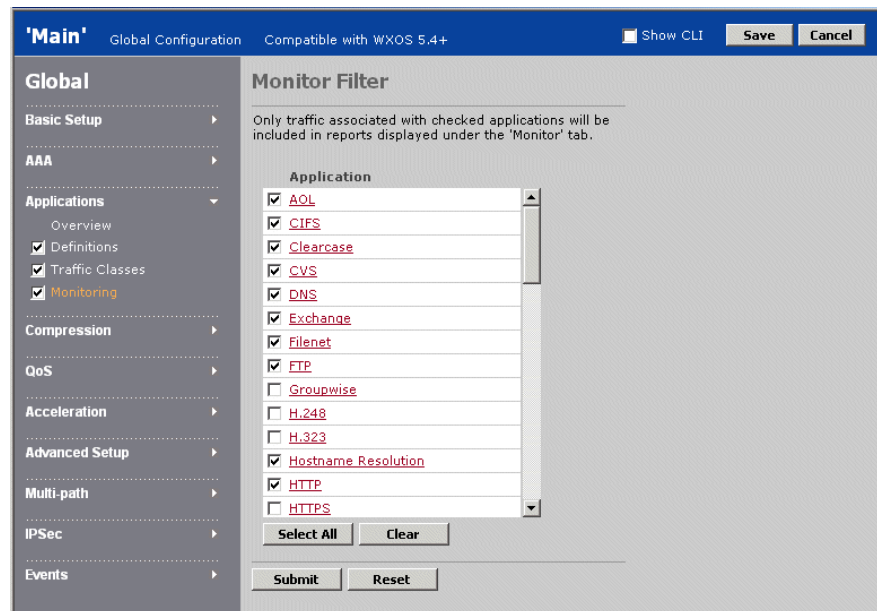
Monitoring an application lets you view compression and acceleration statistics for the application. You can monitor up to 40 applications. All unmonitored applications are placed in the “Others” category on reports.

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 145.

To select applications to be monitored:

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

Figure 76: 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.
3. Click Submit to enter the changes, or click Reset to discard them.

Configuring Compression Settings

The following sections describe the global compression parameters:

- “Configuring Endpoints for Compression” on this page
- “Configuring Network Sequence Caching” on page 154
- “Compressing Applications” on page 156
- “Configuring Remote Routes” on page 158
- “Configuring Tunnel Load Balancing Policies” on page 160
- “Configuring Default Decompressors” on page 162
- “Defining Preferred Decompressors” on page 165
- “Configuring Tunnel Mode Settings” on page 167

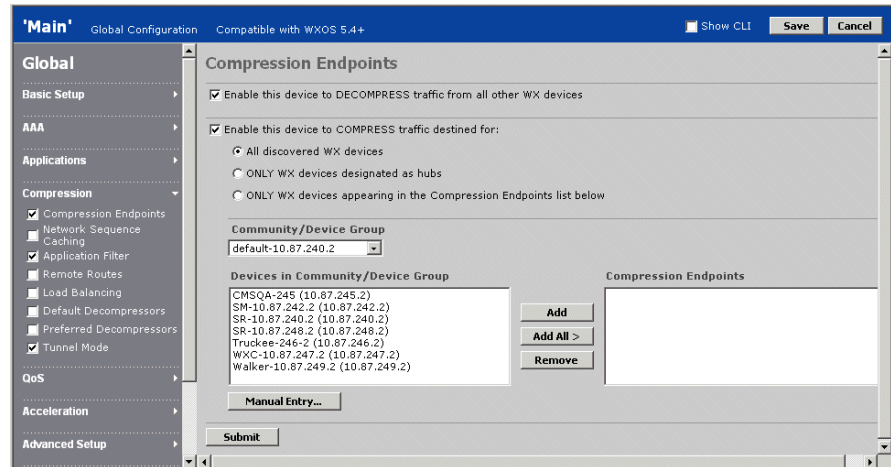
Configuring Endpoints for Compression

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

Compression and decompression begins automatically for the compression subnets that are advertised (refer to “Advertising Compression Subnets” on page 109). If necessary, you can disable compression or decompression for all remote devices, and/or compress data only for specific devices in each community. Each device can belong to multiple communities.

To configure the endpoints for compression:

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

Figure 77: Configuring Endpoints for Compression

2. To stop decompressing data from other devices, clear the Enable this device to DECOMPRESS traffic ... check box. All devices in the community will stop compressing data for devices that have this setting.
3. To stop 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 devices in the same community (default).
 - **ONLY WX devices designated as hubs.** Data is compressed only for devices in the same community that are designated as a hub.
 - **ONLY WX devices appearing in the Compression Endpoints list below.** Data is compressed only for the devices in the Compression Endpoints list.
4. To add devices to the Compression Endpoints list:
 - a. Select a community from the Community/Device Group list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.
 - b. Select the devices you want to enable compression for, and click Add. To remove devices from the Compression Endpoints list, select the devices and click Remove.

- c. Repeat Steps **a** and **b** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - d. If one or more devices you want to add are not listed for the community/device group, you can add the devices manually. Click Manual Entry, enter the device IP addresses (one per line), and click Submit.
5. Click Submit to enter the changes.



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

Configuring Network Sequence Caching

Network Sequence Caching (NSC) is an enhanced 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 the tunnel is down). NSC is most effective where large files are often sent over the WAN, such as for database backups.

To use NSC between two WXC devices, you must enable the following on both devices:

- Compression (refer to “Configuring Endpoints for Compression” on page 152),
- TCP Acceleration and outbound QoS (refer to “Enabling Acceleration by Application” on page 219)

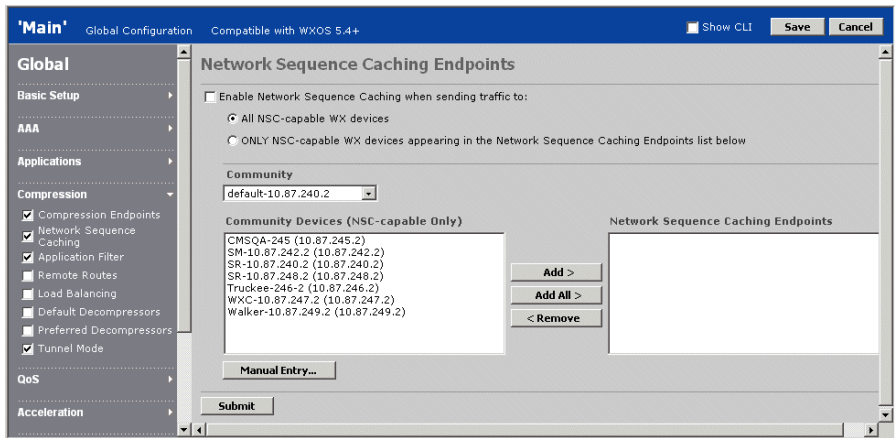
Applications that are enabled for TCP Acceleration can then be enabled for NSC (refer to “Compressing Applications” on page 156).

When you install a new WXC, compression, 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. In the Configuration window, click Compression in the left-hand navigation frame, click Network Sequence Caching, and select the check box.

Figure 78: Configuring Endpoints for Network Sequence Caching



2. To disable NSC on this device so that standard compression is used for all remote devices, clear the Enable Network Sequence Caching... 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 in the list below.** NSC is used only for devices in the Network Sequence Caching Endpoints list.



NOTE: NSC takes effect between two WXC devices only if it is enabled on both devices. NSC settings are ignored if you download the configuration to a WX.

3. To add devices to the Network Sequence Caching Endpoints list:
 - a. Select a community from the Community/Device Group list. The device name and IP address are shown for each WXC device in the selected community/device group. The IP address is enclosed in parentheses.
 - b. Select the devices you want to enable NSC for, and click Add. To remove devices from the list, select the devices and click Remove.

- c. Repeat Steps **a** and **b** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - d. If one or more devices you want to add are not listed for the community/device group, you can add the devices manually. Click Manual Entry, enter the device IP addresses (one per line), and click Submit.
4. Click Submit to enter the changes.

Compressing Applications

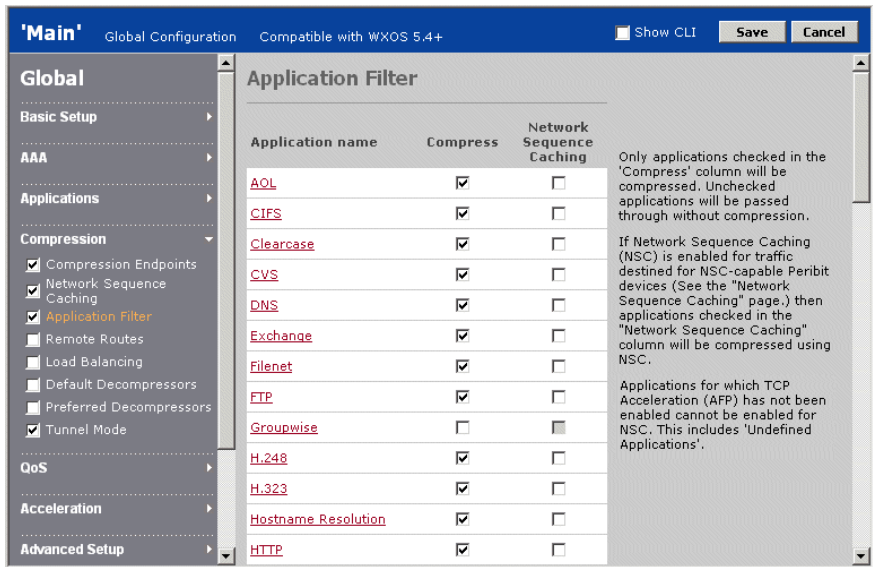
For each application, you can enable or disable compression and Network Sequence Caching (NSC). To conserve system processing capacity, you should disable compression for applications whose traffic is encrypted or already compressed. However, you must compress all TCP applications that you want to accelerate.

Application definitions are provided for applications with well-known port numbers. All other applications are grouped together as “Undefined”. To define additional applications, refer to “Configuring Application Settings” on page 141.

To select the applications to be compressed and monitored:

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

Figure 79: Selecting Applications for Compression and Monitoring



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

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

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 “Configuring Traffic Acceleration” on page 206).</p>
Network Sequence Caching	<p>On WXC devices, you can enable Network Sequence Caching (NSC) for compressed applications. If NSC is enabled for one or more remote WXC devices (refer to “Configuring Network Sequence Caching” on page 154), 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 WXs or to WXCs 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 220).</p>

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



NOTE: NSC settings are ignored if you download the configuration to a WX device.

Configuring Remote Routes

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

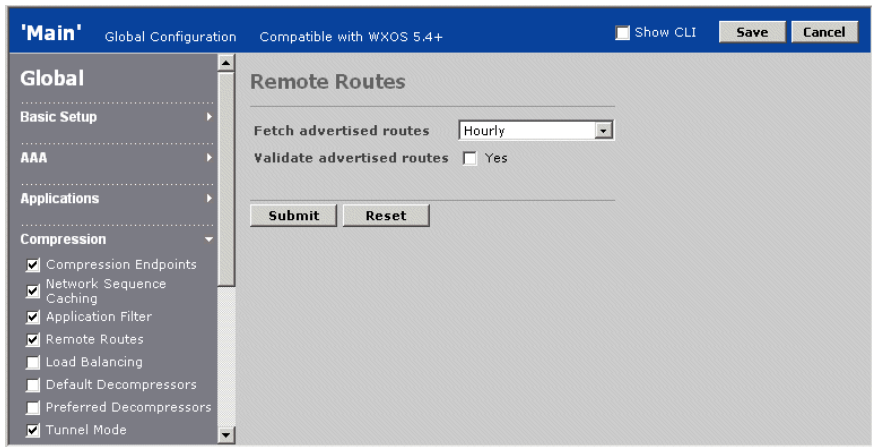


NOTE: Enable the test only if the validity of the remote routes is in question. You should not use this option if load balancing is enabled (refer to “Configuring Tunnel Load Balancing Policies” on page 160).

To configure the remote route settings:

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

Figure 80: Configuring Remote Routes Parameters



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

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

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



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

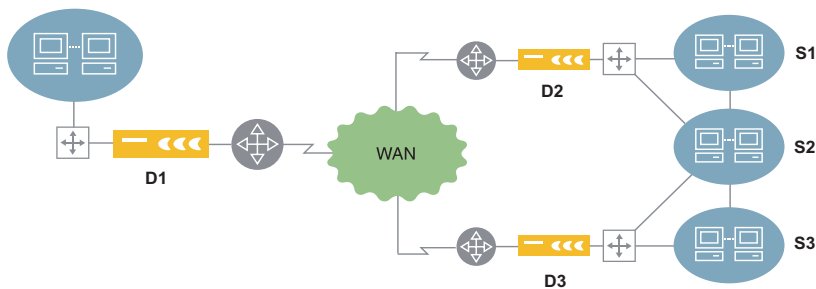
4. Click Submit to enter the changes, or click Reset to discard them.

Configuring Tunnel Load Balancing Policies

If two or more devices in the same community have equal cost paths to the same subnet, you can use load balancing to share the load of decompressing the compressed data. Alternatively, you can specify preferred decompressors, as described in “Defining Preferred Decompressors” on page 165. If neither load balancing nor preferred decompressors are used, the path selection is arbitrary.

For example, in Figure 81, devices D2 and D3 advertise a local route to subnet S2. On D1, the two routes to S2 have equal cost paths.

Figure 81: Sample Load Balancing Scenario

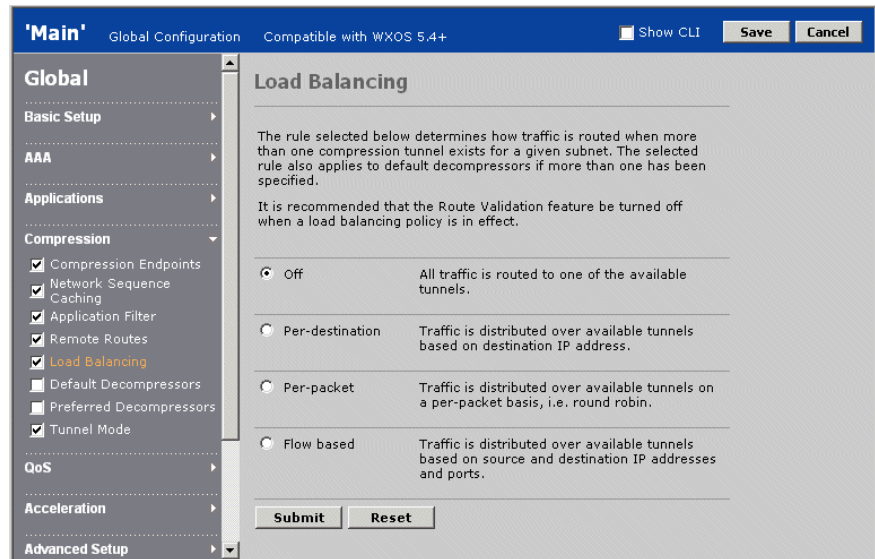


NOTE: If you enable load balancing policies, you should not enable the validate advertised routes feature (refer to “Configuring Remote Routes” on page 158).

To configure load balancing policies:

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

Figure 82: Configuring Load Balancing



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

Configuring Default Decompressors

You can sometimes simplify route administration by designating a 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. Outbound QoS and IPSec encryption also use default decompressors, regardless of whether compression is enabled.



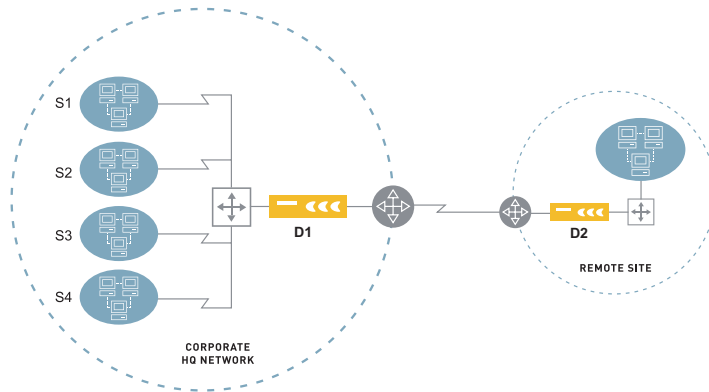
NOTE: Default decompressors cannot be used for an off-path device that uses RIP for packet interception (compression will fail).

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 defined on the device.

Figure 83 shows a simple example of a remote site with one outbound connection to the corporate network. If device D1 is the default decompressor for D2, all traffic that uses the default route on D2 is compressed and sent to D1. To disable compression for traffic sent to subnet S4, you can add S4 to the exclusion list on D2.

Figure 83: Sample Default Decompressor Scenario

You can specify up to six default decompressors on a device. If you specify more than one default decompressor, the current load balancing policies are applied (refer to “Configuring Tunnel Load Balancing Policies” on page 160).

To create a list of default decompressors:

1. In the Configuration window, click Compression in the left-hand navigation frame, click Default Decompressors, and select the check box.

Figure 84: Configuring Default Decompressors

The screenshot shows the 'Main' configuration page with a sidebar on the left and a main content area on the right. The sidebar contains a 'Global' section with a 'Compression' subsection where '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 and a 'Submit' button. The 'Exclude List' box has a text area for entering IP addresses/subnets and a 'Reset' button. The top of the page has a blue header with 'Main', 'Global Configuration', 'Compatible with WXOS 5.4+', 'Show CLI', 'Save', and 'Cancel' buttons.

'Main' Global Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Global

Basic Setup

AAA

Applications

Compression

☒ Compression Endpoints

☒ Network Sequence

☒ Caching

☒ Application Filter

☒ Remote Routes

☒ Load Balancing

☒ **Default Decompressors**

☐ Preferred Decompressors

☒ Tunnel Mode

QoS

Acceleration

Advanced Setup

Multi-path

IPSec

Events

Default Decompressors

If traffic is destined for a subnet that is NOT included in the Remote Routes list, it is normally passed through without compression. However, if a default decompressor (a Juniper WX device) is entered below, this traffic will be compressed and routed to the default decompressor.

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

Default Decompressors

Enter IP addresses of Juniper WX devices, one per line. A maximum of 6 default decompressors may be entered. If more than one default decompressor is entered, then the load balancing policy will be applied. If load balancing is set to "Off", then the precedence of the default decompressors will be based on their order in the list.

Exclude List

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

Submit Reset

2. 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.
3. In the Exclude List box, enter an IP address or an IP address and subnet mask separated by a slash (/) for the hosts or subnets whose traffic is not compressed before being sent to the default decompressor. If you enter an address or subnet that belongs to some other WX, the exclusion is ignored.
4. Click Submit to enter the changes, or click Reset to discard them.
5. Do the following for each default decompressor (log in to the device Web console or load new Device Settings partial configurations from CMS):
 - If dynamic routing is not used, add a static route to each device in the community. The gateway for each route is the default gateway on the Remote interface (the WAN side).
 - Change the default gateway to the IP address of the next-hop router on the Local interface (the LAN side).

Defining Preferred Decompressors

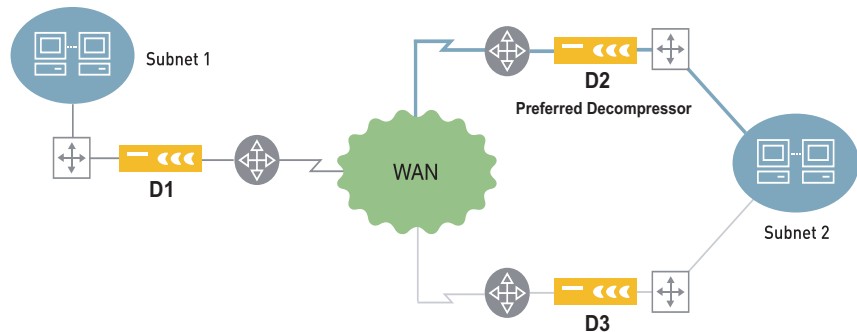
If two or more devices in the same community have equal cost paths to the same subnet, you can control the selected path by specifying a preferred decompressor. Alternatively, you can use load balancing to vary the selected path, as described in “Configuring Tunnel Load Balancing Policies” on page 160. 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 85, data from Subnet 1 has two network paths to Subnet 2. If 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 85: Designating a Preferred Decompressor



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

To create a list of preferred decompressors:

1. In the Configuration window, click Compression in the left-hand navigation frame, click Preferred Decompressors, and select the check box.

Figure 86: Defining Preferred Decompressors

The screenshot shows the 'Main' configuration window with the 'Global Configuration' tab selected. The left-hand navigation pane is expanded to 'Compression', and 'Preferred Decompressors' is selected. The main content area is titled 'Preferred Decompressors' and contains the following text:

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

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

Preferred Decompressors

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

Below the text is a large text input field. At the bottom of the configuration area are 'Submit' and 'Reset' buttons.

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

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

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

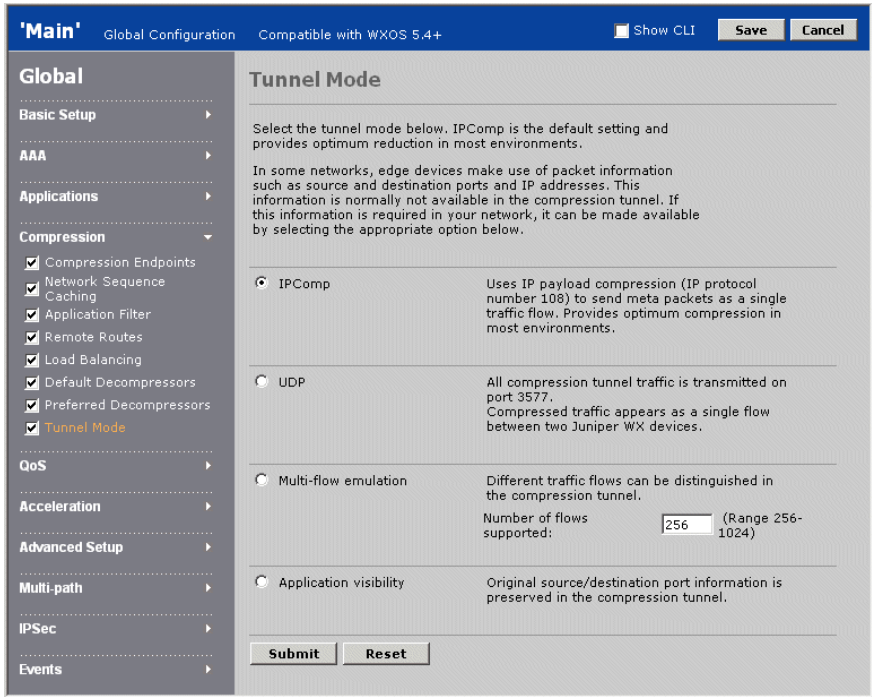
Configuring Tunnel Mode Settings

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

To configure the tunnel mode settings:

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

Figure 87: 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. |

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 tunnel. Your tools must be configured to monitor UDP traffic.



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

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

Configuring QoS Settings

The following sections describe the global outbound QoS parameters:

- “Using Outbound QoS to Enhance Performance” on page 169
- “Understanding Outbound QoS” on page 170
- “Procedure for Configuring Outbound QoS Policies” on page 180
- “Using the Outbound QoS Setup Wizard” on page 182
- “Defining Outbound QoS Settings by Endpoint” on page 191
- “Defining Outbound QoS Templates” on page 193
- “Defining Outbound QoS Endpoints” on page 195
- “Changing Outbound ToS/DSCP Values” on page 199
- “Starting and Stopping Outbound QoS” on page 202
- “Configuring Inbound QoS Policies” on page 203

Using Outbound QoS to Enhance Performance

Outbound QoS provides two key benefits:

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



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

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

To enable basic bandwidth allocation:

1. Specify the outbound WAN speed, as described in “Defining Outbound QoS Endpoints” on page 195. Adding the remote devices and specifying the WAN circuit speed for each device is also recommended.

For guidance on adjusting the WAN speeds to account for router overhead, refer to “WAN Circuit Speeds and Router Overhead” on page 173.

2. Start outbound QoS using Weighted Fair Queuing (WFQ) or Weighted Strict Priority (WSP), as described in “Starting and Stopping Outbound QoS” on page 202. Unless you need strict priority treatment for traffic classes, WFQ is recommended.

Understanding Outbound QoS

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

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

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



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

The following topics provide an overview of outbound QoS:

- “Traffic Classes and Bandwidths” on page 171
- “QoS Templates and Endpoints” on page 172
- “WAN Circuit Speeds and Router Overhead” on page 173
- “Dedicated, Oversubscribed, and Variable Rate WANs” on page 174
- “Direct Setup Versus Wizard Configuration Results” on page 176
- “Class Priorities and Excess Bandwidth Allocation” on page 178
- “ToS/DSCP Values” on page 180
- “Unadvertised Subnets” on page 180

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

Figure 88: 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 178).

Note that the Default class, which cannot be deleted, includes all undefined traffic. You must create an application definition for any traffic whose bandwidth you want to manage separately (refer to “Configuring Application Definitions” on page 145).

QoS Templates and Endpoints

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



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

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

- **Wizard-PrimeTime.** Applies to prime time hours, or to all hours if prime time is not defined. To specify the prime time, refer to “Defining the Prime Time” on page 241.
- **Wizard-NonPrimeTime.** Applies to non-prime time hours (if prime time hours are defined), and can be modified to allocate more bandwidth to applications that run during off-peak hours, such as database backups. To view bandwidth reports for prime time vs. non-prime time hours, use the device Web console.

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

WAN Circuit Speeds and Router Overhead

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

- **Outbound speed.** The sum of the WAN circuit speeds on the adjacent router that conduct traffic from the WX device. You must specify the outbound speed only if it is less than the sum of the remote WAN speeds—that is, if the WAN is “oversubscribed” (refer to “Dedicated, Oversubscribed, and Variable Rate WANs” on page 174).
- **Endpoint circuit speeds.** The maximum WAN circuit speed associated with each remote WX device for which you want to manage the outbound bandwidth. You can use the Ethernet speed for a remote WX device if you enable bandwidth detection for the remote 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, all WAN circuit speeds specified for outbound QoS must be set slightly lower than the WAN router’s full interface speed to allow for router overhead (Frame Relay LMI updates, CDP, SNMP, routing updates, and so on). Setting the bandwidth about 2 % below the link speed should work well in most cases. However, the router overhead is highly variable, and depends on the network configuration.

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.

Table 7: Recommended WAN Circuit Speed Adjustments

WAN Interface	Recommended QoS Speed	Description
Frame Relay	CIR minus 2 %	Reduce the Committed Information Rate (CIR) by 2 % . Higher speeds, up to the Peak Information Rate (PIR), may be acceptable, depending on the traffic load and whether "discard eligible" traffic is actually discarded. If the device exceeds the CIR, and discard eligible traffic is dropped, the QoS behavior may be unpredictable.
1.544 Mbps (T1)	1500 Kbps	The T1 line rate is 1.544 Mbps, but the data rate is 1.536 Mbps. The 8 Kbps difference is used for framing and encapsulation. Subtracting 2 % from 1.536 yields about 1.5 Mbps.

Table 7: Recommended WAN Circuit Speed Adjustments

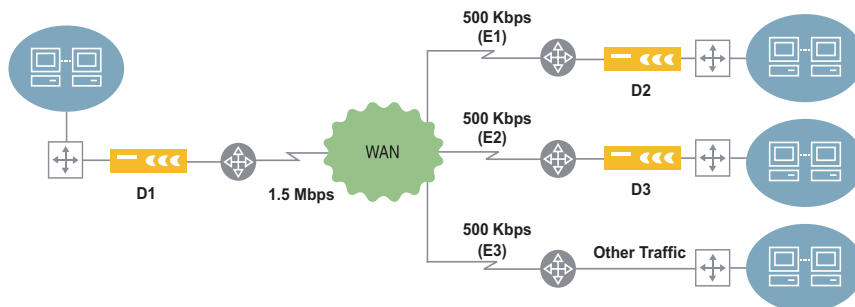
WAN Interface	Recommended QoS Speed	Description
512 Kbps (Fractional T1)	500 Kbps	Use one third of the T1 setting.
64 Kbps	60 Kbps	On low-speed links, router overhead may take up a greater percentage of the WAN link speed. Using 60 Kbps assumes that 6 % of the link is used for router control traffic.

Dedicated, Oversubscribed, and Variable Rate WANs

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

- **Dedicated.** The sum of the WAN circuit speeds on the adjacent router (the outbound speed) is equal to or greater than the sum of the remote WAN speeds. In this case, no adjustments to the bandwidth percentages are needed. In Figure 89, the aggregate WAN speed for device D1 (the outbound speed) is 1.5 Mbps, which equals the total speed of the three remote endpoints—D2, D3, and Other Traffic.

If D1 specifies a guaranteed bandwidth of 60 % for all traffic classes for each endpoint, the guaranteed capacity is 300 Kbps for D2, D3, and Other Traffic (.6 x 500 Kbps). However, in dedicated mode, Other Traffic is unconstrained by QoS.

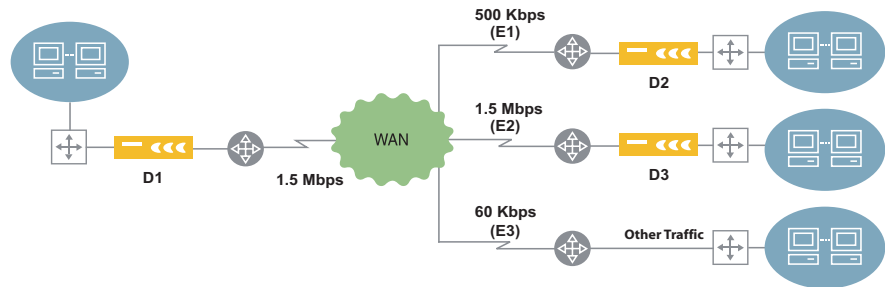
Figure 89: Dedicated WAN

If you want the Other Traffic to be managed by QoS, you can specify a dedicated WAN as “oversubscribed” or you can define virtual endpoints for the remote subnets you want to manage.

- **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 90, the WAN is oversubscribed from the perspective of device D1 because the outbound speed is 1.5 Mbps and the sum of the remote speeds is 2060 Mbps.

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

Figure 90: Oversubscribed WAN



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

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 173).



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 195).

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 8 shows the Wizard and direct setup results when D1 in Figure 90 is configured with two traffic classes and the same guaranteed bandwidths for each endpoint.

Table 8: 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 8 shows the guaranteed bandwidth in Kbps allocated to each traffic class on each endpoint. The guaranteed rate is calculated as follows:

(Remote Circuit Speed) * (Class Guaranteed Percentage)

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

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

Since the total guaranteed bandwidth (1133 Kbps) does not exceed 80 % of the D1 outbound speed ($.8 * 1500 = 1200$ Kbps), you can enter all the QoS settings directly without having to adjust the guaranteed percentages. Figure 91 shows the “Oversubscribed” template specifying the 15% and 40% guarantees, and Figure 92 shows the guaranteed bandwidths in Kbps displayed on the Outbound QoS Overview page when the template is applied to each endpoint.

Figure 91: 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 92: Direct Setup Results on the 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 8. The Wizard template guarantees are calculated as follows:

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

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

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

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

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

Note that the Wizard total guaranteed bandwidth (821 Kbps) is 55 % (15 % + 40 %) of the outbound speed (1500 Kbps) for D1. Figure 93 shows the guaranteed bandwidths in Kbps generated by the Setup Wizard and displayed on the Outbound QoS Overview page.

Figure 93: Wizard Results on the 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 90 on page 175, 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 9. The allocation depends on whether the WAN is dedicated or oversubscribed.

Table 9: 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 8 on page 176, the sum of the bandwidths is 1133 Kbps using direct setup or 821 Kbps using the Wizard.</p>

ToS/DSCP Values

The ToS/DSCP values in the packet headers can be set by traffic class for use by other devices in your network. You can also preserve the incoming ToS/DSCP values in the “meta-packets,” so that each meta-packet encapsulates only packets that have the same ToS/DSCP value. This allows other QoS devices in the path to manage the meta-packets in the same manner as the individual packets. For more information about setting ToS/DSCP values, refer to “Changing Outbound ToS/DSCP Values” on page 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 device should advertise all the subnets it can access. The source/destination filter can be used to prevent compression for specific destinations, as needed (refer to “Configuring Source/Destination Filters” on page 238).

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

Procedure for Configuring Outbound QoS Policies

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

1. For best results, verify that each 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 compression for specific destinations (refer to “Configuring Source/Destination Filters” on page 238).
2. Run the Setup Wizard or specify the outbound QoS policies directly:

- To run the Setup Wizard in CMS, refer to “Using the Outbound QoS Setup Wizard” in the next section). The Setup Wizard creates and applies the Wizard-PrimeTime and Wizard-NonPrimeTime templates to the selected endpoints.



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

- To specify the outbound QoS policies directly:
 - a. Specify the traffic classes and the applications in each class (refer to “Assigning Applications to Traffic Classes” on page 149).
 - b. Define one or more templates to specify the priorities and bandwidths for each traffic class (refer to “Defining Outbound QoS Templates” on page 193).
 - c. Specify the local outbound speed and the maximum circuit speeds for each remote endpoint (refer to “WAN Circuit Speeds and Router Overhead” on page 173 and “Defining Outbound QoS Endpoints” on page 195).
 - d. Assign a template to each endpoint (refer to “Defining Outbound QoS Settings by Endpoint” on page 191).
 - e. Enable QoS and select a queuing model (refer to “Starting and Stopping Outbound QoS” on page 202).
- 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 191).
 - Change traffic class names (refer to “Assigning Applications to Traffic Classes” on page 149).
 - Add new templates, change a template name, or change just one of the Wizard templates (refer to “Defining Outbound QoS Templates” on page 193)
 - Define virtual endpoints (refer to “Defining Outbound QoS Endpoints” on page 195).

- Change the ToS/DSCP values for one or more traffic classes (refer to “Changing Outbound ToS/DSCP Values” on page 199).
- Exclude address or subnet pairs from bandwidth management (refer to “Defining Outbound QoS Exclusions” on page 111).

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:

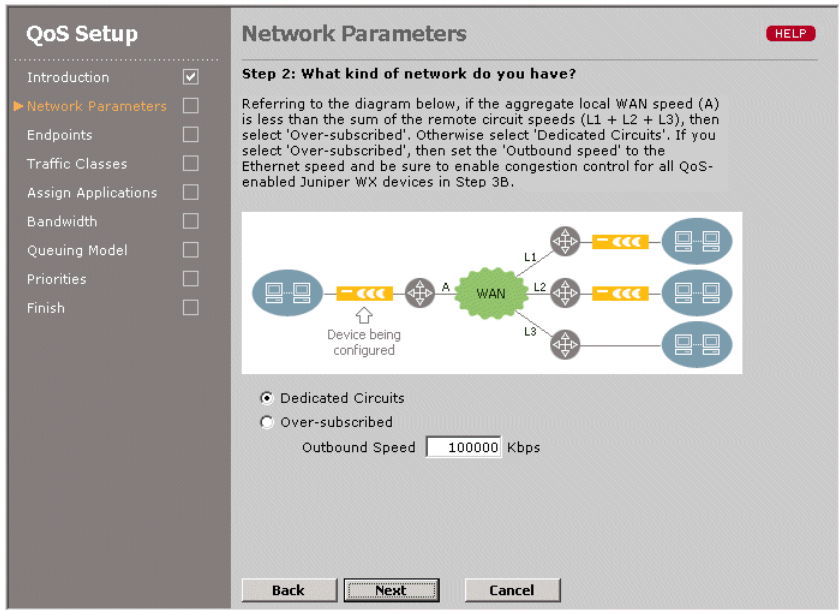
- **Wizard-PrimeTime.** Applies to the prime time hours (critical business hours). To specify the prime time, refer to “Defining the Prime Time” on page 241.
- **Wizard-NonPrimeTime.** Applies to nonprime time hours. To view QoS reports for prime time or nonprime time hours, use the device Web console.

Each time you run the Setup Wizard, both of the templates and all customized settings are overwritten. To change just one of the templates, refer to “Defining Outbound QoS Templates” on page 193.

To run the outbound QoS Setup Wizard:

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

Figure 94: Configuring Outbound QoS Network Parameters



3. Specify the WAN mode for the device(s) where you intend to load the configuration, 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. |
| 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 173). |

Figure 95: Configuring QoS Endpoints

QoS Setup

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

Endpoints

Step 3: Which endpoints do you want to manage bandwidth?

Select devices from the list on the left. Then click the **Add** button to add the selected devices to the QoS Endpoints list. Hold down the **Ctrl** key on your keyboard to make multiple selections. To add devices that do not appear in a community/device group list, click the **Manual Entry** button.

Community/Device Group
 Inline

Devices in Community/Device Group

SR-192.168.0.195 (192.168.0.195)
SR-192.168.0.195 (192.168.0.244)
SR-192.168.5.131 (SR-50A) (192.168.5.131)
SR-192.168.5.130 (192.168.5.132)
10.250.1.97 (10.250.1.97)
192.168.0.223 (192.168.0.223)
5.200 -> 0.x (192.168.5.200)

QoS Endpoints

SR-192.168.0.195 (192.168.0.195)

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

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable QoS for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 116.

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



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

e. When you are done, click Next.

Figure 96: Configuring Endpoint Circuit Speeds

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. If you need to add or remove Juniper WX endpoints, click the Back button.

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

Endpoints enabled for acceleration cannot be disabled for QoS.

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



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 173). Exceeding the actual WAN speed effectively shifts bandwidth management to the router, and may cause the router to drop traffic.

Note the following:

- 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 Juniper WX” virtual endpoints have been defined to manage the “Other traffic” sent to specific remote subnets (refer to “Defining Outbound QoS Endpoints” on page 195), you can change their circuit speeds or disable them. The settings for “Other traffic” and virtual endpoints can be changed in the same manner as other endpoints (refer to “Defining Outbound QoS Settings by Endpoint” on page 191).

Click Next.

Figure 97: Configuring Bandwidth Detection

The screenshot shows the 'QoS Setup' wizard with a sidebar on the left containing the following steps: Introduction (checked), Network Parameters (checked), Endpoints (selected), Traffic Classes, Assign Applications, Bandwidth, Queuing Model, Priorities, and Finish. The main panel is titled 'Endpoints' and contains the following text:

Step 3B: Do your circuit speeds vary?

For some endpoints, actual circuit speeds may vary. In order to optimize Bandwidth Management for these endpoints, you should enable 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 reduced traffic to:

☐ All QoS-enabled Juniper WX devices that also have reduction enabled

☒ ONLY checked Juniper WX devices below

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

Below the table are 'Select All' and 'Clear' buttons. At the bottom are 'Back', 'Next', and 'Cancel' buttons.

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. Also, if you entered the device's Ethernet speed as the outbound speed, enable bandwidth detection for all endpoints.

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 Juniper WX devices.** Applies bandwidth detection to all remote devices for which QoS is enabled (default).
 - **ONLY checked Juniper 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 tunnels are enabled. In oversubscribed mode, if you have substantial passthrough traffic for other destinations, you may want to reduce the maximum speed for the “Other traffic” endpoint to limit the bandwidth allocated to passthrough traffic. After you complete the Wizard configuration, refer to “Defining Outbound QoS Endpoints” on page 195.

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

Figure 98: Configuring Traffic Classes

QoS Setup

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

Traffic Classes HELP

Step 4: Create Traffic Classes

You can create up to 15 traffic classes. Enter the traffic class name below and click 'Add'. Click 'Delete' to delete checked traffic classes.

Traffic Class

Default

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

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

Figure 99: Assigning Applications to Traffic Classes

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
CVS	Default
Clearcase	Default
DNS	Default
Exchange	Default
FTP	Default
Filenet	Default
Groupwise	Default

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

9. Select the appropriate traffic class for each application. If one of your network applications is not shown, you must create an application definition for it, as described in “Configuring Application Settings” on page 141. Click Next.

Figure 100: Defining Guaranteed and Maximum Bandwidths

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	Percentage of the bandwidth that is guaranteed to be allocated to the applications in the traffic class. Lower values indicate that the traffic in the class is more likely to be delayed. Traffic may be dropped when the guaranteed bandwidth is exceeded, such as during a burst of higher-priority traffic. The total guaranteed bandwidth across all traffic classes cannot exceed 80 %. Also, the total guaranteed bandwidth across all endpoints cannot exceed 80 % of the local outbound WAN speed.
Maximum Bandwidth	Maximum percentage of the bandwidth that can be allocated to the applications in the traffic class. Traffic is dropped when the maximum bandwidth is exceeded. A zero indicates that all traffic in the class is dropped.



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

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

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

You can later change the queuing method, as described in “Starting and Stopping Outbound QoS” on page 202. Click Next.

Figure 101: Defining Priorities by Traffic Class

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 available (unguaranteed) bandwidth to the competing traffic classes. These priorities are used only by the WX device, and are not related to ToS priorities.
13. Click Next, click Submit, and then click Close.

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

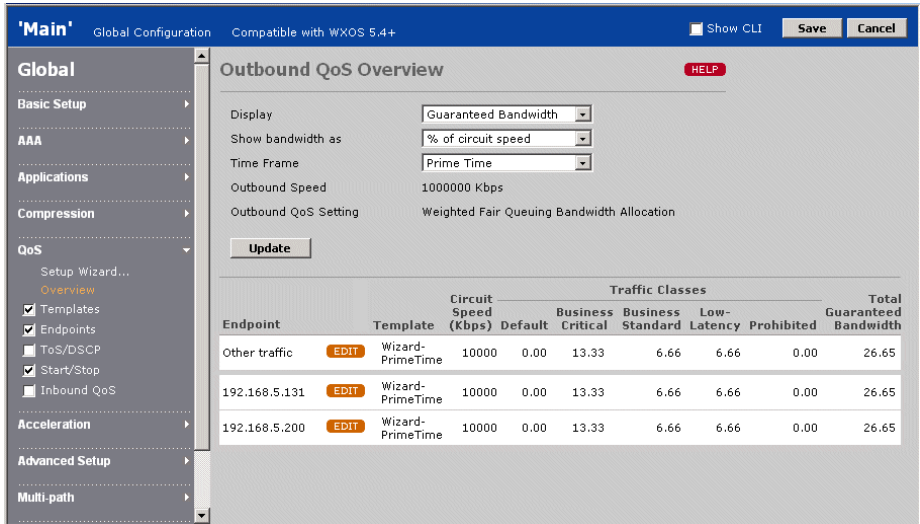
Defining Outbound QoS Settings by Endpoint

After you run the Setup Wizard to create the initial outbound QoS settings, you can manually change the prime-time or nonprime-time template assigned to each endpoint or override the template values (class priorities or bandwidths) for a single endpoint. To change the WAN circuit speed for an endpoint, refer to “Defining Outbound QoS Endpoints” on page 195.

To view or change the outbound QoS settings by endpoint:

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

Figure 102: Outbound QoS Overview



The Outbound QoS Overview page shows the outbound speed and 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 view the bandwidth percentages as circuit speeds.



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 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 241).
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 103: Changing Endpoint Templates or Template Settings

'Main' Global Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Global

- Basic Setup
- AAA
- Applications
- Compression
- QoS
 - Setup Wizard...
 - Overview
 - ☒ Templates
 - ☒ Endpoints
 - ☒ ToS/DSCP
 - ☒ Start/Stop
 - ☐ Inbound QoS
- Acceleration
- Advanced Setup
- Multi-path
- IPSec

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
Time Frame: Prime Time

☐ Use QoS template
☒ 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 %	100.00 %
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 193.

- To override the current template settings for this endpoint, click Use custom setting and change the priority or bandwidth settings for one or more traffic classes, and click Submit.

Note that to increase the guaranteed bandwidth for a traffic class on an oversubscribed WAN, you must first decrease the bandwidth on another class (on the same endpoint or a different endpoint), reduce the circuit speed, or increase the outbound speed. The Setup Wizard adjusts the guaranteed bandwidths for you (refer to “Using the Outbound QoS Setup Wizard” on page 182).

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

The Outbound QoS Overview page is displayed. When you override the template settings for an endpoint, the template name is changed to None. You can later reapply the template to restore the original settings.

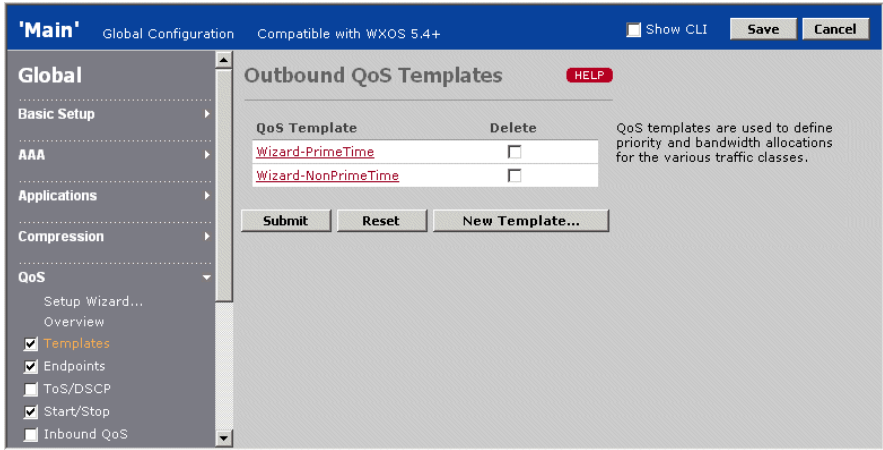
Defining Outbound QoS Templates

Templates specify the priority, and guaranteed and maximum bandwidths for each traffic class. You can change the templates created by the Setup Wizard or create new templates. To apply a template to an endpoint, refer to “Defining Outbound QoS Settings by Endpoint” on page 191.

To define outbound QoS templates:

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

Figure 104: Defining Outbound QoS Templates



From the Outbound QoS Templates page, you can:

- Add a new template, as described in Step 2 through Step 3.
- Change a template name or settings. Click the template name, change the template name and/or the settings for each traffic class, and click Submit.
- Delete a template. Click the check box next to the template name, and click Submit. If the template is applied to an endpoint, all priority and guaranteed bandwidth values are set to zero for that endpoint. Maximum bandwidth values are set to 100 %.

2. To add a new template, click New Template and enter the following information:

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



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

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

Defining Outbound QoS Endpoints

Each device can manage the outbound bandwidth for one or more remote WX devices or other (virtual) endpoints. After you run the Setup Wizard, you can:

- Specify the local WAN as dedicated or oversubscribed
- Add or remove endpoints for bandwidth management.
- Define virtual endpoints to manage the traffic to specific remote subnets that do not have a WX device.
- Change the remote WAN circuit speeds.
- Enable bandwidth detection for one or more endpoints.

To exclude specific LAN/WAN address or subnet pairs from bandwidth management, refer to “Defining Outbound QoS Exclusions” on page 111.

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 182).

To define the outbound QoS endpoints:

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

Figure 105: Enabling Bandwidth Management by Endpoint

MainGlobal ConfigurationCompatible with WXOS 5.4+Show CLISaveCancel

Global

Basic SetupAAAApplicationsCompressionQoSSetup Wizard...OverviewTemplatesEndpointsToS/DSCPStart/StopInbound QoSAccelerationAdvanced SetupMulti-pathIPSecEvents

Outbound QoS EndpointsHELP

You can enable Outbound QoS by using the Setup Wizard or from the 'Start/Stop' page. If Outbound QoS is enabled, then only outbound traffic destined for the checked endpoints below will be affected.

Juniper WX devices are automatically included in the list below. If you want to enable QoS to endpoints that are NOT reachable through a Juniper 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 Juniper WX device, click this link.

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.)

Dedicated Circuits

Over-subscribed

Outbound Speed1000000kbps

Enable Bandwidth Detection when sending traffic to:

All QoS-enabled Juniper WX devices

ONLY Juniper WX devices checked under "Bandwidth Detection"

Endpoint	Name	Circuit Speed (Kbps)	Congestion Control	Min. Speed (Kbps)
Other traffic		default	ADD	
<input checked="" type="checkbox"/> No Remote Peribit	Non-SR-Chicago	500	DELETE	
<input checked="" type="checkbox"/> 192.168.207.200	192.168.207.200	0	<input checked="" type="checkbox"/>	512
<input type="checkbox"/> 192.168.206.200	SR-192.168.206.200	0	<input type="checkbox"/>	

Add/Remove Endpoints

SubmitReset

2. Specify the WAN mode for the device(s) where you intend to load the configuration:

- 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 add virtual endpoints (see Step 6) or select Oversubscribed and use the default outbound speed.
- Over-subscribed

Indicates that the local outbound WAN speed is less than the sum of the remote WAN speeds. Add up the speeds of all the WAN interfaces on the adjacent router, and enter the total in the Outbound Speed field (in Kbps). Be sure to account for router overhead (refer to "WAN Circuit Speeds and Router Overhead" on page 173).

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 191).

3. To change the maximum circuit speed associated with each remote endpoint, enter the new values (in Kbps) and click Submit. For a description of the outbound WAN speed, refer to “Dedicated, Oversubscribed, and Variable Rate WANs” on page 174. The outbound and remote circuit speeds are required.



NOTE: If bandwidth detection is not enabled (see Step 5), be sure to adjust the WAN speed to account for router overhead (refer to “WAN Circuit Speeds and Router Overhead” on page 173).

4. To add or remove remote endpoints for outbound QoS:

- a. Click Add/Remove Endpoints.
- b. Select a community from the Community/Device Group list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable QoS for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 116.

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



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

- g. Enter the remote WAN circuit speed (in Kbps) for each endpoint that you added, and click Submit.

When you add a new endpoint, all the endpoint's traffic classes have a priority and guaranteed bandwidth of zero, and a maximum bandwidth of 100%. To change the default settings, refer to "Defining Outbound QoS Settings by Endpoint" on page 191.

5. 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. Also, if you entered the device's Ethernet speed as the outbound speed, enable bandwidth detection for all endpoints.

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 Juniper WX devices.** Applies bandwidth detection to all remote devices for which QoS is enabled (default).
 - **ONLY Juniper 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 the compressed traffic sent to other WX endpoints. In oversubscribed mode, if you have substantial passthrough traffic for other destinations, you may want to reduce the maximum speed for the "Other traffic" and virtual endpoints to limit the bandwidth allocated to passthrough traffic.

6. 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. To view the subnets associated with the current virtual endpoints, click **view a list of remote networks...**

To add a virtual endpoint, click ADD, specify the following information, and click Submit. The maximum number of virtual endpoints depends on the device type (2 for the WX 15, 5 for the WX 20 and WXC 250, 60 for the WXC 500, 120 for the WX 50 and WX 60, and 320 for the WX 100).

Name	Enter the endpoint name (up to 20 characters).
Circuit Speed	Enter the 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 change a virtual endpoint’s name or subnets, click the endpoint name, make the changes, and click Submit. To delete a virtual endpoint, click DELETE next to the endpoint. Traffic to deleted virtual endpoints is managed by the “Other-traffic” endpoint.

Changing Outbound ToS/DSCP Values

The ToS/DSCP values on incoming traffic from the LAN can be modified to support other QoS devices in your network. For each traffic class, you can specify a Type of Service (ToS) IP precedence value or a Differentiated Services Code Point (DSCP) value, depending on the QoS scheme in use. The specified ToS/DSCP values apply to all traffic in the class, regardless of whether the traffic is compressed or outbound QoS is enabled.

You can also preserve the incoming ToS/DSCP values in the “meta-packets,” so that each meta-packet encapsulates only packets that have the same ToS/DSCP value. This allows other QoS devices in the path to manage the meta-packets in the same manner as individual packets. By default, meta-packets have a ToS/DSCP value of zero and can encapsulate packets with varying ToS/DSCP values.

ToS IP precedence values (0 to 7) use the upper three bits of the Diffserv field; DSCP values (0 to 63) use the upper six bits. The upper three bits of DSCP are used like ToS to indicate the priority (7 is the highest priority). Table 10 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.

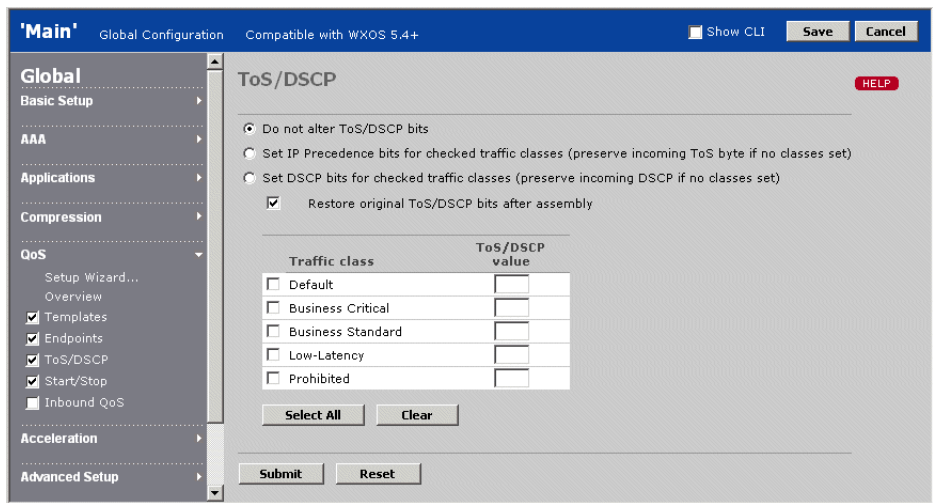
Table 10: ToS and DSCP Values

Name	DSCP	IP Precedence
Default or BE (best effort)	0	0
CS1	8	1
CS2	16	2
CS3	24	3
CS4	32	4
CS5	40	5
CS6	48	6
CS7	56	7
AF11	10	–
AF12	12	–
AF13	14	–
AF21	18	–
AF22	20	–
AF23	22	–
AF31	26	–
AF32	28	–
AF33	30	–
AF41	34	–
AF42	36	–
AF43	38	–
EF	46	–

To set ToS/DSCP values by traffic class:

1. In the Configuration window, click QoS in the left-hand navigation frame, click ToS/DSCP, and select the check box.

Figure 106: Setting ToS/DSCP Values



2. To set ToS/DSCP values by traffic class, select Set IP Precedence bits... or Set DSCP bits... to specify whether you want to enter ToS or DSCP values.

The default selection, Do not alter ToS/DSCP bits, indicates that meta-packets have a ToS/DSCP value of zero. If you want to preserve all the incoming values, and have each meta-packet reflect the ToS/DSCP value of its encapsulated packets, select Set IP Precedence bits... or Set DSCP bits... and do not check any of the traffic classes.

3. Select the check boxes next to the traffic classes whose ToS/DSCP values you want to set (or click Select All).

4. Enter a ToS value (0 to 7) or a DSCP value (0 to 63) in the ToS/DSCP value field for each of the selected classes. The value specified for each class is applied to the traffic for all applications in the selected class. To assign applications to a traffic class, refer to “Assigning Applications to Traffic Classes” on page 149.



NOTE: These ToS/DSCP values are overridden by the ToS/DSCP settings defined for Multi-Path (refer to “Enabling Policy-Based Multi-Path” on page 265), and by ToS values set for router-based balancing (refer to the “configure route” CLI command).

5. After compressed traffic from remote devices is decompressed, the **Restore original ToS/DSCP bits after decompression** option resets the ToS/DSCP value to its original value (if the remote device changed it).
6. Click Submit to enter the changes, or click Reset to discard them.

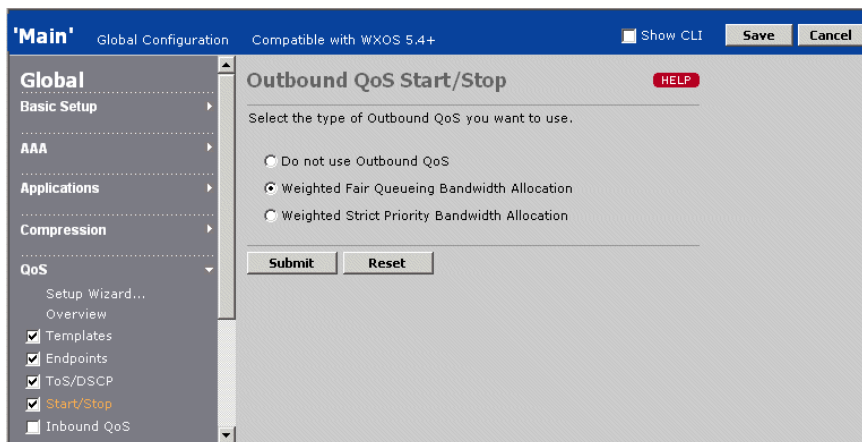
Starting and Stopping Outbound QoS

You can start or stop outbound QoS or change the queuing method without using the Setup Wizard. The queuing method determines how the available (unguaranteed) bandwidth is allocated among the contending applications. The selected queuing model applies to all the managed endpoints.

To stop the outbound QoS service or change the queuing model:

1. In the Configuration window, click QoS in the left-hand navigation frame, click Start/Stop, and select the check box.

Figure 107: Starting and Stopping Outbound QoS



2. To stop the outbound QoS service, click Do not use Outbound QoS.
3. To restart the service or change the queuing method used for each endpoint, select one of the following:
 - **Weighted Fair Queuing Bandwidth Allocation.** Queues are created for each traffic class, and the available bandwidth is allocated by processing the queues based on their priority and guaranteed bandwidth.
 - **Weighted Strict Priority Bandwidth Allocation.** Queues are created for each priority, and the available bandwidth is allocated by processing the queues based only on priority.



NOTE: When you save a global configuration, an error occurs if PFA is enabled and outbound QoS is off.

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

Configuring Inbound QoS Policies

You can specify the maximum bandwidths for four classes of incoming WAN traffic destined for the Local Area Network (LAN). Setting maximum bandwidths for each class ensures that low-priority traffic, such as Web traffic, does not interfere with mission-critical applications. Bandwidths are specified as percentages of the inbound WAN speed, and traffic that exceeds the maximum bandwidths is dropped.



NOTE: Inbound QoS applies only to traffic received on the Remote interface. Off-path devices use only the Local interface. In hierarchical deployments where both the Local and Remote interfaces are connected to a WAN router, inbound QoS has no effect on incoming WAN traffic on the Local interface.

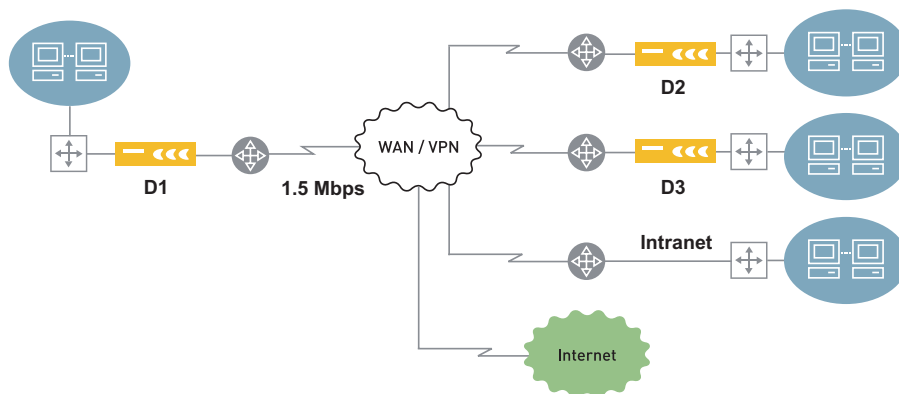
The following table describes the four traffic classes for inbound bandwidth management.

Table 11: 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 Top Traffic report to help create the list of subnets (refer to “Top Traffic Statistics” on page 355).
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 108, set the inbound speed to 1500 Kbps (1.5 Mbps). You then set maximum bandwidth percentages for one or more of the four traffic classes. In this example, you might set the maximum bandwidth percentage for the Default class to 10 % to limit low-priority traffic from the public Internet.

Figure 108: Configuring Inbound Bandwidth Management



To configure inbound QoS:

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

Figure 109: Configuring Maximum Inbound QoS Bandwidths

'Main' Global Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Global

- Basic Setup
- AAA
- Applications
- Compression
- QoS**
 - Setup Wizard...
 - Overview
 - Templates
 - Endpoints
 - ToS/DSCP
 - Start/Stop
 - Inbound QoS**
- Acceleration

Inbound QoS

If 'Enable Inbound QoS' is checked, traffic from the following four predefined traffic classes will be limited to the specified maximum bandwidths.

☐ Enable Inbound QoS

Inbound Speed Kbps

Traffic Class	Maximum Bandwidth	Description
Reduced	100 %	Any traffic that has been compressed by a WX device.
Intranet	100 %	TCP traffic originating from the corporate network that has NOT been compressed.
TCP	100 %	TCP traffic NOT originating from the corporate network.
Default	100 %	All other protocols (e.g. UDP, streaming)

Submit Reset

2. To start the inbound QoS service, click Enable Inbound QoS.
3. Add up the speeds of all the WAN interfaces on the router that conducts WAN traffic to the device where you intend to load the configuration, and enter the value (in Kbps) in the Inbound Speed field.
4. Enter the maximum bandwidth of each traffic class as a percentage of the inbound speed.
5. Click Submit to enter the changes, or click Reset to discard them.
6. To specify the remote subnets whose traffic belongs to the Intranet class, click Intranet and enter the remote subnets (one per line) whose traffic belongs to the Intranet traffic class, and click Submit. The subnet format is:

<IP address>/<subnet mask>

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

Configuring Traffic Acceleration

The following topics describe how to configure traffic acceleration:

- “Overview of Packet Flow Acceleration” in the next section
- “Overview of Application Flow Acceleration” on page 210
- “Enabling Acceleration by Endpoint” on page 214
- “Enabling Acceleration by Application” on page 219

Overview of Packet Flow Acceleration

While compression effectively increases the available WAN bandwidth, application performance may be further constrained by network latency. Packet Flow Acceleration (PFA) improves the performance of compressed TCP application flows across high-speed, high-latency WAN links. For devices that support Multi-Path, you can enable PFA for the primary and/or secondary paths.



NOTE: PFA is most effective in networks with high-speed connections and high latency. However, PFA may have no effect if the traffic must cross low-speed or high-latency connections that are one or more hops beyond the receiving WX device.

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.



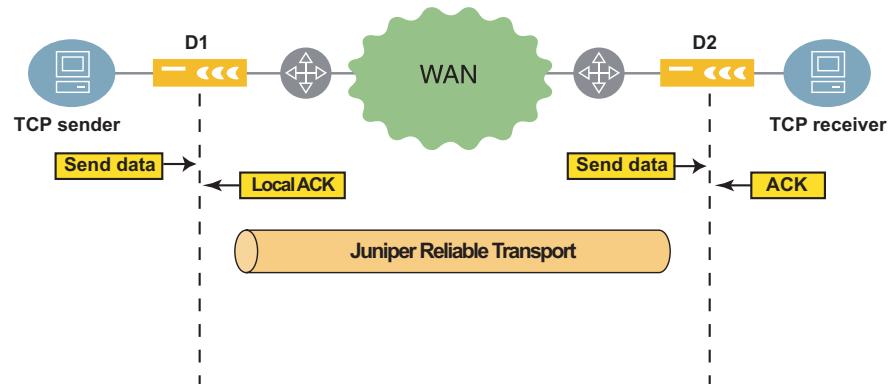
NOTE: TCP Acceleration is required to use Network Sequence Caching (NSC) on WXC devices, or to accelerate Microsoft CIFS, Microsoft Exchange, and HTTP traffic using Application Flow Acceleration.

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 device, between the two WX devices, and between the receiving device and the destination.

Since all transmissions are acknowledged locally, more data can be put “in flight” at once. The ACKs are returned to the sender at a rate governed by the speed of the link.

To avoid the TCP congestion mechanism, which is very inefficient over the WAN, a reliable transport protocol ensures in-order delivery between the two WX devices, and provides retransmission when necessary. Congestion is managed by outbound QoS.

Figure 110: 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 50 % 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 flow in both directions must be handled by the same two WX devices. In a load-balancing environment, the two TCP setup packets for a new flow (SYN and SYN ACK) may be seen by different devices. In this case, you can define clusters of devices that advertise their SYN packets so that any device in the cluster that sees the SYN ACK can establish the flow to the sending device.

For more information about using asymmetric routing support, refer to the *WX/WXC Operator's Guide*.

Forward Error Correction

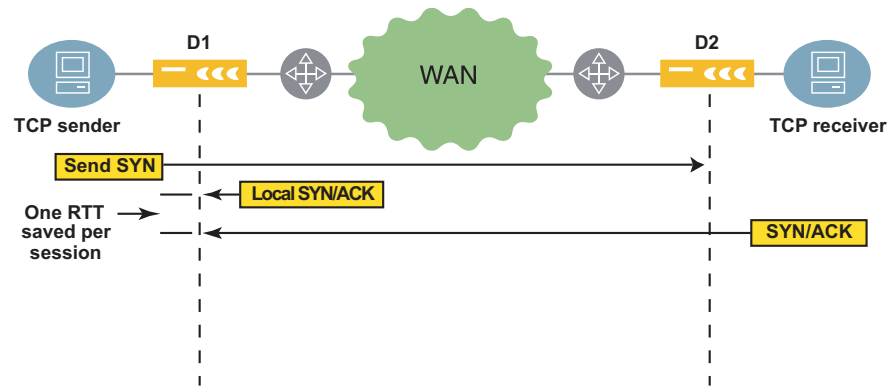
Forward Error Correction (FEC) enables the sending device to send recovery packets along with all data packets, so that the receiving device can reconstruct lost packets without requesting a retransmission. You can specify the number of recovery packets per block of data packets.

FEC is intended for use in high-loss, high-latency environments, such as satellite connections. However, FEC should be disabled if the satellite modem also provides forward error correction. Note that when FEC is enabled for a device, recovery packets are generated for all traffic sent to that device.

After you enable FEC, check the monitoring report periodically. If losses are not persistent, disable FEC to avoid the overhead of processing recovery packets.

Fast Connection Setup

With Fast Connection Setup (FCS), the sending device locally acknowledges the initial session request (the SYN packet) for each new TCP session if the destination is known to be active. FCS saves one round-trip time (RTT) for each session, and is intended for applications that have many short sessions, such as HTTP 1.0. Short sessions are those that last less than ten times the round-trip time.

Figure 111: Fast Connection Setup

FCS is particularly useful in pre-Windows 2000 environments, where NetBios (not CIFS) is used for file transfer. FCS is also beneficial for HTTP 1.0 traffic (pre-Windows 2000) as it creates more short-lived TCP connections than HTTP 1.1. Some custom enterprise WAN applications may also benefit from FCS.

FCS is most effective in high latency environments, because each RTT that is saved per session represents a larger slice of time as the latency increases. If latency is very low (LAN latencies for example), FCS will not provide much benefit.

FCS is applied only to sessions that last less than ten times the round-trip time (RTT). If a specific application traffic flow has five consecutive short sessions, FCS is applied to all subsequent identical traffic flows. The average session acceleration is calculated as follows:

$$100 - [100 (\text{Accelerated session time}) / (\text{Session time without acceleration})]$$

Note that performance improvements will be more noticeable to users as the percentage of accelerated sessions increases.

Overview of Application Flow Acceleration

Though technologies such as compression (MSR and NSC) and TCP Acceleration can greatly increase the performance for applications across the WAN, these benefits may be undermined by inefficient protocols above TCP. To achieve the best end-user performance, specific protocols need to be optimized for the WAN.

The primary purpose for Application Flow Acceleration (AppFlow) is to improve end-user performance for specific business-critical protocols that traverse the WAN. Application Flow Acceleration not only improves performance for existing WAN applications but also facilitates the centralization of branch servers to central data centers.

Currently three business-critical, but WAN-inefficient protocols are accelerated: Microsoft Common Internet File System (CIFS), which is the underlying protocol for Microsoft File Services, traffic between Microsoft Exchange servers and Outlook clients (MAPI over RPC), and Web traffic (HTTP).

If TCP Acceleration is enabled for one or more remote endpoints, you can enable application-level acceleration for Microsoft CIFS, Microsoft Exchange, and HTTP traffic sent to those endpoints. You can accelerate all such traffic, or you can create application definitions that let you accelerate traffic to specific servers. Application Flow Acceleration must be enabled on the devices closest to the clients.



NOTE: AppFlow and tunnel switching cannot be enabled on the same device. When AppFlow is enabled, an error occurs if the tunnel switching CLI commands are added to the CLI section of the configuration, or if you load the configuration on a device that has tunnel switching enabled.

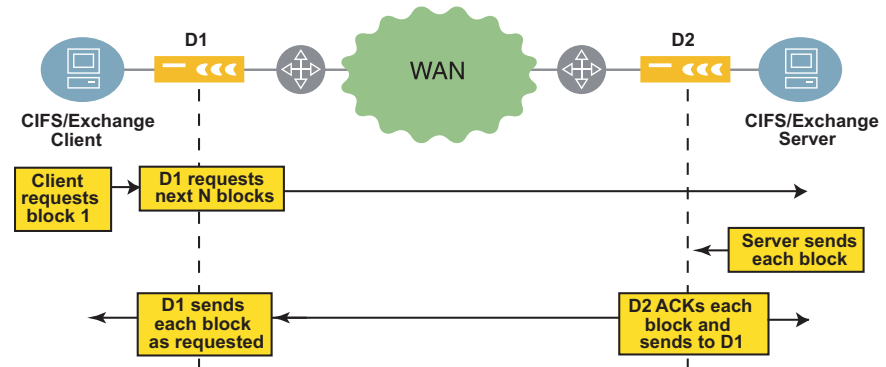
Microsoft CIFS and Microsoft Exchange Acceleration

Microsoft CIFS and Microsoft Exchange traffic is accelerated by having the WX device locally acknowledge each block of traffic sent during bulk read/write operations, such as copying files (for CIFS) and sending or receiving Emails with attachments. This allows many data blocks to be in flight at the same time, which speeds up the data transfer. Acceleration benefits begin at relatively low latencies (about 30 ms. round-trip time).

CIFS and Exchange are TCP protocols that transfer bulk data (files or attachments) by breaking up the object into smaller data blocks. CIFS and Exchange write or read one block of data at a time before proceeding to the next block. This serial transmission of small data blocks is a major contributor to slow performance over the WAN.

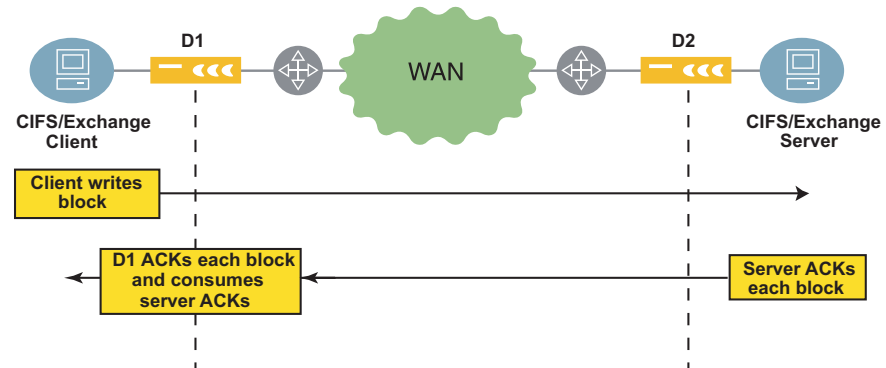
In read operations (Figure 112), the client requests one block of data at a time. The device closest to the client (D1) requests the next N blocks. The device closest to the server (D2) locally acknowledges each block from the server and sends them to D1. D1 serves each block to the client as requested.

Figure 112: Microsoft CIFS/Exchange Read Operations



In write operations (Figure 113), the client writes one block at a time. The device closest to the client (D1) acknowledges each block locally, and discards the acknowledgements from the server.

Figure 113: Microsoft CIFS/Exchange Write Operations



NOTE: CIFS acceleration is not effective if Server Message Block (SMB) signing is enabled. Signing should be disabled on all servers and clients, and on all domain controllers that are also used as file servers (refer to “Enabling Microsoft CIFS Acceleration” on page 222). For more information about SMB signing, go to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;887429>

HTTP Acceleration

HTTP acceleration should be enabled only in environments where:

- Network latency exceeds 100 ms.
- Visited Web pages have many static objects, and cookies are not used

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:

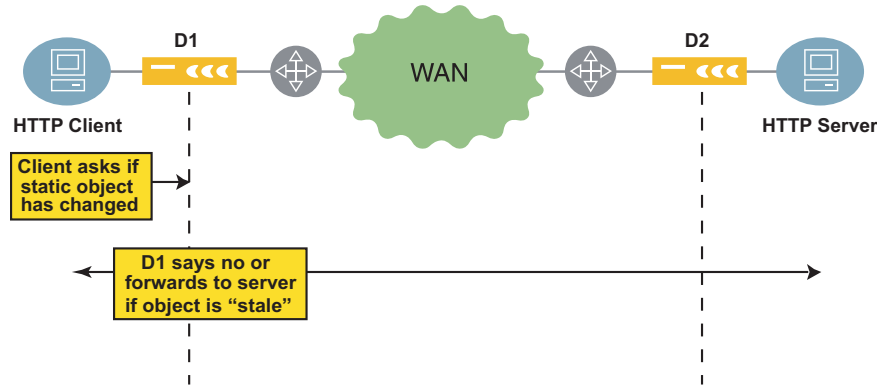
- If the object is fresh, a 304 NOT MODIFIED is sent, which saves a round-trip time.
- If the object is not fresh, the request is forwarded to the originating HTTP server.

In header-and-body mode, a WXC locally responds to both GET and GET IF-MODIFIED-SINCE requests for the static objects in its cache. Serving cached objects saves at least one round-trip time for each object.

- **Pre-fetching.** After a page is requested once (by any client), a request for the first static object on a page triggers requests for all the page's static objects, which saves one round-trip time for each pre-fetched object. On WXC devices, only objects that are considered "stale" by the cache are pre-fetched.

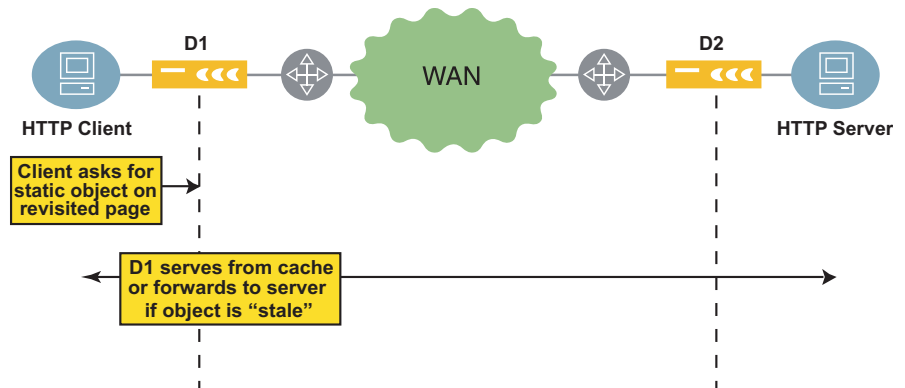
In HTTP cache header-only mode (Figure 114), the client sends HTTP GET IF-MODIFIED-SINCE queries before reloading an object from the browser cache. Based on its own caching timer, the device closest to the client (D1) indicates the object has not changed or forwards the query to the server. Even if the query is not forwarded, D1 sends its own query to verify the object's last-modified date.

Figure 114: HTTP Caching—Header-Only Mode



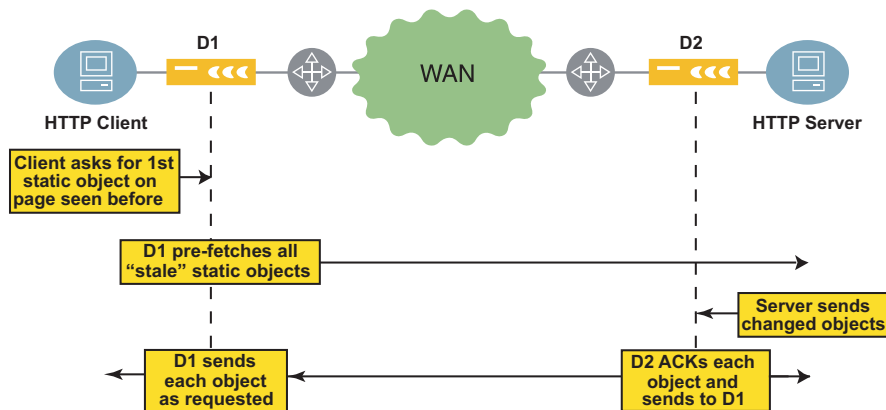
In HTTP cache header-and-body mode (Figure 115), the client sends HTTP GET requests for static objects on a page that has been visited before. The WXC device closest to the client (D1) serves the objects directly from its own cache (if they are still fresh) or forwards the requests to the HTTP server.

Figure 115: HTTP Caching—Header-and-Body Mode (WXC Devices)



If pre-fetch is enabled (Figure 116), the static objects associated with each page (".css", ".gif", ".jpeg", and ".js") are recorded when the page is first requested. When the first object of a previously seen page is requested again, the device (D1) requests all the static objects that are considered stale. The objects returned by the server are acknowledged locally by D2.

Figure 116: HTTP Pre-Fetch



NOTE: HTTP traffic is not accelerated if a proxy server exists between the server-side device and the actual HTTP server. However, if the proxy server is between the Web client and the client-side device, HTTP traffic will be accelerated.

Enabling Acceleration by Endpoint

You can enable each method of Packet Flow Acceleration for all remote devices (endpoints), or for specific endpoints. 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 tunnels in both directions between the devices (refer to "Configuring Endpoints for Compression" on page 152).
- Enable compression for the applications you want to accelerate (refer to "Compressing Applications" on page 156).

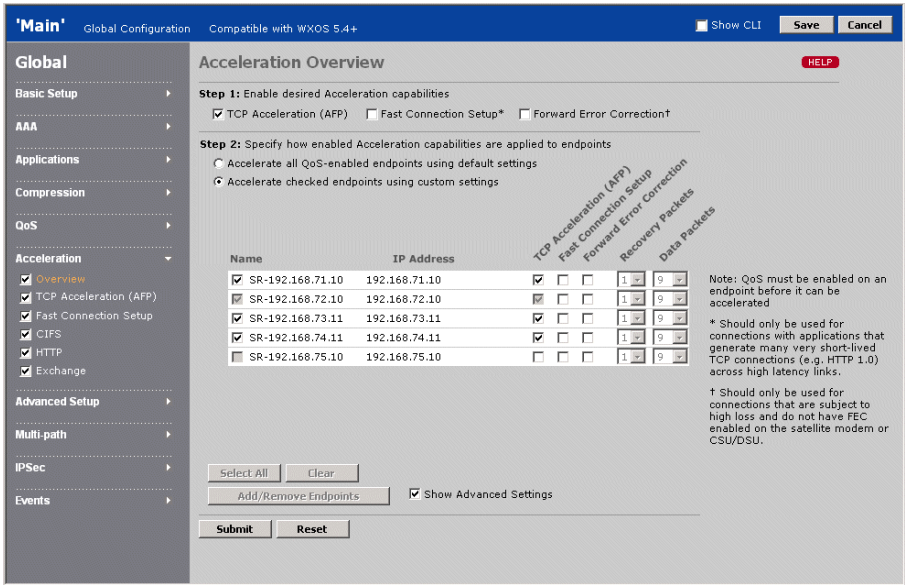
- To enable TCP Acceleration in WXOS 5.3/5.2 configurations, enable TCP Acceleration on the Features/Topology page in the Advanced Setup section.
- Enable outbound QoS, and specify the WAN circuit speed for each remote device for which you want to accelerate traffic (refer to “Using Outbound QoS to Enhance Performance” on page 169).

If you enable TCP Acceleration or Fast Connection Setup, you must also select the applications that each method is applied to (refer to “Enabling Acceleration by Application” on page 219).

To enable Packet Flow Acceleration by endpoint:

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

Figure 117: Enabling Packet Flow Acceleration



2. At the top of the page, select the check box next to each of the PFA methods that you want to use for one or more of the remote endpoints.
3. Select one of the following options:

- **Accelerate all QoS enabled endpoints using default settings.** Traffic is accelerated to all remote devices for which a tunnel exists and outbound QoS is configured correctly. The PFA methods you select apply to all qualifying endpoints, and to all qualifying endpoints added to the same community in the future.
- **Accelerate checked endpoints using custom settings.** Traffic is accelerated only to the selected devices, and different PFA methods can be used for each endpoint. Click the check box next to the IP address of the appropriate devices. An endpoint is greyed out if QoS is not enabled for the device.

To add or remove specific remote endpoints for PFA:

- a. Click Add/Remove Endpoints.
- b. Select a community from the Community/Device Group list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable PFA for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 116.

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



NOTE: When you save a global configuration, an error occurs if QoS and compression are not enabled for all endpoints using PFA. If you enable PFA for an endpoint in an Acceleration partial configuration, an error occurs if you load the configuration on a device where QoS or compression is not enabled for that endpoint.

4. Select the PFA 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 the “configure acceleration” CLI command in the operator’s guide):</p> <ul style="list-style-type: none"> ■ Adjust the buffer size for optimum performance. ■ Increase the number of lost heartbeat packets allowed on high-loss links (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 160) ■ For device speeds of 20 Mbps or more, enable fast compression tunnels for greater throughput if acceleration is more important than compression (refer to the “config reduction set fast-reduction-tunnel” CLI command).
Fast Connection Setup	<p>Intended for applications that have many short sessions, such as HTTP 1.0 and NetBios. The sending device locally acknowledges session requests for destinations known to be active. Short sessions are those that last less than ten times the round-trip time (RTT).</p>
Forward Error Correction	<p>Intended for high-loss environments. The sending device sends recovery packets with the data to reduce the number of retransmissions required when data packets are lost. By default, one recovery packet is sent for every nine data packets. To change the number of data and recovery packets, click Show Advanced Settings at the bottom of the page.</p> <p>After you enable FEC, check the monitoring report periodically. If losses are not persistent, disable FEC to avoid the overhead required to process recovery packets.</p>

- Recovery Packets and Data Packets
- Select the number of recovery packets (1 through 5) for the number of data packets (4 through 25). The settings should be based on the WAN error rate, as shown in Table 12.
- Note the following:
- Increasing the ratio of recovery packets to data packets reduces retransmissions, but requires more overhead. May be useful for losses caused by congestion.
 - Data packets must be a multiple of the recovery packets. For one recovery packet, the data packets can be 4 through 25; for 2 recovery packets, the data packets can be 4, 6, 8, and so on through 24.

Table 12: Recommended Data and Recovery Packets for FEC

Error Rate	Recovery Packets	Data Packets	Recovery Packet Overhead
6.25 %	1	4	25 %
5.00 %	1	5	20 %
4.25 %	1	6	17 %
3.50 %	1	7	14 %
3.00 %	1	8	13 %
2.75 %	1	9	11 %
2.50 %	1	10	10 %
2.25 % or less	1	11	9 %

5.
- Click Submit to enter the changes, or click Reset to discard them. You can now enable PFA for specific applications, as described in the next section.

Enabling Acceleration by Application

The following topics describe how to accelerate specific applications.

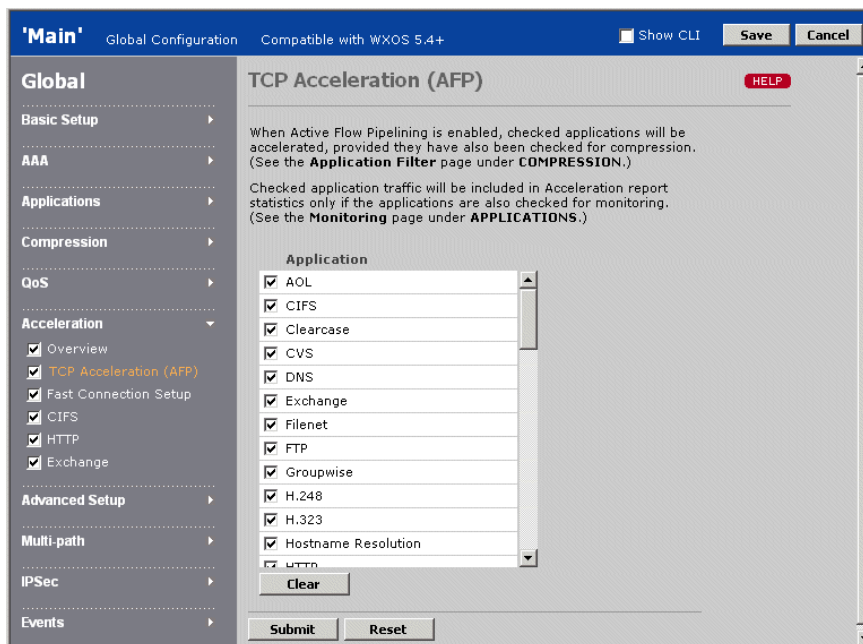
- “Enabling TCP Acceleration by Application” in the next section
- “Enabling Fast Connection Setup by Application” on page 221
- “Enabling Microsoft CIFS Acceleration” on page 222
- “Enabling Microsoft Exchange Acceleration” on page 225
- “Enabling HTTP Acceleration” on page 227

Enabling TCP Acceleration by Application

After you enable TCP Acceleration as described in “Enabling Acceleration by Endpoint” on page 214, you can select the applications whose outgoing traffic you want to accelerate. To enable TCP Acceleration for one or more applications:

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

Figure 118: Enabling TCP Acceleration by Application



2. Select the check box next to each application that you want to accelerate using TCP Acceleration. To disable TCP Acceleration for all applications, click Clear. The selected applications are accelerated only if they are also being compressed (refer to “Compressing Applications” on page 156).



NOTE: TCP Acceleration must be enabled on both the sending and receiving devices.

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

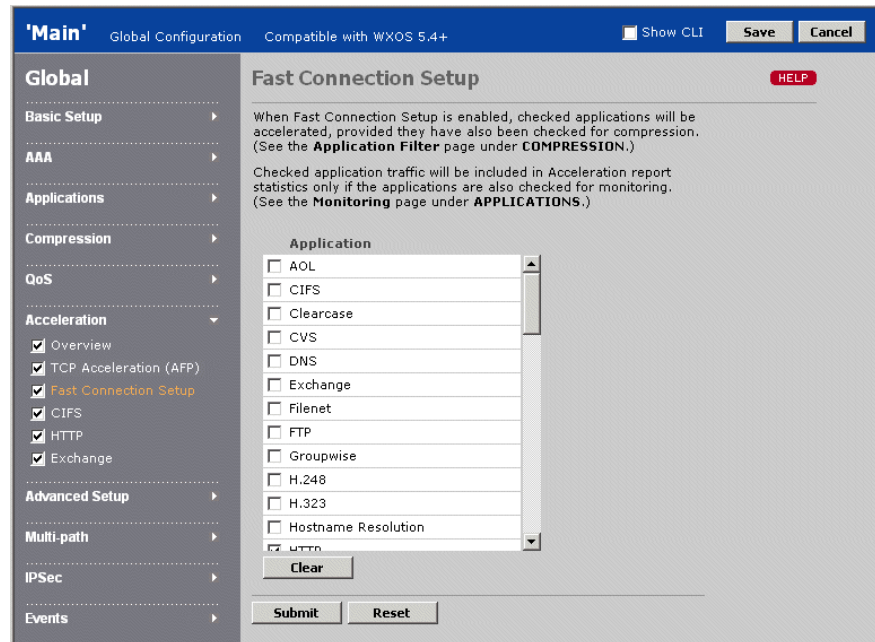
Enabling Fast Connection Setup by Application

After you enable Fast Connection Setup, as described in “Enabling Acceleration by Endpoint” on page 214, you can select the applications whose outgoing traffic you want to accelerate. Fast Connection Setup is intended for applications that have many short sessions, such as HTTP 1.0.

To enable Fast Connection Setup for one or more applications:

1. In the Configuration window, click Acceleration in the left-hand navigation frame, click Fast Connection Setup, and select the check box.

Figure 119: Enabling Fast Connection Setup by Application



2. Select the check box next to each application that you want to accelerate using Fast Connection Setup. To disable Fast Connection Setup for all applications, click Clear. The selected applications are accelerated only if they are also being compressed (refer to “Compressing Applications” on page 156)
3. Click Submit to enter the changes, or click Reset to discard them.

The selected applications have no effect if you load the configuration on a device where Fast Connection Setup is not enabled.

Enabling Microsoft CIFS Acceleration

You can accelerate all CIFS traffic using the default CIFS application definition, or you can create multiple application definitions to accelerate selected CIFS traffic, such as the traffic to or from a specific server.

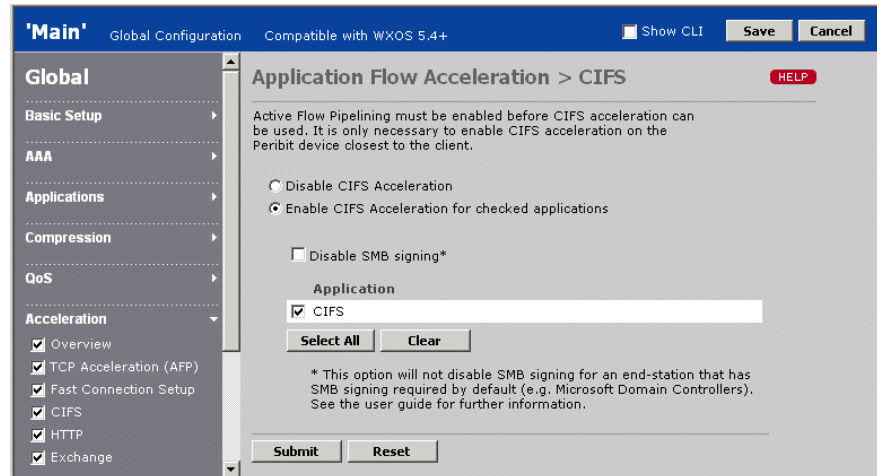
Microsoft CIFS traffic between Windows 2000 or XP clients and Windows 2000 or 2003 servers is accelerated. Enable CIFS acceleration on the devices closest to the clients (not required on the devices closest to the servers).

Note the following:

- Any new CIFS application definitions created must have an application type of CIFS and port numbers 139 and 445 (refer to “Configuring Application Definitions” on page 145)
- TCP Acceleration must be enabled on both the client- and server-side devices (refer to “Enabling TCP Acceleration by Application” on page 220).
- For WXOS 5.3/5.2 configurations, you must select All features on the Features/Topology page (refer to “Topology Settings for WXOS 5.3/5.2” on page 237). On the server-side device, you can conserve system resources by selecting All features except Application Flow Acceleration on the Features/Topology page.

To enable CIFS acceleration for one or more applications:

1. To add new CIFS application definitions to accelerate specific CIFS traffic:
 - a. Under Applications in a global or Applications partial configuration, click Definitions in the left-hand navigation frame, select the check box, and then click New Applications.
 - b. Select the CIFS application type, and be sure to specify port numbers 139 and 445. Complete the definition, and click Submit.
 - c. On the Application Definitions page, the new definition receives the order number of the generic CIFS definition. For example, if the order number of the generic definition was 6, the new definition becomes 6 and all subsequent definitions are incremented.
2. To enable acceleration for CIFS applications, under Acceleration in a global or Acceleration partial configuration, click CIFS in the left-hand navigation frame, and select the check box.

Figure 120: Enabling CIFS Acceleration

3. Select Enable CIFS Acceleration... 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.

The Disable SMB signing check box allows CIFS transactions to be accelerated for servers that have SMB signing enabled, but not required. To accelerate CIFS transactions for Windows 2000 and Windows 2003 domain controllers that require SMB signing, you must change the Group Policy on the domain controller (see Step 5).

4. Click Submit to enter the changes, or click Reset to discard them.
5. Verify that SMB signing is disabled on Windows 2000 and Windows 2003 domain controllers that are also file servers. On a Windows 2000 domain controller:
 - a. Open Active Directory Users and Computers on the domain controller.
 - b. Right click Domain Controllers and select Properties.
 - c. Click the Group Policy tab.
 - d. Click Default Domain Controllers Policy and select Edit.
 - e. Click Default Domain Controllers Policy/Computer Configuration/Windows Settings/Security Settings/Local Policies/Security Options.

- f. Disable the four signing options:
 - Digitally sign client communication (always)
 - Digitally sign client communication (when possible)
 - Digitally sign server communication (always)
 - Digitally sign server communication (when possible)
 - g. Reboot all domain controllers, member servers, and clients for which you want to accelerate CIFS traffic.
6. To verify that SMB signing is disabled on Windows 2000 clients:
- a. Click Start/Settings/Control Panel and select Administrative Tools.
 - b. Select Local Security Policy, and then select Local Policies/Security Options.
 - c. Disable the four signing options:
 - Digitally sign client communication (always)
 - Digitally sign client communication (when possible)
 - Digitally sign server communication (always)
 - Digitally sign server communication (when possible)

Enabling Microsoft Exchange Acceleration

You can accelerate all Exchange traffic using the default Exchange application definition, or you can create multiple application definitions to accelerate selected Exchange traffic, such as the traffic to or from a specific server.

Microsoft Exchange traffic between the following platforms is accelerated:

- Windows 2000 or XP clients and Windows 2000 or 2003 servers
- Outlook 2000, 2002 or 2003 clients and Exchange 5.5, 2000 or 2003 servers



NOTE: Traffic between an Outlook 2003 client and Exchange 2003 server is not accelerated, but WXC devices using NSC disk-based compression provide substantial benefits for such traffic without acceleration. Also, Exchange 2003/Outlook 2003 use Microsoft Exchange compression by default, which should be disabled (refer to <http://support.microsoft.com/?kbid=825371>).

Enable Exchange acceleration on the devices closest to the clients (not required on the devices closest to the servers).

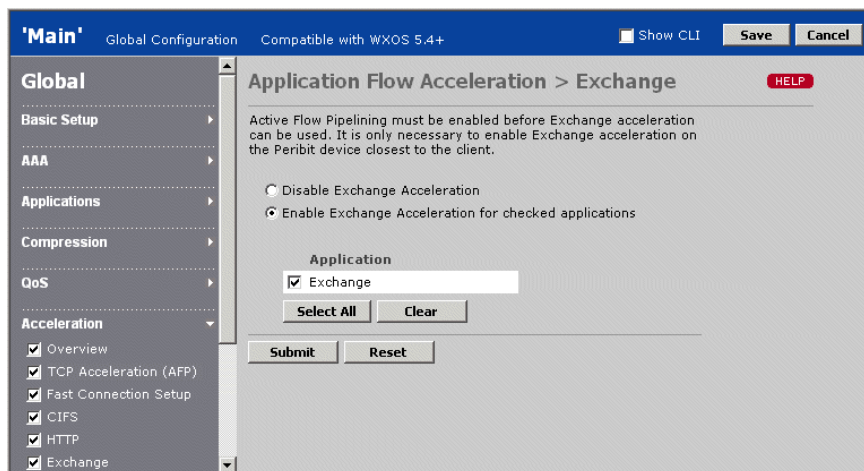
Note the following:

- Any new Exchange application definitions created must have an application type of Exchange and port number 135 (refer to “Configuring Application Definitions” on page 145)
- TCP Acceleration must be enabled on both the client- and server-side devices (refer to “Enabling TCP Acceleration by Application” on page 220).
- For WXOS 5.3/5.2 configurations, you must select All features on the Features/Topology page (refer to “Topology Settings for WXOS 5.3/5.2” on page 237). On the server-side device, you can conserve system resources by selecting All features except Application Flow Acceleration on the Features/Topology page.

To enable Exchange acceleration for one or more applications:

1. To add new Exchange application definitions to accelerate specific Exchange traffic:
 - a. Under Applications in a global or Applications partial configuration, click Definitions in the left-hand navigation frame, select the check box, and then click New Applications.
 - b. Select the Exchange application type, and be sure to specify port number 135. Complete the definition, and click Submit.
 - c. On the Application Definitions page, the new definition receives the order number of the generic Exchange definition. For example, if the order number of the generic definition was 20, the new definition becomes 20 and all subsequent definitions are incremented.
2. To enable acceleration for Exchange applications, under Acceleration in a global or Acceleration partial configuration, click Exchange in the left-hand navigation frame, and select the check box.

Figure 121: Enabling Exchange Acceleration



3. Select Enable Exchange Acceleration... 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 enter the changes, or click Reset to discard them.

Enabling HTTP Acceleration

You can accelerate all HTTP traffic using the default HTTP application definition, or you can create multiple application definitions to accelerate selected HTTP traffic, such as the traffic to or from a specific server.

Enable HTTP acceleration on the devices closest to the clients (not required on the devices closest to the servers).

Note the following:

- Any new HTTP application definitions created must have an application type of HTTP and the correct port number (refer to “Configuring Application Definitions” on page 145)
- TCP Acceleration must be enabled on both the client- and server-side devices (refer to “Enabling TCP Acceleration by Application” on page 220).
- For WXOS 5.3/5.2 configurations, you must select All features on the Features/Topology page (refer to “Topology Settings for WXOS 5.3/5.2” on page 237). On the server-side device, you can conserve system resources by selecting All features except Application Flow Acceleration on the Features/Topology page.

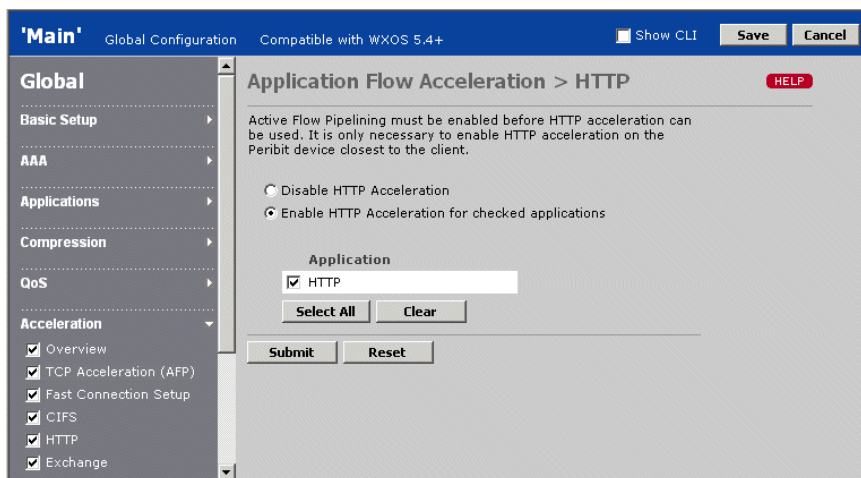


NOTE: HTTP traffic is not accelerated if a proxy server exists between the server-side device and the actual HTTP server. However, if the proxy server is between the Web client and the client-side device, HTTP traffic will be accelerated.

To enable HTTP acceleration for one or more applications:

1. To add new HTTP application definitions to accelerate specific HTTP traffic:
 - a. Under Applications in a global or Applications partial configuration, click Definitions in the left-hand navigation frame, select the check box, and then click New Applications.
 - b. Select the HTTP application type, and be sure to specify the HTTP port number (usually 80). Complete the definition, and click Submit.

- c. On the Application Definitions page, the new definition receives the order number of the generic HTTP definition. For example, if the order number of the generic definition was 4, the new definition becomes 4 and all subsequent definitions are incremented.
2. To enable acceleration for HTTP applications, under Acceleration in a global or Acceleration partial configuration, click HTTP in the left-hand navigation frame, and select the check box.

Figure 122: Enabling HTTP Acceleration

3. Select Enable HTTP Acceleration... 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 enter the changes, or click Reset to discard them.
 5. To enable caching of static objects on a WXC device:

On WXC devices, static objects are cached by default (header-and-body mode). To change the cache setting, refer to the “configure acceleration CLI command in the operator’s guide.

Configuring Advanced Setup Parameters

The following sections describe the global advanced setup parameters:

- “Configuring Topology Settings” on page 229
- “Configuring Source/Destination Filters” on page 238
- “Defining the Prime Time” on page 241
- “Configuring Packet Interception” on page 242
- “Configuring WAN Performance Monitoring” on page 258
- “Configuring Multiple Tunnels Between WX 100 Servers” on page 260
- “Adding CLI Commands to Configurations” on page 262

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 related to the topology of your network.

Selecting a Topology

Table 13 describes the topology settings and how they should be used.

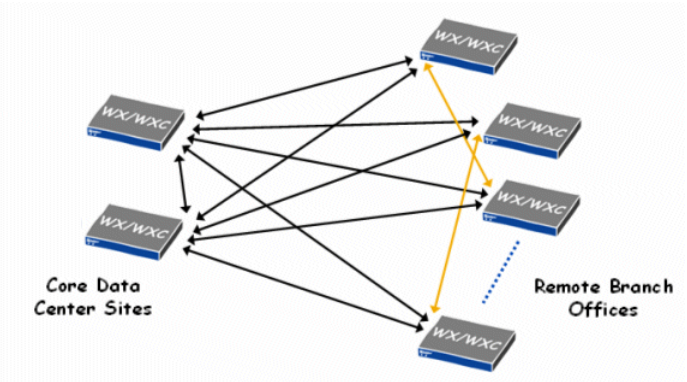
Table 13: Topology Settings and Recommended Use

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—the range of devices to be supported by the hub (refer to “Selecting a Community Size” on page 232).	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 topology size as the hubs.	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 152).
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 range of devices that each device supports.	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. This is called the “max-mem” setting in releases prior to WXOS 5.4.	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 123, 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.

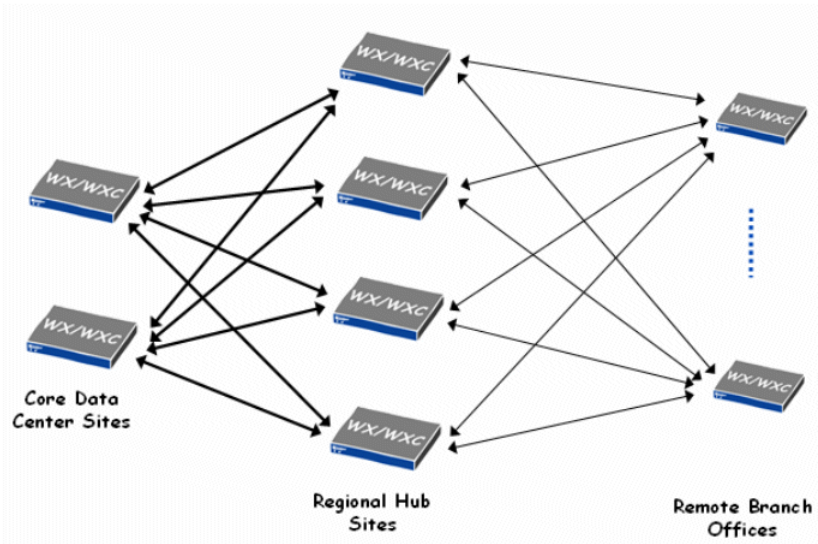
Figure 123: Partial Mesh Example



Tiered Network Example

In general, mixing hub, spoke, and mesh devices in the same community is not recommended. In Figure 124, 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 124: Tiered Network Example



Selecting a Community Size

The following sections describe how to set the community size for devices running WXOS 5.4, and for devices running WXOS 5.3 or 5.2.

- “WXOS 5.4 Configurations” on page 232
- “WXOS 5.3/5.2 Configurations” on page 234

WXOS 5.4 Configurations

Table 14 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 two tunnels, but both devices must be the same model and have the Point-to-Point setting.

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.

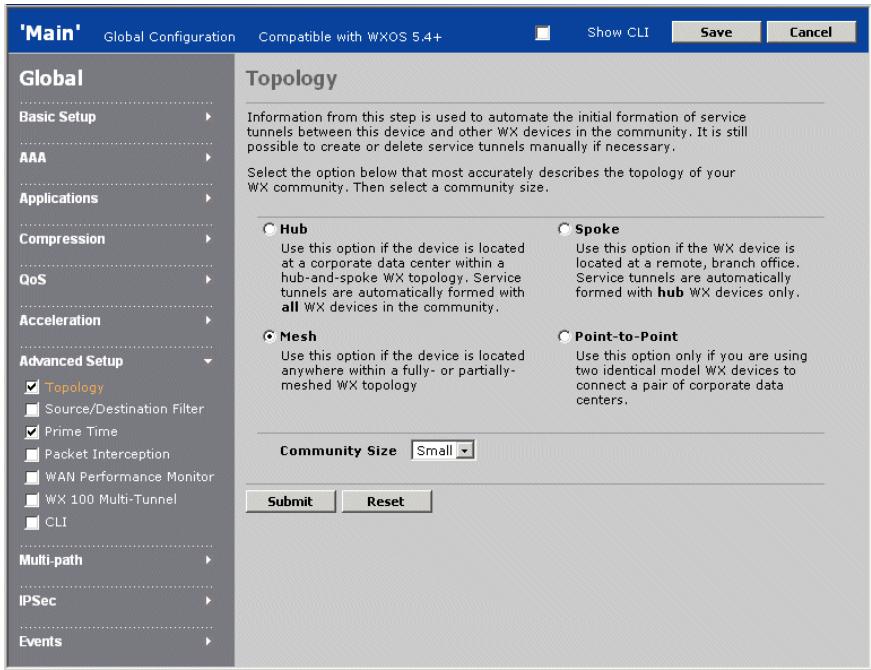
Table 14: Community Size by Device Type for WXOS 5.4

Device	Small	Large	Point-to-Point
WX 15	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 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 Configuration window, click Advanced Setup in the left-hand navigation frame, click Topology, and select the check box.

Figure 125: Topology Settings for WXOS 5.4



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 a Community Size” on page 232). 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 152. Note that a WX 15 must be a spoke, but a WX 100 can never be a spoke. |
| 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 a Community Size” on page 232). 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.

WXOS 5.3/5.2 Configurations

Table 15 and Table 16 show the numbered ranges of devices supported by each device type, based on the topology (hub or mesh), and whether Application Flow Acceleration (APP) and TCP Acceleration are enabled. Note that in WXOS 5.4, all features are enabled by default.

The **max-mem** option allocates all available memory for a limited number of tunnels, but all devices must be the same model and have the same topology setting (hub or mesh). When selecting a community (topology) size, select a size that is one range higher than needed when:

- The number of devices in the community is near the top of the range (within one or two devices).
- The community has multiple model types or more than two hubs.
- The community has mesh, hub, and spoke devices.

Table 15: Mesh Device Ranges for WXOS 5.3/5.2

Device	Mesh Ranges		
	No APP/TCP Accel	No APP	All Features
WX 15	0 = Up to 2	0 = Up to 2	0 = Up to 1
WX 20	0 = Up to 3 1 = Up to 6 2 = Up to 9 max-mem = 5	0 = Up to 3 1 = Up to 6 2 = Up to 9 max-mem = 2	0 = Up to 2 1 = Up to 5 2 = Up to 7 max-mem = 1
WX 60	0 = Up to 30 1 = Up to 50 2 = Up to 70 3 = Up to 79 4 = Up to 88 5 = Up to 99 max-mem = 7	0 = Up to 30 1 = Up to 45 2 = Up to 60 3 = Up to 75 4 = Up to 85 5 = Up to 95 max-mem = 6	0 = Up to 25 1 = Up to 35 2 = Up to 50 3 = Up to 65 4 = Up to 75 5 = Up to 85 max-mem = 5
WXC 250	0 = Up to 2 1 = Up to 5 2 = Up to 11 max-mem = 15	0 = Up to 5 1 = Up to 5 2 = Up to 10 max-mem = 5	0 = Up to 2 1 = Up to 5 2 = Up to 10 max-mem = 1

Table 15: Mesh Device Ranges for WXOS 5.3/5.2Table continued on next page

Device	Mesh Ranges		
	No APP/TCP Accel	No APP	All Features
WXC 500	0 = Up to 8	0 = Up to 8	0 = Up to 7
	1 = Up to 13	1 = Up to 13	1 = Up to 11
	2 = Up to 18	2 = Up to 18	2 = Up to 16
	3 = Up to 23	3 = Up to 23	3 = Up to 21
	4 = Up to 28	4 = Up to 27	4 = Up to 26
	5 = Up to 37	5 = Up to 37	5 = Up to 35
	max-mem = 20	max-mem = 3	max-mem = 1
WXC 590	0 = Up to 120	0 = Up to 30	0 = Up to 20
	1 = Up to 120	1 = Up to 40	1 = Up to 30
	2 = Up to 120	2 = Up to 50	2 = Up to 40
	3 = Up to 120	3 = Up to 60	3 = Up to 50
	4 = Up to 120	4 = Up to 80	4 = Up to 60
	5 = Up to 120	5 = Up to 100	5 = Up to 80
	max-mem = 40	max-mem = 10	max-mem = 2
WX 100 (no clients)	0 = Up to 60	0 = Up to 60	0 = Up to 60
	1 = Up to 100	1 = Up to 90	1 = Up to 90
	2 = Up to 140	2 = Up to 130	2 = Up to 120
	3 = Up to 170	3 = Up to 150	3 = Up to 140
	4 = Up to 200	4 = Up to 170	4 = Up to 160
	5 = Up to 220	5 = Up to 190	5 = Up to 180
	max-mem = 20	max-mem = 15	max-mem = 15
WX 100 (with clients)	When a WX 100 has one or more client devices (stack configuration), the topology ranges displayed in the Web console are for the clients, not the WX 100. Since hosting tunnels on a WX 100 server is not recommended, the maximum range of devices is the sum of the ranges for each client device. For example, a WX 100 in a mesh that uses all features and has two WXC 590 clients can support up to 160 devices (80 + 80).		

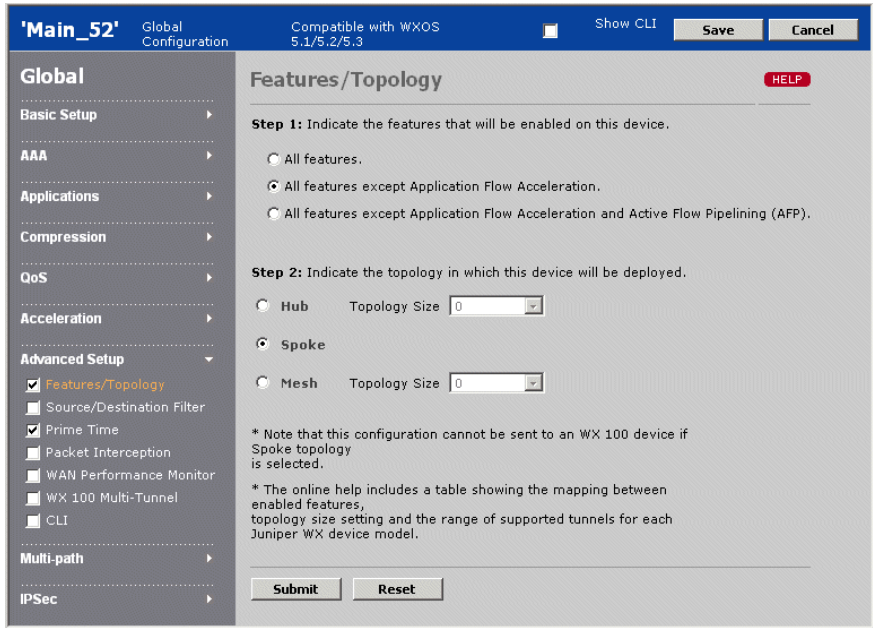
Table 16: Hub Device Ranges for WXOS 5.3/5.2

Device	Hub Ranges		
	No APP/TCP Accel	No APP	All Features
WX 15	0 = Up to 2	0 = Up to 2	0 = Up to 1
WX 20	0 = Up to 5 1 = Up to 7 2 = Up to 8 max-mem = 3	0 = Up to 3 1 = Up to 8 2 = Up to 10 max-mem = 2	0 = Up to 2 1 = Up to 5 2 = Up to 7 max-mem = 1
WX 60	0 = Up to 30 1 = Up to 60 2 = Up to 69 3 = Up to 88 4 = Up to 105 5 = Up to 128 max-mem = 7	0 = Up to 30 1 = Up to 50 2 = Up to 69 3 = Up to 87 4 = Up to 105 5 = Up to 128 max-mem = 6	0 = Up to 25 1 = Up to 45 2 = Up to 65 3 = Up to 85 4 = Up to 103 5 = Up to 125 max-mem = 5
WXC 250	0 = Up to 6 1 = Up to 6 2 = Up to 11 max-mem = 15	0 = Up to 5 1 = Up to 6 2 = Up to 11 max-mem = 5	0 = Up to 2 1 = Up to 5 2 = Up to 10 max-mem = 1
WXC 500	0 = Up to 14 1 = Up to 14 2 = Up to 18 3 = Up to 23 4 = Up to 32 5 = Up to 49 max-mem = 20	0 = Up to 13 1 = Up to 13 2 = Up to 17 3 = Up to 22 4 = Up to 32 5 = Up to 49 max-mem = 3	0 = Up to 7 1 = Up to 12 2 = Up to 16 3 = Up to 21 4 = Up to 30 5 = Up to 46 max-mem = 1
WXC 590	0 = Up to 120 1 = Up to 120 2 = Up to 120 3 = Up to 120 4 = Up to 120 5 = Up to 120 max-mem = 40	0 = Up to 30 1 = Up to 40 2 = Up to 50 3 = Up to 60 4 = Up to 80 5 = Up to 120 max-mem = 10	0 = Up to 20 1 = Up to 30 2 = Up to 40 3 = Up to 50 4 = Up to 70 5 = Up to 100 max-mem = 2
	No APP/TCP Accel	No APP	All Features
WX 100 (no clients)	0 = Up to 60 1 = Up to 120 2 = Up to 170 3 = Up to 220 4 = Up to 270 5 = Up to 320 max-mem = 20	0 = Up to 60 1 = Up to 110 2 = Up to 150 3 = Up to 190 4 = Up to 230 5 = Up to 280 max-mem = 15	0 = Up to 60 1 = Up to 110 2 = Up to 140 3 = Up to 180 4 = Up to 220 5 = Up to 260 max-mem = 15
WX 100 (with clients)	The maximum range of devices is the sum of the ranges for each client device. For example, a WX 100 hub that uses all features and has two WXC 590 clients can support up to 200 devices (100 + 100).		

To change the community topology settings:

1. In the Configuration window, click Advanced Setup in the left-hand navigation frame, click Features/Topology, and select the check box.

Figure 126: Topology Settings for WXOS 5.3/5.2



2. Select the set of features to be used. If this setting is changed on a device, all tunnels will be reset.

All Features	Allows all licensed features to be used.
All features except Application Flow Acceleration	Allows the basic features and TCP Acceleration to be used, but not CIFS, Exchange, and HTTP acceleration. Note that WXC devices require TCP Acceleration for optimum performance.
All features except Application Flow Acceleration and TCP Acceleration	Allows all licensed features to be used, except TCP Acceleration and CIFS, Exchange, and HTTP acceleration. Excluded features cannot be accessed from the Web console.

3. Select one of the following topology settings:

Hub	A hub can compress and accelerate traffic for all devices in the community. Select the range of devices in the community (refer to “Selecting a Community Size” on page 232). If a community has multiple hubs, each hub must specify the same range of devices.
Spoke	By default, a spoke can compress and accelerate traffic only for devices that are designated as hubs. To enable compression between spoke devices, refer to “Configuring Endpoints for Compression” on page 152. Note that a WX 15 must be a spoke, but a WX 100 can never be a spoke.
Mesh	A mesh device can compress and accelerate traffic for all other devices in the community. Select the range of devices in the community (refer to “Selecting a Community Size” on page 232).

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

Configuring Source/Destination Filters

You can create a list of source and destination addresses or subnet pairs that are either included or excluded from compression. This source/destination filter applies to all application traffic sent from the LAN to the WAN. To enable or disable compression by application, refer to “Compressing Applications” on page 156. The source/destination filter is applied before the application filter, and is more efficient.



NOTE: If you disable compression between a source and destination, Packet Flow Acceleration between those points is also disabled.

For example, to disable compression for all traffic from a local subnet, create a “Do not compress” entry and specify the subnet as the source and enter an asterisk (*) as the destination. To disable compression for all traffic sent to the subnet from other devices, you must disable the advertisement of the subnet (refer to “Advertising Compression Subnets” on page 109).

Note the following:

- If you disable compression between a source and destination, traffic between those points cannot be accelerated. Also, the traffic is managed by the Outbound QoS policies defined for the Default traffic class under the “Other traffic” endpoint.
- Source/destination filters are disallowed on off-path devices that use RIP for packet interception. Also, they should not be used with the External packet interception mode.

To define source and destination subnets:

1. In the Configuration window, click Advanced Setup in the left-hand navigation frame, click Source/Destination Filter, and select the check box.

Figure 127: Filtering Compression by Source and Destination

'Main' Global Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Global

- Basic Setup
- AAA
- Applications
- Compression
- QoS
- Acceleration
- Advanced Setup**
 - ☒ Topology
 - ☒ **Source/Destination Filter**
 - ☒ Prime Time
 - ☐ Packet Interception
 - ☐ WAN Performance Monitor
 - ☐ WX 100 Multi-Tunnel
 - ☐ CLI

Source/Destination Filter

☒ Off (default)
☐ Compress data between the following source/destination pairs ONLY
☐ DO NOT compress data between the following source/destination pairs

Source	Destination	Bidirectional
		<input type="checkbox"/>

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

Submit Reset

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 the other WX 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 (the traffic cannot be accelerated, and is managed by the outbound QoS policies defined for the Default traffic class under the "Other traffic" endpoint).

3. Specify the following information:

Source	<p>Enter a source IP address or subnet. The general format is: address/subnet_mask</p> <p>The default subnet mask is "255.255.255.255". An asterisk (*) with no subnet mask indicates any source IP address.</p>
Destination	<p>Enter a destination IP address or subnet (same format as the source address). An asterisk (*) indicates any destination IP address.</p>
Bidirectional	<p>Select the check box to include traffic sent from the destination to the source. This option is particularly useful for creating "do not compress" lists in Demo Mode.</p> <p>In Profile Mode, you should exclude all traffic sent to the subnet where the device is installed. For more information about Profile Mode, refer to the <i>WX/WXC Operator's Guide</i>.</p>

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

Defining the Prime Time

The prime time setting lets you specify the days of the week and hours of the day when network performance is most important. The prime time can be used to filter performance statistics, and to specify bandwidth management policies for prime-time and non prime-time hours.



NOTE: The prime time can be used to filter reports in the device Web console, but not in CMS.

For example, to view 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 Configuration window, click Advanced Setup in the left-hand navigation frame, click Prime Time, and select the check box.

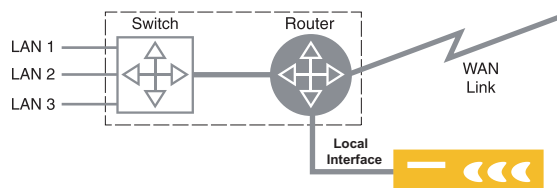
Figure 128: Defining the Prime Time

2. To set the prime time, select the Use Prime Time check box, select a time range, and select the days of the week.
3. Click Submit to enter the changes, or click Reset to discard them.

Configuring Packet Interception

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 129). 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 129: 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.

The following topics describe how to configure packet interception. 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 244
- “RIP Router/Switch Configuration Commands” on page 247
- “WCCP Router Configuration Commands” on page 250
- “External Policy-Based Router Commands” on page 255
- “Alternatives to Packet Interception” on page 256

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 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 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 247.

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 tunnelled) 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 250.



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 the cluster CLI commands in the *WX/WXC Operator's Guide*).

External

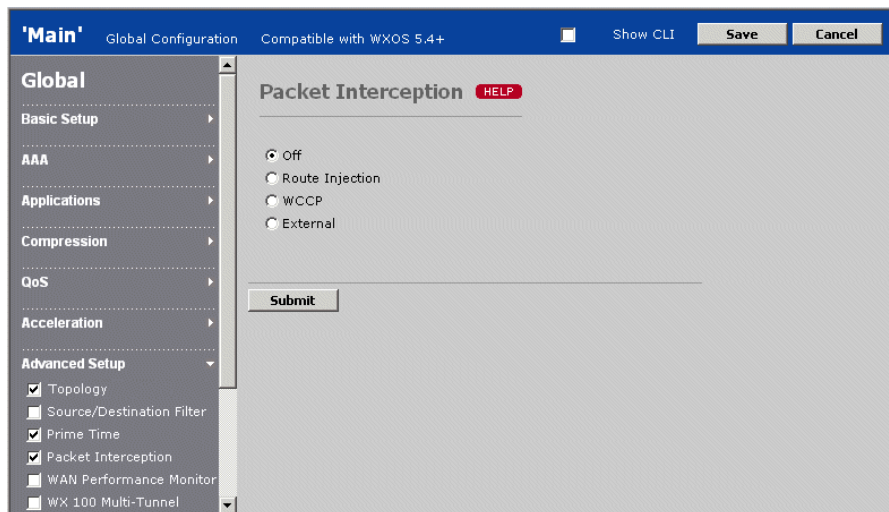
The WAN edge router is configured to route traffic to the off-path device. The off-path device should be connected directly to the router, and must be the only device on the port. You can also connect the off-path device to a dedicated VLAN on a Layer 3 switch. Refer to the sample router commands in “External Policy-Based Router Commands” on page 255.

Configuring Packet Interception for Off-Path Devices

To configure packet interception for an off-path device:

1. In the Configuration window, click Advanced Setup in the left-hand navigation frame, click Packet Interception, and select the check box.

Figure 130: 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.
Inter-packet delay	To reduce the load on slower routers, enter the number of milliseconds between each packet when multiple packets are generated for a single RIP update (0 through 50). The default is 0.

You can lower the RIP update timers to reduce the failover time (not recommended if RIP is used for network-wide routing). To change the frequency of RIP updates or the cost assigned to each advertised route, refer to the “configure packet-interception” CLI command.

- **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 the “configure acceleration cluster” CLI command in the <i>WX/WXC Operator’s Guide</i>).</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 device, specify a lower value to avoid diverting traffic from the Web cache.</p>
WCCP Auth. Password	If the WAN edge router uses WCCP authentication, enter the WCCP password specified on the router.

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

IP Protocol	<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 off-path 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 255.

3. Click Submit to enter the changes.
4. 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 109). Since only the Local interface is connected to the network, the device cannot distinguish between LAN- and WAN-side subnets.



CAUTION: If you use RIP for packet interception, and 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.

The following sections provide sample router configuration commands to support each method of packet interception.

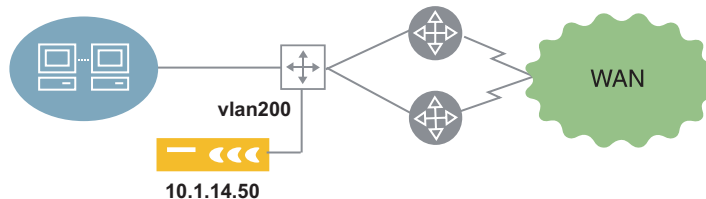
RIP Router/Switch Configuration Commands

In general, an off-path device should be connected to a dedicated port on a router or Layer 3 switch. RIP is then configured on the router or switch where the device is connected. If the off-path device is connected to a Layer 2 switch, RIP is configured on the router. In each case, the RIP configuration is essentially the same.

Single Layer 3 Switch

The following commands provide an example of how to configure RIP on a Layer 3 Cisco switch (Figure 131). Installing the off-path device on a dedicated VLAN is recommended to reduce the routing failover time if the device fails. The port where the off-path device is connected should be the only port in the VLAN. Note that the load balancing done by the switch across the two routers is not affected.

Figure 131: Off-Path Device Connected to a Layer 3 Switch



1. Enable RIP version 2:

```
router rip
version 2
```

2. If RIP is used only for packet interception, you can lower the RIP timers to reduce the failover time (may cause instability if RIP is used for network-wide routing):

```
timers basic 5 15 15 30
```

3. Enable RIP to listen passively on all interfaces:

```
passive-interface default
```

4. Specify the subnet where the off-path device is installed:

```
network 10.0.0.0
```

5. Specify the RIP administrative distance to be lower than all other methods used by the router or switch to discover routes (such as OSPF):

```
distance 30
```

6. Disable auto-summarization of routes:

```
no auto-summary
```

Do not redistribute the RIP routes to any other routing protocol, such as OSPF. The advertised RIP routes apply only to the configured router or switch and the off-path WX device. If RIP is used only for packet interception, no other routers should be affected.



NOTE: If you change the number of seconds between RIP updates (the default is 30), you must specify the same value on the off-path device. To match this example, enter the following CLI command on the off-path device:

```
config packet-interception rip set update-timer 5
```

To view the RIP routes advertised by the off-path device, enter the following command:

```
show ip route rip
```

If packet interception is working correctly, you should see routes like the following. In this example, 10.1.14.50 is the off-path device, and the IP address of the remote WX device (10.1.203.50) has been carved out.

```
10.1.0.0/16 is variably subnetted, 24 subnets, 9 masks
R 10.1.203.128/25 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.51/32 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.48/31 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.52/30 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.56/29 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.32/28 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.0/27 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
R 10.1.203.64/26 [30/2] via 10.1.14.50, 00:00:23, Ethernet0/1
```

To view debugging information for RIP events on a Cisco router:

```
debug ip rip events
```

Sample debugging information:

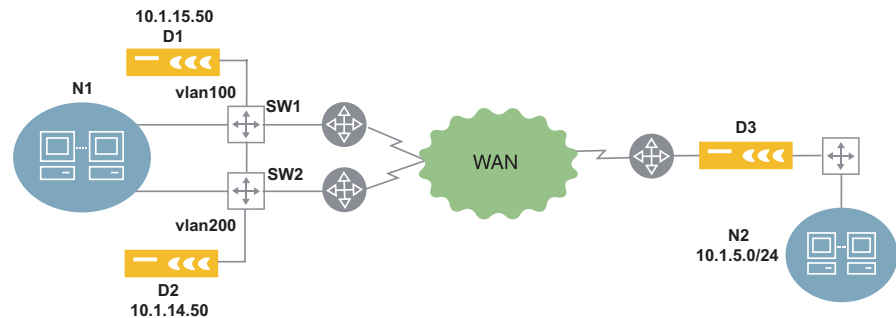
```
1w1d: RIP: received v2 update from 10.1.14.50 on Ethernet0/1
1w1d: RIP: Update contains 8 routes
```

You can also enter "debug ip rip database" or "debug ip rip trigger" for more details.

Dual Off-Path Devices on Two Layer 3 Switches

In Figure 132, 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 132: Dual Off-Path Devices on Two Layer 3 Switches



7. 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
```

8. Enable RIP on SW2 so that a five-hop offset is added to the RIP routes received from D2 (which are the routes advertised by D3):

```
access-list 10 permit host any
```

```

router rip
version 2
timers basic 5 15 15 30
offset-list 10 in 5 interface vlan200
network 10.0.0.0
distance 30
no auto-summary

```

Thus, the routes from D2 have six hops on SW2, and seven hops on SW1, while the same routes from D1 have one hop on SW1 and two hops on SW2. The routes from D2 are used only if D1 fails.

If the D1 and D2 are on the same subnet, you can specify the offset on D2:

```
config packet-interception rip set metric 7
```



NOTE: If you change the number of seconds between RIP updates (the default is 30), you must specify the same value on the off-path devices. To match this example, enter the following CLI command on the off-path device:

```
config packet-interception rip set update-timer 5
```

WCCP Router Configuration Commands

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:

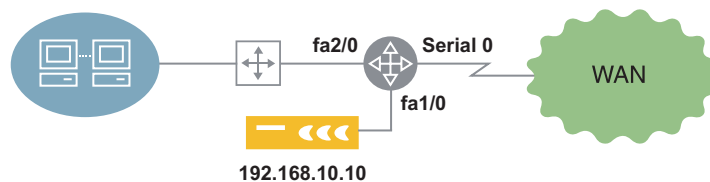
- **Router ID.** When using the WCCP multicast group function, make sure that the loopback address that the Cisco device chooses as the source address for GRE traffic is active and reachable. A known Cisco WCCP bug is that a configured, but down, IP address is sometimes used as the source address.
- **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 133.

Figure 133: Off-Path Device Connected to a Router



1. Define an access list that specifies the traffic that is eligible for redirection to the off-path device:

```
access-list 120 permit ip any any
```

2. If the off-path device assigns WCCP service IDs 85 and 87 to TCP and UDP, respectively, create the two service IDs on the router. Include the password if authentication is enabled.

```
ip wccp 85 redirect-list 120 password <password>
ip wccp 87 redirect-list 120 password <password>
```

3. To redirect traffic from the outbound WAN interface, specify the WCCP service IDs to be redirected:

```
interface Serial 0
ip address 192.168.5.103 255.255.255.0
ip wccp 85 redirect out
ip wccp 87 redirect out
```

Alternatively, to redirect traffic from the inbound interface from the switch:

```
interface FastEthernet 2/0
ip address 192.168.5.103 255.255.255.0
ip wccp 85 redirect in
ip wccp 87 redirect in
```

Verify that Cisco Express Forwarding is enabled:

```
ip cef
```

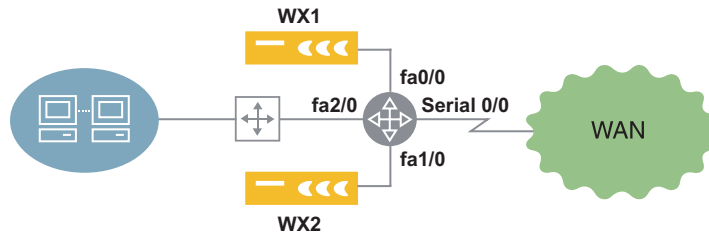


NOTE: If you define a service ID on the router, but omit the redirect commands, no traffic is redirected to the WX, but entering a “show packet-interception” command on the WX will indicate the service is connected.

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 134: Multicast Example for the Cisco 3640



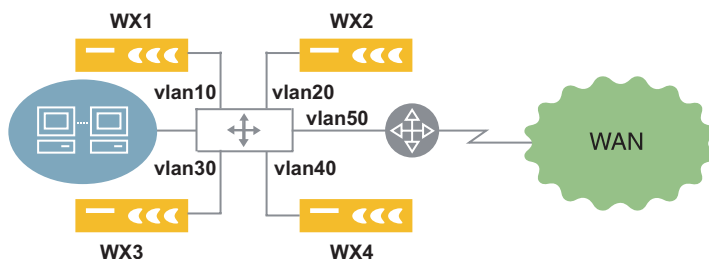
```

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 135: 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 off-path device at IP address 192.168.10.10. The access list shown here redirects all packets, but it can be as specific as necessary.

```

interface FastEthernet 0/0
ip address 192.168.9.1 255.255.255.0
ip policy route-map juniper
access-list 120 permit ip any any

route-map juniper permit 50
match ip address 120
set ip next-hop 192.168.10.10

```

If the off-path device is connected to a dedicated VLAN on a Layer 3 switch, the commands are almost the same, except that the policy is applied to the switch on the inbound interface from the LAN:

```

interface Vlan200
ip address 192.168.9.1 255.255.255.0
ip policy route-map juniper

```



NOTE: Use the “set ip next-hop” command to redirect packets to the IP address of the off-path device. Do not use the “set interface” command to redirect traffic to the interface where the off-path device is connected.

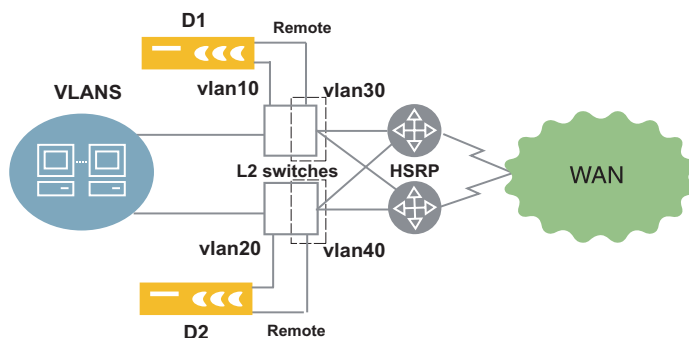
Alternatives to Packet Interception

In some environments, you can install an off-path device by connecting the Local and Remote interfaces to different VLANs on the same switch. Packet interception is not used.

Layer 2 Switch Sandwich

In the high-availability environment in Figure 136, D1 and D2 are connected in “two-legged” VLANs on two Layer 2 switches. All traffic is switched through the devices as it passes to and from the WAN routers.

Figure 136: Layer 2 Switch Sandwich



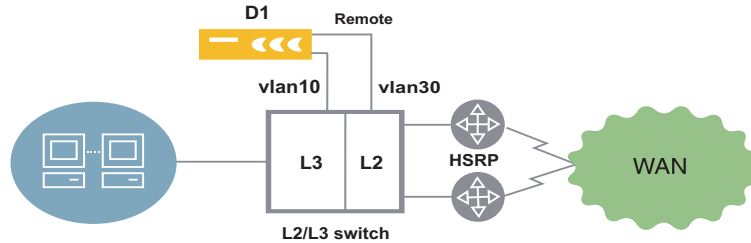
Note the following:

- The Local interface is placed in the original VLAN that previously connected the switch port to the WAN router.
- The Remote interface is placed in a new VLAN along with the switch ports that feed the WAN routers.
- The default gateway of D1 and D2 is the HSRP address of the WAN routers. If one router fails, traffic is directed to the other router.
- Use a cross-over cable to connect the Local interface to the switch so that traffic is blocked if one device fails. The Layer 3 switches can then route the traffic through the other device.

Layer 3 Switch Sandwich

Figure 137 shows a single device connected across Layer 2 and Layer 3 VLANs on an L2/L3 switch. All traffic is switched through D1 as it passes to and from the WAN routers.

Figure 137: Layer 3 Switch Sandwich



Note the following:

- Hosts on the local LAN must point to the HSRP default gateway on same subnet.
- The Local interface is placed in the original VLAN that previously connected the switch port to the WAN router.
- The Remote interface is placed in a new Layer 2 VLAN along with the switch ports that feed the WAN routers.
- The default gateway of D1 is the HSRP address of the WAN routers. If one router fails, traffic is directed to the other router.

Configuring WAN Performance Monitoring

WAN performance monitoring lets each WX device measure the latency and loss to one or more remote WX devices. Probes are sent at an adjustable rate to each selected endpoint, and the loss and latency calculated for each WAN path is shown on the WAN Performance report. If the loss or latency exceeds the specified thresholds, an informational SNMP trap and Syslog entry are generated, and an event icon is shown on the report.

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 Multi-Path Addresses” on page 116).

To enable WAN performance monitoring:

1. In the Configuration window, click Advanced Setup in the left-hand navigation frame, click WAN Performance Monitor, and select the check box.

Figure 138: Configuring WAN Performance Monitoring

Device Name	IP Address	Latency Threshold (msec)
SR-10.87.240.2	10.87.240.2	100
SM-10.87.242.2	10.87.242.2	100
CMSQA-245	10.87.245.2	100
10.87.245.200	10.87.245.200	100

2. Select the Enable WAN Performance Monitoring... check box.

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

Latency Threshold	<p>Enter the round-trip time (RTT) threshold in milliseconds (20 to 5000). Traps, Syslog entries, and report events are generated when the threshold is exceeded, and again when latency drops below the threshold.</p> <p>By default, a probe tests the path 12 times per minute. Traps are generated when the median latency exceeds the threshold for four consecutive minutes or if two or more probes are lost per minute for four consecutive minutes. To change these settings, refer to the “configure wan-performance-monitor” CLI command).</p> <p>Note that availability on the WAN Performance report is measured as the percentage of minutes for which at least one probe was acknowledged.</p>
-------------------	---

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

Configuring Multiple Tunnels Between WX 100 Servers

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



NOTE: Disable bandwidth detection between the WX 100 servers, as described in “Defining Outbound QoS Endpoints” on page 195 (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 Configuration window, click Advanced Setup in the left-hand navigation frame, click WX 100 Multi-Tunnel, and select the check box.

Figure 139: Configuring Multiple Tunnels Between WX 100 Servers

WX 100 Multi-Tunnel

This page allows you to specify the maximum number of tunnels that can be formed between this and other remote WX 100 endpoints in the community. Enter the maximum number of tunnels (1-6) to the right of the WX 100 endpoint. When you are finished, click **Submit**.

Note: WX 100 endpoints must be enabled for compression in order to take advantage of the Multi-Tunnel feature.

WX 100 Endpoint	IP Address	Maximum Tunnels
Stack-10.87.249.2	10.87.249.2	2

Add/Remove Endpoints...

Submit **Reset**

2. To establish multiple outbound tunnels to one or more remote WX 100 servers:
 - a. Click Add/Remove Endpoints.
 - b. Select a community from the Community/Device Group list, or select all devices. The device name and IP address are shown for each device. The IP address is enclosed in parentheses.

- c. Select the devices you want to configure, and click Add. To remove devices from the Multi-Tunnel Endpoints list, select the devices and click Remove.
- d. Repeat Steps b and c for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
- e. If one or more WX 100 servers are not listed, click Manual Entry and enter the IP addresses for each device (one per line), and click Submit.
- f. When you are done, click Submit.
- g. Enter the number of tunnels (1 to 6) to each remote server. If you specify two or more tunnels, the local and remote servers must have the same number of clients as the number of specified tunnels.
- h. Click Submit.

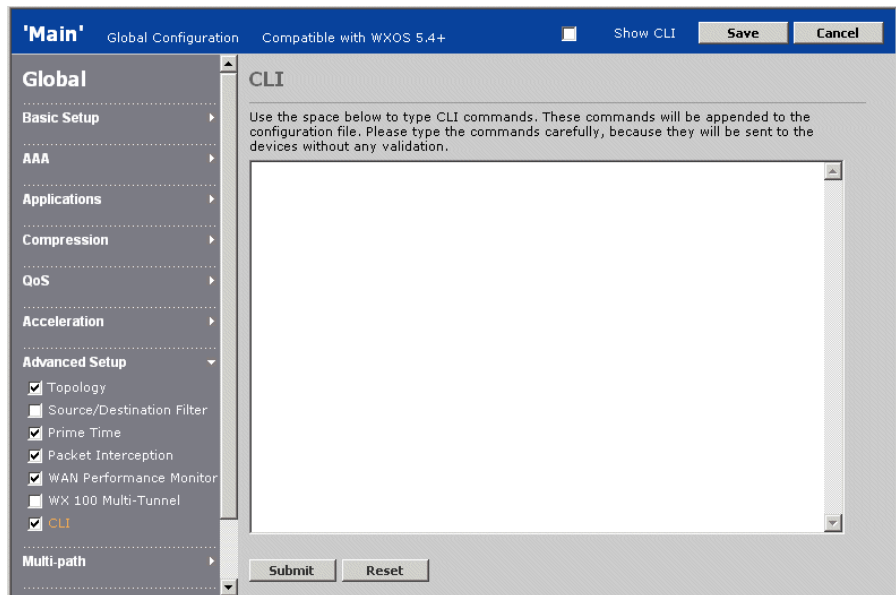
Adding CLI Commands to Configurations

You can append CLI commands to a global configuration or an Advanced Setup partial configuration. This is intended primarily for use by support representatives to troubleshoot problems.

To append CLI commands to a global configuration:

1. In the Configuration window, click Advanced Setup in the left-hand navigation frame, click CLI, and select the check box.

Figure 140: Appending CLI Commands to a Global Configuration



2. Enter CLI commands provided by the support representative.
3. Click Submit to enter the changes, or click Reset to discard them.

Configuring Policy-Based Multi-Path

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

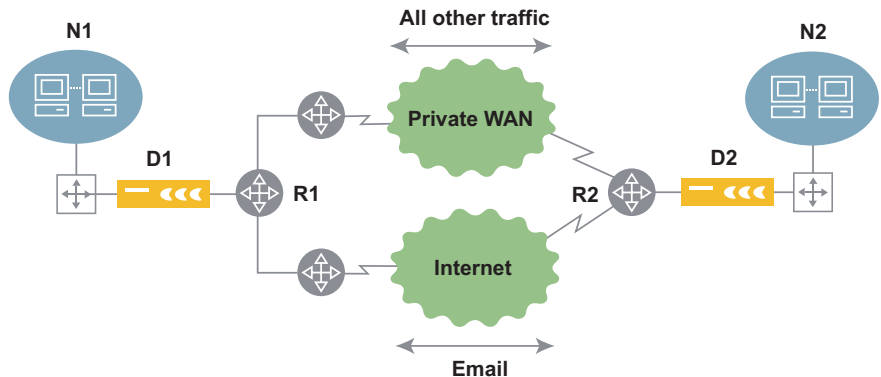
To use Multi-Path, you configure both devices so that outgoing packets intended for the secondary path are marked with a secondary source IP address and, optionally, with a specific gateway address or ToS/DSCP value. In most cases, you must configure the WAN routers to route the marked packets to the appropriate path. The traffic for the preferred path (primary or secondary) is specified by traffic class, where each class contains one or more applications.



NOTE: The secondary IP address on each device must be specified using the device Web console or a Device Settings partial configuration.

For example, in Figure 141, most traffic might normally be sent over the private WAN, while email traffic is sent over the Internet. D1 and D2 mark email traffic with a secondary IP address, and R1 and R2 are configured to route the marked traffic to the Internet. If the private WAN fails, selected application traffic can be diverted automatically to the Internet; if the Internet latency exceeds a specified threshold, email traffic can be diverted to the private WAN. Traffic is switched back to the preferred path when conditions return to normal.

Figure 141: Multi-Path Deployment



The following topics describe how to configure Multi-Path:

- “Procedure for Configuring Multi-Path” in the next section
- “Enabling Policy-Based Multi-Path” on page 265
- “Defining Multi-Path Templates” on page 266
- “Defining Multi-Path Endpoints” on page 268
- “Configuring Routers to Support Multi-Path” on page 270

Procedure for Configuring Multi-Path

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

1. Verify that compression is enabled between the two devices (refer to “Configuring Endpoints for Compression” on page 152).
2. Verify that the appropriate traffic classes are defined (refer to “Assigning Applications to Traffic Classes” on page 149). Outbound QoS can be on or off.
3. For each device, specify a secondary IP address and primary and secondary gateway addresses (if applicable) using the device Web console or a Device Settings partial configuration (refer to “Configuring Multi-Path Addresses” on page 116). The Device Settings configuration must be loaded on each device before you can configure the Multi-Path endpoints.
4. Enable Multi-Path and, optionally, specify primary and secondary ToS/DSCP values (refer to “Enabling Policy-Based Multi-Path” on page 265).
5. Define templates that specify the preferred path (primary or secondary) for each traffic class and the conditions when the traffic for each class can be switched to the alternate path (refer to “Defining Multi-Path Templates” on page 266).
6. Apply a template to each remote device that supports Multi-Path, and specify the congestion and latency thresholds for each path (refer to “Defining Multi-Path Endpoints” on page 268).
7. If necessary, configure the WAN router to route traffic to the appropriate path (refer to “Configuring Routers to Support Multi-Path” on page 270).

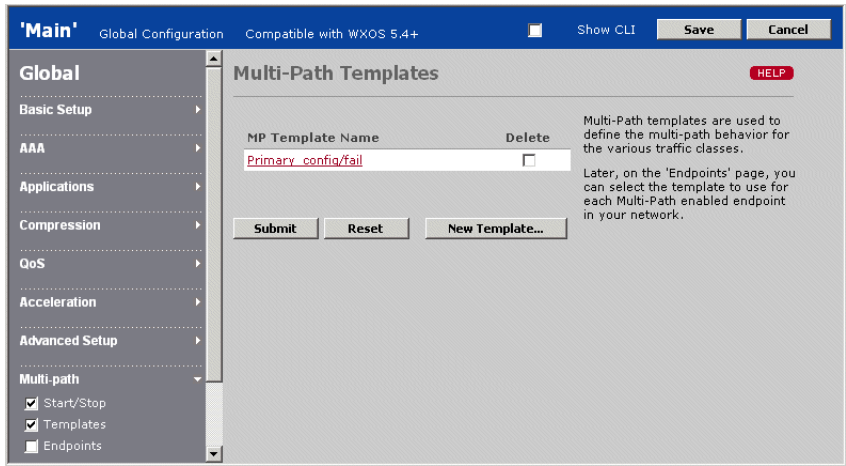
Enabling Policy-Based Multi-Path

To enable Multi-Path on a device from CMS, you must first specify a secondary IP address and primary and secondary gateway addresses (if applicable) using the device Web console or a Device Settings partial configuration (refer to “Configuring Multi-Path Addresses” on page 116). The Device Settings configuration must be loaded on each device before you can configure the Multi-Path endpoints.

To enable Multi-Path:

- 1. In the Configuration window, click Multi-path in the left-hand navigation frame, click Start/Stop, and select the check box.

Figure 142: Multi-Path Start/Stop Page



- 2. Specify the following information:

Multi-Path	Select Enabled to activate the Multi-Path feature.
IP Precedence/DSCP	<p>Optionally, you can mark packets sent on the primary and secondary paths with different ToS/DSCP values. Select IP Precedence or DSCP and enter a ToS IP precedence value (0 to 7) or DSCP value (0 to 63) for packets sent on the primary and/or secondary paths.</p> <p>NOTE: 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 199)■ Control packets (refer to the “configure reduction” CLI command) <p>Multi-path DSCP values also override ToS marking for router-based balancing (refer to the “configure route” CLI command).</p>

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

Defining Multi-Path Templates

At least one Multi-Path template is required to specify the preferred path for each traffic class, and the conditions under which the traffic for each class can be switched to the alternate path. To assign a template to each remote device that supports Multi-Path, refer to “Defining Multi-Path Endpoints” on page 268.

To define Multi-Path templates:

1. In the Configuration window, click Multi-path in the left-hand navigation frame, click Templates, and select the check box.

Figure 143: Defining Multi-Path Templates

The screenshot shows the 'Main' configuration window with the 'Global Configuration' tab selected. The left-hand navigation pane is expanded to 'Multi-path', and the 'Templates' checkbox is checked. The main content area is titled 'Multi-Path Templates' and contains a table with one row: 'Primary_config/fail'. To the right of the table is a 'Delete' checkbox. Below the table are buttons for 'Submit', 'Reset', and 'New Template...'. A 'HELP' button is located in the top right corner of the main content area. A text box on the right explains that Multi-Path templates are used to define multi-path behavior for various traffic classes and that users can select a template on the 'Endpoints' page.

MP Template Name	Delete
Primary_config/fail	<input type="checkbox"/>

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

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

2. To add a new template, click New Template, specify the following information, and click Submit:

Template Name	Enter the template name (up to 20 characters).
For each traffic class, select the following (to add new traffic classes, refer to “Assigning Applications to Traffic Classes” on page 149).	
Preferred Path	Select Primary or Secondary to indicate the path used for each traffic class under normal conditions.
Divert	<p>Select the conditions under which each traffic class can be switched to the alternate path:</p> <ul style="list-style-type: none">■ Never. The traffic class is never diverted from the preferred path.■ Failure Only. The traffic class is diverted to the alternate path only if the tunnel for the preferred path goes down and the tunnel for the alternate path is active.■ Congestion/Failure. The traffic class is diverted to the alternate path if the loss or latency threshold is exceeded on the preferred path or the tunnel goes down. A diversion for loss or latency occurs only if the alternate path’s loss and latency are not exceeded. <p>If Congestion/Failure is selected for any traffic class, probe packets are sent to the remote devices to measure the loss and latency of each path. To specify a latency threshold for each remote device, refer to “Defining Multi-Path Endpoints” on page 268. By default, the loss threshold is exceeded if two or more probes are lost per minute for four consecutive minutes.</p> <p>All of the threshold settings can be changed using the CLI (refer to the “configure multi-path” command).</p>



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

3. To change a template name or settings, click the template name, change the template name and/or the settings for each traffic class, and click Submit.
4. To delete a template, click the check box next to the template name, and click Submit. If a template is applied to an endpoint, it cannot be deleted.

Defining Multi-Path Endpoints

For each device that has a secondary IP address, you can select a multi-path template and supplemental marking method (if any), and specify a latency threshold for the primary and secondary paths to the device.

To specify a secondary IP address for a device, use the device Web console or load a Device Settings partial configuration on the device (refer to “Configuring Multi-Path Addresses” on page 116).

To define Multi-Path endpoints:

1. In the Configuration window, click Multi-path in the left-hand navigation frame, click Endpoints, and select the check box.

Figure 144: Defining Multi-Path Endpoints

'Main' Global Configuration Compatible with WXOS 5.4+ Show CLI Save Cancel

Global

- Basic Setup
- AAA
- Applications
- Compression
- QoS
- Acceleration
- Advanced Setup
- Multi-path**
 - ☒ Start/Stop
 - ☒ Templates
 - ☒ Endpoints

Multi-Path Endpoints HELP

Click on **Add/Remove Endpoints** to add or remove endpoints from the list. 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.

Device Name	Latency Threshold (msec)		Multi-Path Template	Supplemental Marking Method
	Primary	Second.		
SM-10.87.247.2	5000	5000	Primary_config/fail	None (Sec. IP Only)

Add/Remove Endpoints

Submit Reset

2. To add or remove remote endpoints for Multi-Path:
 - a. Click **Add/Remove Endpoints**.
 - b. Select a community from the Community/Device Group list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have their primary and secondary addresses enclosed in parentheses, separated by a slash.

 - c. Select the Multi-Path devices you want to configure, and click Add. To remove devices from the Multi-Path Endpoints list, select the devices and click Remove.

- d. Repeat Steps **b** and **c** for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
- e. If one or more devices are not listed, click Manual Entry and enter the primary and secondary IP addresses for each device, separated by a slash (one address pair per line), and click Submit.
- f. When you are done, click Submit.



NOTE: Compression is required for Multi-Path. When you save a global configuration, an error occurs if compression is disabled for an endpoint using Multi-Path. If you add an endpoint to a Multi-Path partial configuration, an error occurs if you load the configuration on a device where compression is disabled for that endpoint.

3. For each Multi-Path endpoint, specify the following:

Latency Threshold	<p>Enter the latency threshold in milliseconds (20 to 5000) for the primary and secondary paths. Traffic is switched to the alternate path when the threshold is exceeded, and is switched back when latency drops below the threshold. This setting is ignored for traffic classes where the selected template disallows switching between paths.</p> <p>NOTE: If you set the threshold too low, minor fluctuations in latency may cause constant switching between paths.</p> <p>By default, a probe tests the path 12 times per minute, and traffic is switched when the median latency exceeds the threshold for four consecutive minutes. Traffic is also switched if two or more probes are lost per minute for four consecutive minutes. To change these settings, refer to the “configure multi-path” command</p>
Multi-Path Template	<p>Select a template for this endpoint that specifies the preferred path and the conditions under which traffic can be switched to the alternate path. To add a new template, refer to “Defining Multi-Path Templates” on page 266.</p>
Supplemental Marking Method	<p>Optionally, select one of the additional marking methods for the packets sent on each path (refer to “Configuring Multi-Path Addresses” on page 116 and “Enabling Policy-Based Multi-Path” on page 265).</p> <p>By default, all packets to be sent on the secondary path have the source address set to the secondary IP address.</p>

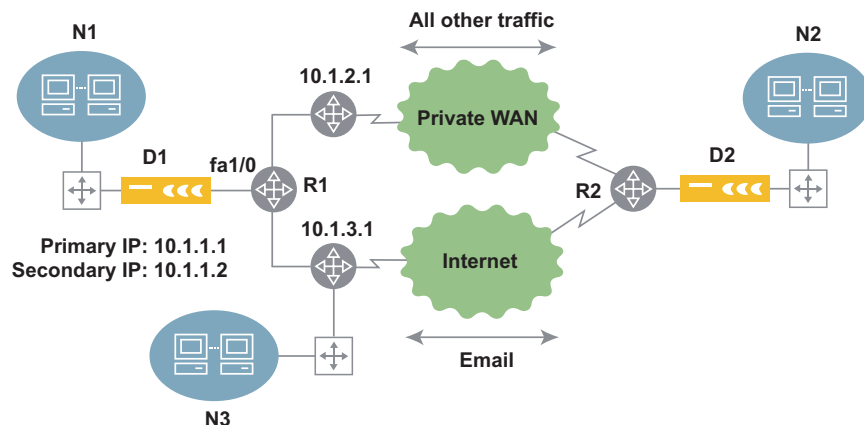
4. Click Submit to enter the changes, or click Reset to discard them.

To view the status of the primary and secondary paths from a specific device, access the device Web console and open the Multi-Path Endpoints page and the Multi-Path monitoring report.

Configuring Routers to Support Multi-Path

You can configure a WAN router to select a gateway for multi-path traffic based on the source IP address, or based on the source address and a ToS or DSCP value. The following configuration examples apply to router R1 in Figure 145. A similar configuration is needed for R2.

Figure 145: Multi-Path Router Configuration Example



To configure the WAN router R1 to use only the source IP address:

1. On the inbound interface from the WX device, define a route map for Multi-Path. For example:

```
interface FastEthernet 1/0
  ip address 10.1.1.5 255.255.255.0
  ip policy route-map mpath
```

2. Define access lists for the primary and secondary source IP addresses. For example:

```
access-list 50 permit 10.1.1.1
access-list 51 permit 10.1.1.2
```

3. Match the primary and secondary source IP addresses with the appropriate primary and secondary gateways. For example:

```
route-map mpath permit 10
  match ip address 50
  set ip next-hop 10.1.2.1
```

```
route-map mpath permit 20
  match ip address 51
  set ip next-hop 10.1.3.1
```

To configure R2, use the commands above, but change the interface address and use the primary and secondary address for device D2.

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

1. Define a route map for Multi-Path (see the previous example).
2. Define extended access lists for the primary and secondary source IP addresses and their associated IP precedence or DSCP values. For example, for IP precedence values:

```
access-list 100 permit ip host 10.1.1.1 any precedence 10
access-list 101 permit ip host 10.1.1.2 any precedence 11
```

For DSCP values:

```
access-list 100 permit ip host 10.1.1.1 any dscp 1
access-list 101 permit ip host 10.1.1.2 any dscp 2
```

3. Match the primary and secondary source IP addresses with the appropriate primary and secondary gateways. For example:

```
route-map mpath permit 10
  match ip address 100
  set ip next-hop 10.1.2.1
```

```
route-map mpath permit 20
  match ip address 101
  set ip next-hop 10.1.3.1
```



NOTE: Unless you use a console server to manage devices, you may need to change the access lists to allow management access from some locations using SSH or Web/SSL. For example, in Figure 145, 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 IPSec

IPSec can be used to authenticate and encrypt traffic between any pair of WX devices in the same community. Enabling IPSec allows you to:

- Compress traffic before it is encrypted (encrypted traffic cannot be compressed).
- Encrypt traffic over unprotected networks, such as the Internet.

To configure IPSec, you define templates that specify the security algorithms and key lifetimes for outgoing traffic, and then apply a template to each of the remote endpoints that act as IPSec peers. For a pair of devices to use IPSec, IPSec must be enabled on both devices, and both devices must be configured with the same pass phrase (preshared key) and security algorithms. Each device can encrypt traffic for up to 100 remote peers (the WX 20 is limited to five devices).

The following sections describe how to configure IP security (IPSec) to authenticate and encrypt traffic between any pair of devices:

- “Default IPSec Policy” in the next section
- “IPSec Implementation Details” on page 273
- “Procedure for Configuring IPSec Policies” on page 274
- “Defining IPSec Settings by Endpoint” on page 275
- “Defining IPSec Templates” on page 277
- “Defining the Default IPSec Policy” on page 279

Default IPSec Policy

When two devices are configured as IPSec peers, all compressed and passthrough traffic sent between them is encrypted. For passthrough traffic destined for subnets that are not served by a WX device, a “default IPSec policy” is provided that lets you specify, by subnet, whether the traffic is dropped and logged or sent unencrypted. Initially, the default IPSec policy allows all traffic to be sent unencrypted.

The default IPSec policy also applies to traffic between peer devices where IPSec is enabled, but the key negotiation has failed. Note that an IPSec-enabled device never encrypts traffic destined for a remote device where IPSec is disabled.

After you verify that IPSec is working correctly, all subnets advertised by IPSec-enabled peers should be added to the encryption-required list to avoid sending unencrypted traffic to those subnets if a remote device fails.



NOTE: If an inline device fails, all traffic is passed through without encryption. To block all traffic during a hardware failure, use a cross-over cable (rather than a straight-through cable) to connect the device to the WAN router. This works only if Ethernet auto-MDI negotiation is disabled on the router.

IPSec Implementation Details

IPSec is implemented in compliance with RFCs 2401-2409, and includes the following:

- Encryption algorithms—Advanced Encryption Standard (AES) encryption algorithm, with 128, 192, and 256 bit keys, and Triple DES (3DES)
- Authentication algorithms—HMAC/SHA-1 and HMAC/MD5
- Internet Key Exchange (IKE) protocol for dynamic key exchange
- Encapsulated Security Protocol (ESP) in transport mode used for all encrypted packets

AES with a 256 bit key and HMAC/SHA-1 authentication provides the highest security, while AES with a 128 bit key and HMAC/MD5 authentication provides the highest throughput (primarily because SHA-1 is two to three times slower than MD5). 3DES is supported for environments where AES is not approved, but 3DES is slower and less secure than AES, and is not recommended.

Although the IPSec protocols allow two peers to communicate using different policies, such as having Peer 1 use AES to encrypt for Peer 2, while Peer 2 uses DES to encrypt for Peer 1, both WX devices must use the same encryption and authentication algorithms.

Using IPSec allows you to compress traffic before encrypting it (encrypted traffic cannot be compressed because it contains few recognizable patterns). Since outgoing traffic is both compressed and encrypted, 3rd party IPSec devices cannot support our implementation because they cannot decompress the traffic. However, uncompressed IPSec traffic has been validated against Cisco and Microsoft IPSec implementations to ensure IPSec compliance.



NOTE: The IPSec Authentication Header (AH) is not used, and only Diffie-Hellman Group 5 is supported.

Procedure for Configuring IPSec Policies

To encrypt the traffic sent between two or more devices:

1. Select the devices that you want to encrypt traffic and specify the pass phrase(s) needed to establish a secure connection (refer to “Defining IPSec Settings by Endpoint” on page 275).
2. To change the default Wizard template or define new templates, refer to “Defining IPSec Templates” on page 277.
3. To change the default IPSec policy, refer to “Defining the Default IPSec Policy” on page 279.

Alternatively, you can run the IPSec Setup Wizard on each device from the device Web console.

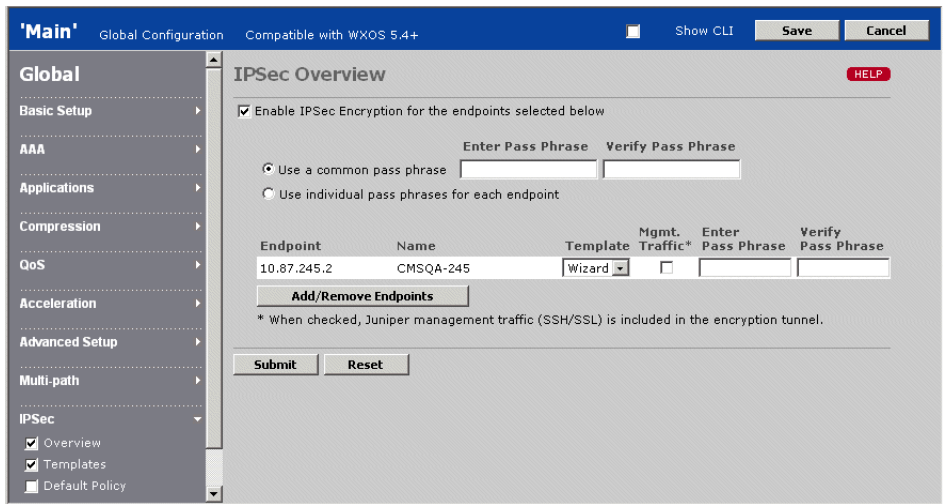
Defining IPSec Settings by Endpoint

On the IPSec Overview page, you can enable or disable IPSec for all endpoints or specific endpoints, change the IPSec template or pass phrase for an endpoint, or enable encryption for management traffic. To add or change IPSec templates, refer to “Defining IPSec Templates” on page 277.

To view or change the IPSec settings by endpoint:

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

Figure 146: IPSec Overview



- 2. To enable IPSec, click the check box next to Enable IPSec Encryption.
- 3. To add or remove remote endpoints for IPSec:
 - a. Click Add/Remove Endpoints.
 - b. Select a community from the Community/Device Group list. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses.

Devices that support Multi-Path have two separate entries for the primary and secondary IP address, which correspond to the primary and secondary paths. You can enable IPSec for one or both paths. To configure Multi-Path, refer to “Configuring Multi-Path Addresses” on page 116.

- c. Select the devices you want to configure, and click Add. To remove devices from the IPSec Endpoints list, select the devices and click Remove. If you remove an endpoint, all subsequent traffic to that endpoint is sent unencrypted.
 - d. Repeat Steps b and c for each community/device group (some devices may belong to multiple communities or groups). When you download the configuration, any devices or communities that do not apply to a device are ignored.
 - e. If one or more devices are not listed, click Manual Entry and enter the primary or secondary IP addresses for each device (one per line), and click Submit.
 - f. When you are done, click Submit.
4. Enter and verify a pass phrase for each endpoint, or select Use a common pass phrase and enter one pass phrase for all endpoints (eight or more characters is recommended). The pass phrase is used to generate a preshared key of the appropriate length.
 5. To change the template for an endpoint, select a template from the Template drop-down list. To create new templates, refer to “Defining IPSec Templates” on page 277. Two endpoints can establish a secure connection only if their IPSec templates specify the same authentication and encryption algorithms. The default Wizard template uses AES-128 and HMAC/SHA-1.
 6. To encrypt all management traffic sent to a remote endpoint, including SNMP, Syslog, and registration server traffic, click the Mgmt. Traffic check box for the endpoint. Encrypting management traffic is recommended after you verify that the IPSec connection is operating normally.

To view the status of the IPSec connections from a specific device, access the device Web console and open the IPSec Overview page.

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

Defining IPSec Templates

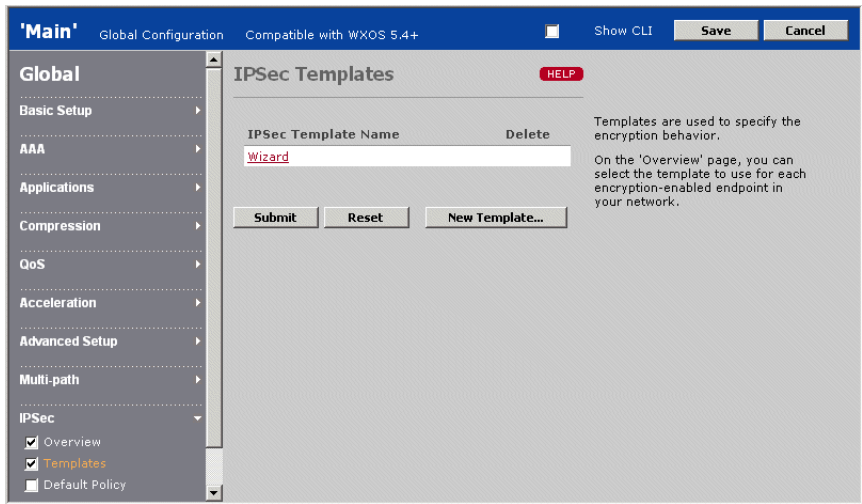
IPSec templates specify the algorithms used to protect traffic between endpoints, and the lifetime of each generated key. You can change the default Wizard template or create new templates. The default Wizard template uses the AES-128 and HMAC/SHA-1 encryption and authentication algorithms, and the generated keys do not expire.

To apply a template to an endpoint, refer to “Defining IPSec Settings by Endpoint” on page 275.

To define IPSec templates:

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

Figure 147: Defining IPSec Templates



2. To add a new template, click New Template, specify the following information, and click Submit:

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

3. To change a template name or settings, click the template name, change the template, and click Submit.
4. To delete a template, click the check box next to the template name, and click Submit. If you load the configuration on a device where the deleted template is applied to an endpoint, the endpoint reverts to the Wizard template.

Defining the Default IPSec Policy

The default IPSec policy is applied to the following types of traffic:

- Passthrough traffic sent to unadvertised subnets
- Traffic between devices where IPSec is enabled, but the key negotiation has failed

By default, all such traffic is unencrypted. However, you can change the default policy so that traffic to specific destinations is dropped and logged, rather than sent unencrypted. The number of packets dropped for each destination is written to the system log every five minutes. Use the device Web console to view the system log.

After you verify that IPSec is working correctly, all subnets advertised by IPSec-enabled peers should be added to the encryption-required list to avoid sending unencrypted traffic to those subnets if a remote device fails.



NOTE: All passthrough traffic between IPSec-enabled devices is encrypted. For example, traffic is encrypted even when compression is disabled.

To change the default IPSec policy:

1. In the Configuration window, click IPSec in the left-hand navigation frame, click Default Policy, and select the check box.

Figure 148: Defining the IPSec Default Policy

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

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

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

Configuring Events

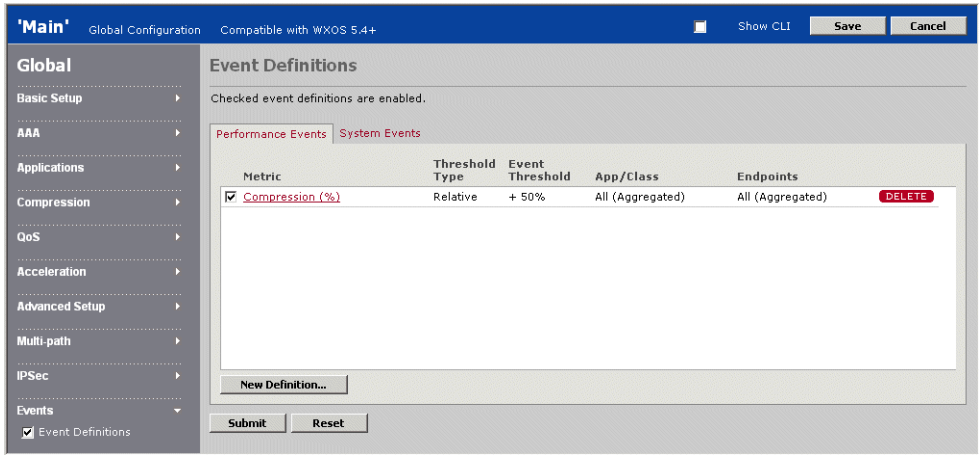
For WXOS 5.4 devices, you can configure performance thresholds for compression, acceleration, throughput, and dropped traffic so that events are triggered when the average performance for the previous hour or day exceeds (or drops below) the specified threshold. You can also enable or disable the generation 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 Reports” on page 362. Note that the WXOS 5.4 devices must specify the CMS server as a Syslog server (refer to “Defining Syslog Servers” on page 124). Performance events are also sent to the defined SNMP trap destinations (refer to “Enabling SNMP” on page 123).

To configure events:

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

Figure 149: 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 17 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>

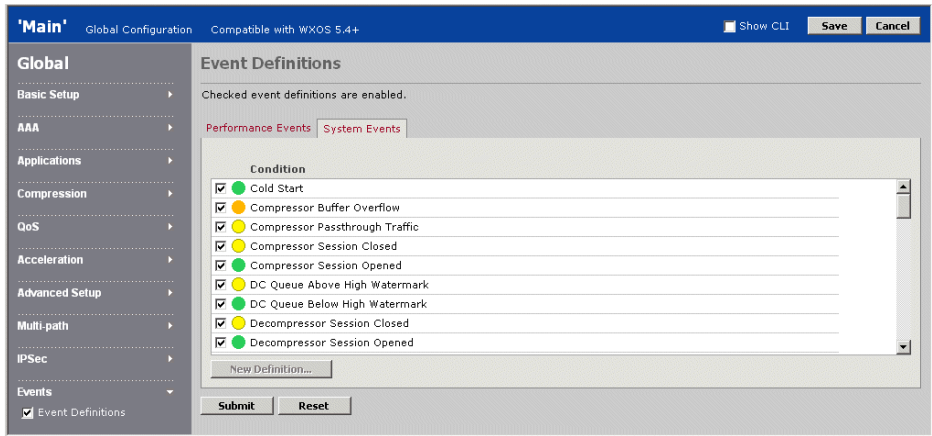
Destination	<p>Click Edit and specify which endpoints can generate the event:</p> <ul style="list-style-type: none"> ■ All (Aggregated). An event occurs based on the overall performance of all remote endpoints (the default). ■ Any. An event occurs if any endpoint violates the threshold. ■ Select WX device. Select a community or device group, and select a specific device name. An event occurs only if the selected WX device violates the threshold. ■ Enter WX Device Address or non-WX Endpoint Name. Enter the IP address of a WX device not listed above or the name of a non-WX endpoint, such as “Other Traffic”. Non-WX endpoints can generate events only for the WAN Throughput, QoS Throughput, Bytes Dropped, and Packets Dropped metrics. Note that the “Other Traffic” endpoint includes only the traffic that is not sent to the WX or 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 241).</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	<p>Select the Yes check box to enable monitoring for this event definition.</p>

Table 17: 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: $(\text{AccelerationFactor} * 100) - 100$ <p>For example, if the acceleration factor is 3, the acceleration percentage is 200 %.</p>
TCP Acceleration Throughput In (Kbps)	Average throughput for all TCP accelerated sessions in the period. $(\text{BytesIn-over-all-sessions} * 8) / (\text{Active-transaction-time-over-all-sessions})$ <p>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.</p>
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 Out (count)	Absolute number of bytes dropped in the period.
Packets Dropped Out (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 150: 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 “WX System Events and SNMP Traps” on page 442.
 - 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.

Chapter 5

Automatic Deployment and License Management

This chapter describes how to use CMS to automatically download configurations and WXOS software to new WX devices, and to generate permanent licenses for devices that need them. It covers the following topics:

- “About Automatic Deployment” in the next section
- “Configuring Auto-Deployment” on page 288
- “Configuring License Management” on page 296

About Automatic Deployment

The first time a WX device running WXOS 5.x (or later) is started, it attempts to contact the CMS server. If you know the subnet where the device is installed, you can configure CMS to download a configuration and an WXOS software image to the new device. On-site personnel simply connect the cables and apply power, and the device becomes operational.

After a successful auto-deployment, you can generate a permanent license for the device (refer to “Configuring License Management” on page 296).

Auto-deployment has two requirements:

- A DHCP server must be reachable from the WX device. When first powered on, the device sends DHCP requests over its Local and Remote interfaces. The DHCP server must reply with an IP address and the address of one or more DNS servers. Up to three DNS servers will be queried.

- For WXOS 5.4 devices, one of the three DNS servers must have an entry for “juniper-cms” in the domain hierarchy (WXOS 5.2/5.3 devices look for “peribit-cms”). For example, if the domain name in the DHCP reply is “sales.company.com”, a WXOS 5.4 device issues DNS requests in the following order to locate the CMS server:

- juniper-cms.sales.company.com
- juniper-cms.company.com
- juniper-cms.com
- juniper-cms

If DHCP does not specify a domain, and a reverse lookup on the DNS server's address does not obtain one, then “juniper-cms” is the only request.

After obtaining the IP address of the CMS server, the device contacts the server over HTTPS. CMS can then download the prepared configuration and software image based on the device's subnet.



NOTE: Only one device can be auto-deployed per subnet. For example, a multi-node configuration, where two devices are connected together, cannot be auto-deployed. Also, an WX 100 can be auto-deployed, but its client devices must be configured locally.

Configuring Auto-Deployment

The following topics describe how to configure auto-deployment:

- “Auto-Deployment Procedure” in the next section
- “Defining Deployment Groups” on page 290
- “Defining Deployment Records” on page 292
- “Viewing the Auto-Deployment Status” on page 294

Auto-Deployment Procedure

Use the following procedure to configure auto-deployment:

1. Prepare full configurations to be downloaded to the auto-deployed devices. You can load a global configuration, a full set of partial configurations, or a combination of both (refer to “Defining Configuration Settings” on page 98).

For example, in a hub and spoke environment, you might create a global configuration for the spokes, and partial configurations that override the topology and other settings for the hubs (refer to “Configuring Topology Settings” on page 229).

2. Define deployment groups that specify the configuration and software image to be downloaded (refer to “Defining Deployment Groups” on page 290).
3. Define a deployment record for each auto-deployed device that specifies the device’s subnet, deployment group, and other settings (refer to “Defining Deployment Records” on page 292).
4. Monitor the status of the auto-deployed devices (refer to “Viewing the Auto-Deployment Status” on page 294).
5. When the deployment is complete, verify that the communities for the auto-deployed devices have been imported from the registration server(s) specified in the device configurations. Initially, all devices are in the Default community (refer to “Importing and Managing Communities” on page 397).
6. Configure licenses for the auto-deployed devices (refer to “Configuring License Management” on page 296).

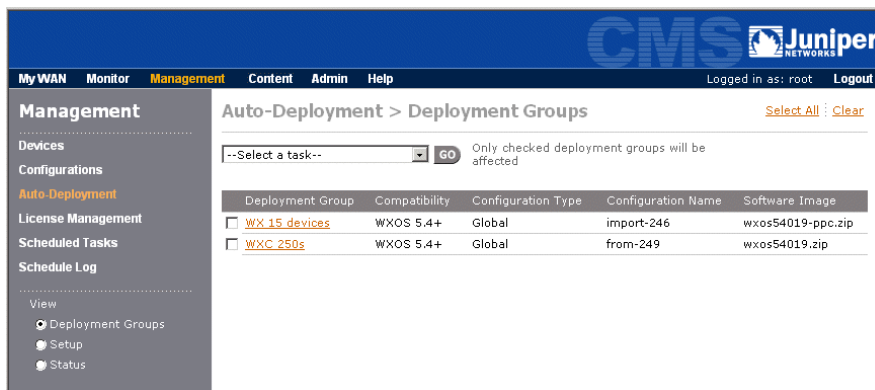
Defining Deployment Groups

After you prepare configurations for the WX devices to be auto-deployed, you must define at least one deployment group. A deployment group specifies the global and/or partial configurations that you want to download to one or more auto-deployed devices. Optionally, you can also specify an WXOS software image to be loaded at the same time.

To define deployment groups:

1. Click Management in the menu frame, and then click Auto-Deployment in the left-hand navigation frame.

Figure 151: Defining Deployment Groups



The Deployment Groups page lists the device compatibility (WXOS 5.2/5.3 or 5.4), configuration(s), and software image specified by each deployment group.

2. To change a deployment group, click the name, make any needed changes and click Submit.
3. To define a new deployment group, select one of the following from the Task menu and click Go:
 - New (5.2/5.3 compatible)
 - New (5.4 + compatible)

Figure 152: Adding a Deployment Group

My WAN Monitor **Management** Content Admin Help

Logged in as: root Logout

Management

- Devices
- Configurations
- Auto-Deployment**
- License Management
- Scheduled Tasks
- Schedule Log

View

- ☒ Deployment Groups
- ☐ Setup
- ☐ Status

Auto-Deployment > Deployment Groups > New

Load the following configuration to the devices assigned to this deployment group.

Deployment Group Name

Global Configuration

☐ Do not load global configuration

☒ factory-54 [History](#)

Partial Configurations

AAA	--	History
Acceleration	--	History
Advanced Setup	--	History
Applications	--	History
Basic Setup	--	History
Event Definitions	--	History
IPSEC	--	History
Multi-Path	--	History
QoS	--	History
Reduction	--	History

Software Image

☐ Do not load image

☒ wxos54019-ppc.zip (for WX15 only)

4. Specify the following information:

Deployment Group Name	Enter a name for the deployment group.
Global Configuration	Select a global configuration or, to load only partial configurations, select Do not load global configuration. Click History to view the selected configuration and its past changes. To create global configurations, refer to “Managing Configurations” on page 85.
Partial Configurations	<p>If you are not loading a global configuration, you must select one of each type of partial configuration. The settings in each partial configuration replace the corresponding settings in the selected global configuration (if any) or are combined into one configuration.</p> <p>Click History to view each selected configuration and its past changes.</p>
Software Image	Optionally, select a WXOS software image to be loaded. The image must first be loaded on the CMS server (refer to “Uploading a Boot Image” on page 408).

- Click Preview to view the resulting configuration. Any settings that are not defined in the global and partial configurations will remain in the factory default state on the device.

All configuration settings are saved as CLI commands. For descriptions of each CLI command, refer to the *WX/WXC Operator's Guide*.

- Click Submit to enter the changes, or click Cancel to discard them.

Defining Deployment Records

After you define the appropriate deployment groups, you must create a deployment record for each device to be auto-deployed. Each deployment record specifies the subnet where the device is installed, various network settings for the device, and a deployment group.

To define deployment records:

- Click Management in the menu frame, click Auto-Deployment in the left-hand navigation frame, and then click Setup.

Figure 153: Defining Auto-Deployment Records

The screenshot displays the 'Auto-Deployment > Setup' page in the Juniper CMS. The left-hand navigation menu is expanded to 'Management', and 'Auto-Deployment' is selected. The main content area shows a table of deployment records. The table has the following columns: Originating Subnet, Static IP Addr., Subnet Mask, Gateway, Device Settings Partial Configuration, and Deployment Group. A single record is listed with the following values: Originating Subnet: 1.1.1.0/24, Static IP Addr.: 1.1.1.5, Subnet Mask: /24, Gateway: 1.1.1.1, Device Settings Partial Configuration: AD1_1_1_5, and Deployment Group: Deployment_group1. Below the table, there are input fields for creating a new record. The fields are: Originating Subnet (1.1.1.0/24), Static IP Addr. (1.1.1.5), Subnet Mask (/24), Gateway (1.1.1.1), Device Settings Partial Configuration (--Create one--), and Deployment Group (Deployment_group1). The Time Zone is set to (GMT -08:00) Pacific Time (US and Canada), Tijuana. There are buttons for Add, Update, Delete, and Submit. A note at the bottom says 'When you are done, click "Submit" to save any changes.'

The Deployment Setup page lists the properties of each deployment record.

2. To add a new deployment record, specify the following information:

Originating subnet	Enter the subnet where a new device is (or will be) installed. The format is: subnet/mask where <i>mask</i> is the number of binary digits used for the network portion of the address.
Static IP Addr.	Enter a static IP address for the new device. It need not be in the originating subnet.
Subnet Mask	Enter the number of binary digits used for the network portion of the static address. The format is: /mask
Gateway	Enter the IP address of the default gateway for the device. It must be in the same subnet as the static IP address.
Device Settings Partial Configuration	If you created a Device Settings partial configuration for the device, you can select it here. The default setting (--Create one--) generates a Device Settings partial configuration named: AD<IP address> where the dots in the static IP address are replaced by underscores. This partial configuration specifies only the settings in the deployment record (static address, subnet mask, gateway address, time zone, and daylight savings indicator).
Deployment Group	Select a deployment group that specifies the configuration(s) to be loaded on the device. To add deployment groups, refer to “Defining Deployment Groups” on page 290.
Time Zone	Select the time zone of the device.
Daylight Saving	Select the check box to enable Daylight Savings Time on the device (if applicable).

3. When you are done, click Add, and click Submit.

To add a new record by copying and modifying an existing record, select the check box next to the subnet for the record you want to copy, and then clear the check box. You can then change the copied record, click Add, and click Submit.

4. To change a deployment record, click the check box next to the subnet, make any needed changes, click Update, and then click Submit.

5. After you click Submit, you can leave the page and return later to add or edit deployment records. You can complete the deployment records over time as you establish the required network information for each device to be auto-deployed.
6. When a deployment record is complete, click the Ready to Deploy check box, and click Submit. The configuration and software image (if any) are now ready to be downloaded when the device checks in. To view the status of the deployment, refer to “Viewing the Auto-Deployment Status” on page 294.

Note that the check box next to the subnet is greyed out. To make any additional changes to the record, you must first clear the Ready to Deploy check box, and click Submit.

Viewing the Auto-Deployment Status

After a deployment record is defined and marked “Ready to Deploy,” you can monitor the status of the auto-deployment to see when the device checks in and whether the deployment is successful.

To view the auto-deployment status:

1. Click Management in the menu frame, click Auto-Deployment in the left-hand navigation frame, and then click Status.

Figure 154: Viewing Auto-Deployment Status

The screenshot displays the CMS interface for viewing auto-deployment status. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (selected), 'Content', 'Admin', and 'Help'. The left sidebar shows 'Management' with sub-items: 'Devices', 'Configurations', 'Auto-Deployment' (selected), 'License Management', 'Scheduled Tasks', and 'Schedule Log'. The main content area is titled 'Auto-Deployment > Status' and includes a 'Refresh' link. A table lists deployment records with the following data:

IP Address	Originating Subnet	Deployment Attempts	Last Attempt	Status
192.168.52.200	192.168.52.192/28	1	2006-08-20 19:18:38.0	Successful
192.168.53.5	192.168.53.0/24	1	2006-08-19 18:42:20.0	Successful
192.168.53.140	192.168.53.128/26	1	2006-08-24 16:55:37.0	In progress
Unknown	10.10.2.0/24	0		

Below the table, there is a 'Remove' button and a note: 'Click this button to remove successfully deployed devices from the table'. A footer note states: 'IP addresses are unknown until the first deployment attempt. * DHCP-assigned IP address'.

2. The following information is provided for each deployment record marked “Ready to Deploy”:

IP Address	<p>The IP address shown for the device depends on the status of the deployment:</p> <ul style="list-style-type: none">■ Unknown. Device has not checked in.■ < DHCP address > . Deployment is in progress.■ < Static address > . Deployment successful.
Originating Subnet	<p>The subnet where the device is installed (specified by the deployment record).</p>
Deployment Attempts	<p>Number of times the deployment has been attempted. After five failed attempts, subsequent requests from the device are rejected. To allow another five attempts, you must the reset the “Ready to Deploy” flag on the deployment record (refer to “Defining Deployment Records” on page 292).</p>
Last Attempt	<p>Date and time of the last deployment attempt.</p>
Status	<p>Indicates the status of the auto-deployment:</p> <ul style="list-style-type: none">■ Blank. Device has not checked in.■ In Progress. Deployment is in progress.■ Successful. The configuration and software image (if any) were successfully downloaded to the device.■ Failed. The last deployment attempt has failed. <p>Click the status for more details, such as the device type and MAC addresses.</p>

3. Click Remove to remove the status entries for successful deployments (the corresponding deployment records are deleted automatically).



NOTE: You can auto-deploy a device only once. If an auto-deployed device is reset to the factory defaults, its attempts to contact the CMS server will be rejected.

Configuring License Management

The following topics describe how to configure the bulk deployment of new WXOS licenses. For an upgrade license, contact your sales representative.

- “Licensing Procedure” in the next section
- “Importing and Validating Authorization Codes” on page 297
- “Generating and Applying Licenses” on page 299
- “Viewing the License Status” on page 302



NOTE: Bulk deployment of licenses from CMS works only for the initial permanent license applied to a device. This feature cannot be used to upgrade permanent licenses.

Licensing Procedure

Use the following procedure to apply permanent WXOS licenses to devices that have evaluation licenses:



NOTE: You must have an account on the Juniper Customer Support Center (<http://www.juniper.net/customers/support>), and the CMS server must be able to establish an HTTP connection with the license server at <http://license.peribit.com>.

1. Create a file of Authorization Codes, and import the file into CMS (refer to “Importing and Validating Authorization Codes” on page 297).
2. Match the Authorization Codes with the devices that have evaluation licenses, generate licenses for the matching devices, and then apply the licenses (refer to “Generating and Applying Licenses” on page 299).
3. Monitor the status of deployed licenses (refer to “Viewing the License Status” on page 302).

Importing and Validating Authorization Codes

When you purchase a software license from Juniper Networks you receive an Authorization Code certificate in PDF format that lists the Authorization Codes used to produce licenses. The WXOS Authorization Code certificates have a section at the end of the document that provides the Authorization Codes in a format that can be imported into WX CMS.

To import and validate the Authorization Codes:

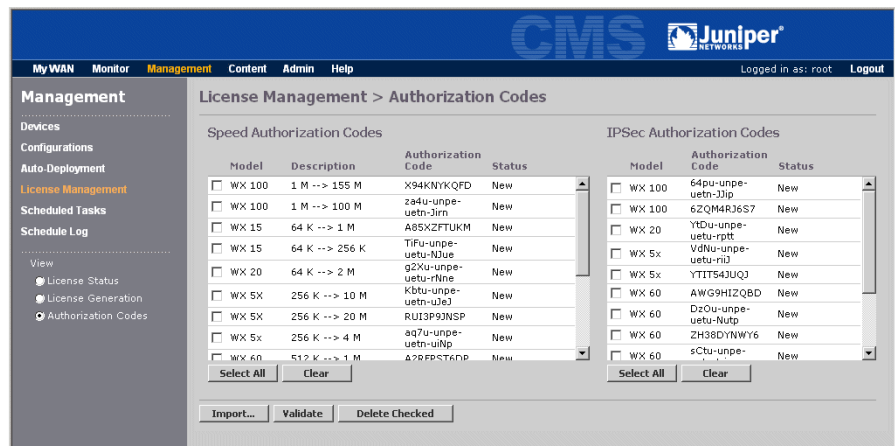
1. Use the Acrobat Text tool to copy and paste the Authorization Codes into a text file (one Authorization Code per line), and save the file in a location accessible from the browser.

The file you create can have any extension, but it must be a simple text file (for example, a “.txt” file created from NotePad, or a “.csv” file created from Excel), Do NOT use a binary file, such as a “.doc” file created from Word or a “.xls” file created from Excel.

2. Click Management in the menu frame, click License Management in the left-hand navigation frame, and then click Authorization Codes.
3. Click Import, enter the Authorization Codes file location or click Browse to locate the file, and then click Import. Authorization Codes that have already been imported are marked as duplicates and excluded automatically. If format errors are displayed, contact the Juniper Technical Assistance Center.

Click Back to import another file, or click Authorization Codes in the navigation frame to view the Authorization Codes that were added to the database.

Figure 155: Importing and Validating Authorization Codes



The following information is shown for each speed and IPSec Authorization Code:

Model	Device type, such as WX 80. An “WX 5x” indicates the Authorization Code can be applied to an WX 50 or WX 55.
Description	Indicates the base and maximum device speed (speed Authorization Codes only).
Authorization Code	Text identifying the Authorization Code (internal use only).
Status	Indicates the Authorization Code status: <ul style="list-style-type: none"> ■ New. Initial status of all imported Authorization Codes that have not been validated. ■ Valid. Validated by the License Server, but not yet assigned to a device. ■ Invalid. Not recognized by the License Server (contact Technical Support). ■ Canceled. No longer valid (contact Technical Support). ■ Temp-Assigned. Matched with a device, but not yet used to generate a license (refer to “Generating and Applying Licenses” on page 299). ■ Perm-Assigned. License generation has started or is complete, so the Authorization Code cannot be assigned to another device. ■ Used. Has been used to generate a license. This status may be set by the License Server when you validate the Authorization Codes. If a New Authorization Code is set to Used, and you have not used it to generate a license, contact Technical Support.

4. Select the check box next to the speed and IPSec Authorization Codes that you want to validate, or click Select All, and then click Validate.

The status for all New Authorization Codes should be changed to Valid. If any new Authorization Codes are set to invalid or canceled, contact Technical Support.

5. To delete the Used Authorization Codes, select the check box next to the appropriate codes, and click Delete.

You can now use the valid Authorization Codes to generate and apply licenses to your devices, as described in the next section.

Generating and Applying Licenses

After you import and validate your Authorization Codes, you can match them with the deployed devices that have evaluation licenses, generate permanent licenses, and then apply the licenses to each device.



NOTE: You must have an account on the Juniper Customer Support Center to generate the licenses.

To generate and apply licenses to your devices:

1. Click Management in the menu frame, click License Management in the left-hand navigation frame, and then click License Generation.
2. To view the devices in a community or device group that have an evaluation license (active or expired), select a community or device group and click Submit. It may take a few minutes to poll a large community. The page displays the progress of the poll (polling continues if you leave the page).

To stop the polling and select another community/device group, click Stop.

Figure 156: Generating and Applying Licenses

Juniper WX CMS automatically matches deployed WX devices with available authorization codes. To exclude a device from the matching process, uncheck it. Then, click "Match Authorization Codes" to view the updated match results. If you are satisfied with the results, click "Generate Licenses" to create licenses based on the matched authorization codes. Click "Apply Licenses" to apply the generated licenses to the corresponding WX devices.

Device Name	Serial No.	Model	Authorization Codes		Permanent License Status		
			Speed	IPSec	Authorization Code Match	Generate	Apply
<input checked="" type="checkbox"/> SR-10.87.0.27	0500002056	WXC 500	---	---	Not Ready	N/A	N/A
<input checked="" type="checkbox"/> SR-10.87.41.2	0050003173	WX 5x	---	---	Not Ready	N/A	N/A
<input checked="" type="checkbox"/> SR-10.87.42.2	0500001995	WXC 500	---	---	Not Ready	N/A	N/A
<input checked="" type="checkbox"/> SR-10.87.45.2	0050003072	WX 5x	---	---	Not Ready	N/A	N/A

Available Valid Speed Authorization Codes		Available Valid IPSec Authorization Codes	
WX 15 (64 K base)	Quantity	Device Model No.	Quantity
Valid authorization codes are not available		Valid authorization codes are not available	
WX 20 (64 K base)	Quantity		
Valid authorization codes are not available			
WX 5x (256 K base)	Quantity		
Valid authorization codes are not available			

The devices with evaluation licenses are listed (if any), followed by the imported speed and IPSec Authorization Codes that are available to be matched with a device. If a device could not be reached, its serial number and model are displayed as "Unknown."



NOTE: The speed Authorization Codes show the new device names (an SR is a "WX" and an SM is a "WXC"), but they can still be applied to SRs and SMs.

To poll another community or device group, click Change in the upper-right corner of the page.

3. To match the available Authorization Codes with the listed devices, you must change the default selection (---) for both the Speed and IPSec fields for each device, and click Match Authorization Codes.

Speed Authorization Codes	Select the speed Authorization Code that you want assigned to the device. Select None to generate a license for the base speed (no Authorization Code required). The base speeds for each device type are shown in parentheses below the device list.
IPSec Authorization Codes	Select whether you want an IPSec Authorization Code assigned to the device (Yes or No).

If you do not have enough Authorization Codes for all devices, clear the check box next to the less-critical devices, and click Match Authorization Codes again. This runs the match for just the selected devices. You can run the match as often as needed.

4. The following licensing information is shown for each device:

Authorization Code Match	<p>Indicates whether an imported Authorization Code matched the device:</p> <ul style="list-style-type: none">■ Excluded. The device is excluded from the matching process (the check box is not selected).■ Not Ready. The speed and/or IPSec Authorization Code have not been selected.■ Successful. Imported Authorization Codes matched the device. The list of available Authorization Code is adjusted accordingly.■ Unavailable. No match for one or both of the selected Authorization Code (unmatched Authorization Code are highlighted in yellow).
Generate	<p>Indicates whether a license has been generated:</p> <ul style="list-style-type: none">■ N/A. No attempt made.■ Successful. License has been generated.■ Failed. License generation failed (contact Technical Support).
Apply	<p>Indicates whether a license has been applied:</p> <ul style="list-style-type: none">■ N/A. No attempt made.■ Successful. License has been applied.■ Failed. License could not be applied. Verify that the device is reachable and try again. If the problem persists, contact Technical Support.

5. To generate licenses for devices that have a successful Authorization Code match:
 - a. Click Generate Licenses.
 - b. Enter the user name and password for your customer account on the License Server, and click Submit. Enter the requested information in all of the fields, and click Submit.
 - c. License generation begins. The Generate column indicates the success or failure of the license generation for each device.
6. When license generation is complete, click Apply Licenses to download the successfully generated licenses to each device. If the last attempt to apply a license failed, the CMS server tries to apply the license again.

Applying the licenses may take some time. You can view the status for each device on the License Status page, as described in the next section.

Viewing the License Status

For each device for which you have successfully generated a license, the License Status page shows the number of attempts to apply the license to the device (if any), and the results of the last attempt.

To view the license status:

1. Click Management in the menu frame, click License Management in the left-hand navigation frame, and then click License Status.

Figure 157: Viewing the License Status

The screenshot shows the CMS (Juniper Networks) interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management' (highlighted), 'Content', 'Admin', and 'Help'. The right side of the bar shows 'Logged in as: root' and a 'Logout' link. The left-hand navigation menu under 'Management' includes 'Devices', 'Configurations', 'Auto-Deployment', 'License Management' (highlighted), 'Scheduled Tasks', and 'Schedule Log'. The main content area is titled 'License Management > License Status' and contains a table with the following data:

Device Name	Download Attempts	Last Attempt	Status	Description
SR-192.168.52.199	1	2006-08-18 16:49:44.0	License Application Successful	
SR-192.168.52.200	1	2006-08-24 14:13:11.0	License Application Successful	
SR-192.168.53.5	1	2006-08-20 12:36:25.0	License Application Successful	
SR-192.168.53.180	1	2006-08-23 19:40:29.0	License Application Successful	
SR-192.168.53.181	0	2006-08-24 15:19:19.0	License Generation Successful	
SR-192.168.55.200	1	2006-08-18 17:34:20.0	License Application Successful	

Below the table, there is a 'Remove' button and a note: 'Click this button to remove devices that have licenses applied successfully'.

2. The following information is provided for each device:

Device Name	Name of the device.
Download Attempts	Number of attempts to apply the license to the device.
Last Attempt	Date and time of the last attempt to apply the license.
Status	Indicates one of the following: <ul style="list-style-type: none">■ License generated. No attempt to apply the license.■ License applied. License applied successfully.■ License application failed. Last attempt to apply the license failed.
Description	Provides additional information if the license application fails. The most common problems are: <ul style="list-style-type: none">■ AUTH_FAILURE. The device belongs to a community that has not been imported. To import the community, refer to “Importing and Managing Communities” on page 397.■ CONNECT_TIMEOUT or CONNECT_FAILURE. Network problem or the device may be down. For other types of errors, contact Technical Support.

3. Click Remove to remove the status entries for the licenses that were applied successfully.

Chapter 6

Monitoring Performance

This chapter describes how to use CMS to monitor the device performance. It covers the following topics:

- “Viewing and Printing Reports” in the next section
- “Configuring the My WAN Page” on page 306
- “Viewing Reports on the Monitor Page” on page 311
- “Scheduling Reports” on page 368
- “Managing Scheduled Reports” on page 369

Viewing and Printing Reports

Note the following about viewing and printing reports:


- Most reports are generated from a local database populated by periodic polling of the WXOS 5.x devices, and report times are based on the local server time, not the device time. On device-specific reports, you can select one of the following options:
 - **Show in Device Time.** Shows the report in the device’s time (for accurate results the device’s time zone must be set correctly).
 - **Show in Destination Time.** Shows the report in the destination device’s time. Applies only to CIFS, Exchange, and HTTP acceleration reports where a specific destination device is selected (acceleration is measured from the destination device).

For example, if the device time is 8:30 AM and the server time is 11:30 AM, a report for “Today” displays 11 hours of data (12:00 AM through 11:00 AM) in the server’s time, and 8 hours of data in the device’s time.

- Hourly aggregation may take up to 28 minutes. For example, performance data from 2:00 PM to 3:00 PM may not be shown on reports until 3:28 PM.



NOTE: The CMS report for a device may not match the WXOS report if the device does not respond to a poll, or if the device is rebooted after a poll. When you reboot a device, the data for the current and past hour are purged from the device.

- Statistics for the last hour cannot be viewed while hourly device polling is in progress.
- If you enter an arbitrary date and time range for a report, 24 hours of data are shown for each day, regardless of the specified start and end times.
- To view, save, or print a report as a PDF file, click  PDF in the upper-right corner of the report (the Adobe Acrobat Reader must be installed). Not available for WAN Performance, Compression Overview, and Tunnel Status reports when **All devices** is selected. To email PDF reports, refer to “Scheduling Reports” on page 368.

Configuring the My WAN Page

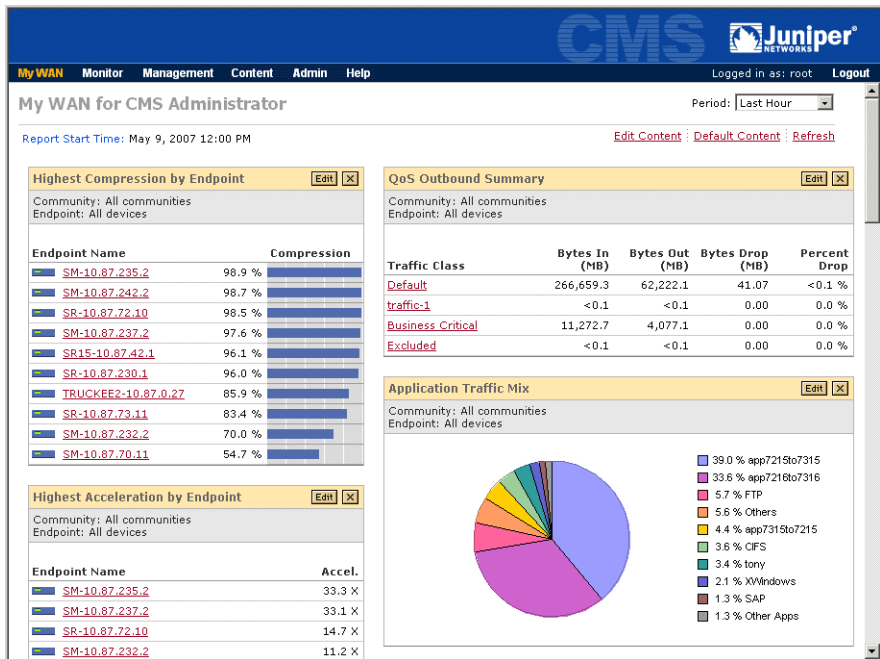
The My WAN page lets each user create a customized mix of charts that depict the overall performance of the devices in one or all communities or device groups. The available charts include:

- The ten applications or endpoints with the highest or lowest compression or acceleration
- The total traffic and dropped traffic for the top ten outbound QoS traffic classes, and the four inbound traffic classes
- The ten monitored applications with the highest percentage of traffic
- The ten monitored endpoints with the worst latency, loss, or availability as measured by WAN performance monitoring
- The ten endpoints, applications, and/or traffic classes with the highest number of WX events, and the ten WX event types with the highest number of occurrences.

To configure the My WAN page:

1. Click My WAN in the menu frame to view the My WAN page for the current user account.

Figure 158: My WAN Page




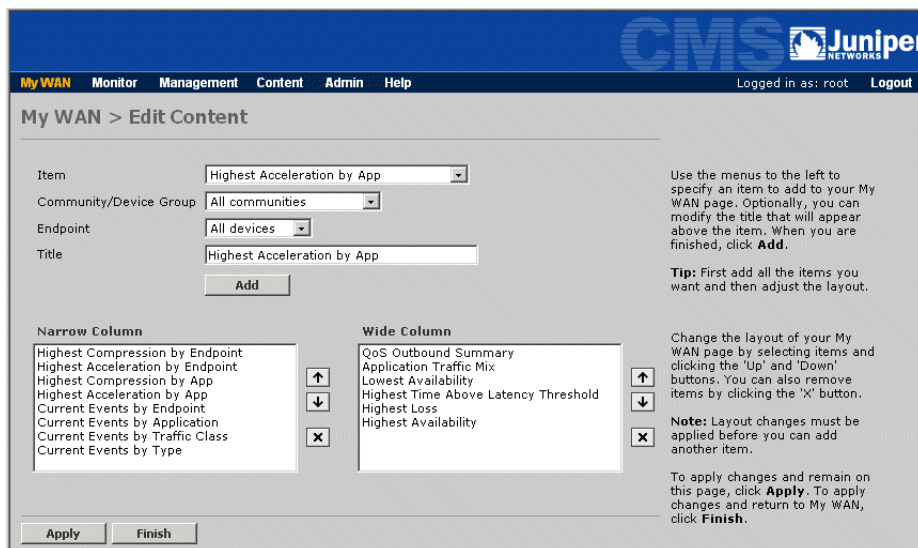
2. To display the default charts, click Default Content.
3. To change the time period for the displayed charts, select a time period from the Period menu in the upper-right corner of the page. You can view the My WAN charts for up to the last month. Note that the event charts show events only for the last 60 minutes, regardless of the selected time period.
4. To change or delete a specific chart on the My WAN page, click the Edit or  buttons in the title bar of the chart.
5. To change the content or layout of the page, click Edit Content.

Figure 159: Edit My WAN Page


6. To add a chart, specify the following information and click Add:
 - a. Select the chart from the Item menu.
 - b. Select a specific community (or device group) and device, as needed, from the Community/Device Group and Endpoint menus (the default is all device groups and communities). You cannot select a specific device for the “by Endpoint” charts.
 - c. Optionally, change the default title of the chart in the Title field.


Table 18 describes the available charts. If you add the same chart multiple times, such as for different communities or device groups, a number is appended to the title automatically. Note that narrow charts are displayed in the left column; wide charts are displayed in the right column.

Table 18: My WAN Charts

Chart	Size	Description
Highest Acceleration by App	Narrow	The ten applications that have the highest acceleration gains from TCP Acceleration, Fast Connection Setup, or CIFS, Exchange, and HTTP acceleration.
Lowest Acceleration by App	Narrow	The ten applications that have the lowest acceleration gains.
Highest Acceleration by Endpoint	Narrow	The ten devices that have the highest acceleration gains from TCP Acceleration or Fast Connection Setup. Click a device name on the chart to view acceleration results by application for the selected device (refer to Figure 183 on page 347 and Figure 184 on page 349).
Lowest Acceleration by Endpoint	Narrow	The ten devices whose applications have the lowest acceleration gains.
Highest Compression by App	Narrow	The ten applications that have the highest percentage of data compression.
Lowest Compression by App	Narrow	The ten applications that have the lowest percentage of data compression.
Highest Compression by Endpoint	Wide	The ten devices that have the highest percentage of data compression. Click a device name on the chart to view the compression details from the selected device to each of the remote devices (refer to Figure 168 on page 325).
Lowest Compression by Endpoint	Wide	The ten devices that have the lowest percentage of data compression.
Application Traffic Mix	Wide	A pie chart of the nine monitored applications that have the highest percentage of the traffic into the selected device(s). The tenth “Other Apps” category indicates the percentage of the traffic for all of the other compressed applications. The “Others” category is for compressed applications that are undefined or unmonitored.
QoS Outbound Summary	Wide	The ten QoS traffic classes with the most outbound traffic. Includes the total number of bytes in and out of the selected devices for each class, and the number and percentage of bytes dropped (if any). Click a class name to view the traffic by device for the selected class, and then click a device to view traffic from the selected device to each QoS endpoint (refer to Figure 176 on page 337).
QoS Inbound Summary	Wide	The total number of bytes into and out of the selected devices for each inbound traffic class, and the number and percentage of bytes dropped (if any). Click a class name to view the traffic by device for the selected class, and then click a device to view traffic into the selected device from all other devices (refer to Figure 179 on page 341).

Table 18: My WAN Charts

Chart	Size	Description
Lowest Availability	Wide	The ten devices that have the lowest percentage of availability as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the low-availability device (refer to Figure 163 on page 317).
Highest Time Above Latency Threshold	Wide	The ten devices where the average latency exceeded the latency threshold for the highest percentage of time, as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the high-latency device (refer to Figure 163 on page 317).
Highest Loss	Wide	The ten devices that have the highest percentage of probe packet loss as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the high-loss device (refer to Figure 163 on page 317).
Highest Availability	Wide	The ten devices that have the highest percentage of availability as measured from another device using WAN Performance Monitoring. Click a "From" device name on the chart to view the WAN Performance report from the selected device to the high-availability device (refer to Figure 163 on page 317).
Current Events by Endpoint	Narrow	The ten devices that have the highest number of unacknowledged events in the last 60 minutes (also indicates the number of critical events). Click a device name on the chart to view the associated list of event occurrences (refer to Figure 191 on page 365 and Figure 192 on page 366).
Current Events by Application	Narrow	The ten applications that have the highest number of unacknowledged performance events in the last 60 minutes (also indicates the number of critical events). Click an application name on the chart to view the associated list of event occurrences (refer to Figure 191 on page 365).
Current Events by Traffic Class	Narrow	<p>The ten traffic classes that have the highest number of unacknowledged performance events in the last 60 minutes (also indicates the number of critical events). Note that only the following performance events apply to traffic classes:</p> <ul style="list-style-type: none"> ■ QoS Throughput Out (Kbps) ■ Bytes Dropped Out (count) ■ Packets Dropped Out (count) <p>Click a traffic class name on the chart to view the associated list of event occurrences (refer to Figure 191 on page 365).</p>
Current Events by Type	Narrow	The ten WX event types (system or performance) that have the highest number of unacknowledged occurrences in the last 60 minutes (also indicates the number of critical events). Click an event type on the chart to view the associated list of event occurrences (refer to Figure 191 on page 365).

7. To delete a chart or change its position on the page:
 - a. To position a chart on the page, select the chart in the Narrow Column or Wide Column lists, and click the up or down arrow keys.
 - b. To delete a chart, select the chart, and click .
 - c. Click Apply to save the changes and stay on the page, or click Finish to return to the My WAN page.

Viewing Reports on the Monitor Page

The following topics describe the reports available on the Monitor page:

- “WAN Statistics” on page 311
- “Data Compression Statistics” on page 321
- “QoS Statistics” on page 336
- “Acceleration Statistics” on page 344
- “Top Traffic Statistics” on page 355
- “Executive Summary” on page 358
- “Trend Reports” on page 360
- “Events Reports” on page 362

WAN Statistics

This section describes the WAN reports.

- “WAN Performance Statistics” on page 312
- “WAN Throughput Statistics” on page 317
- “WAN Application Summary” on page 319

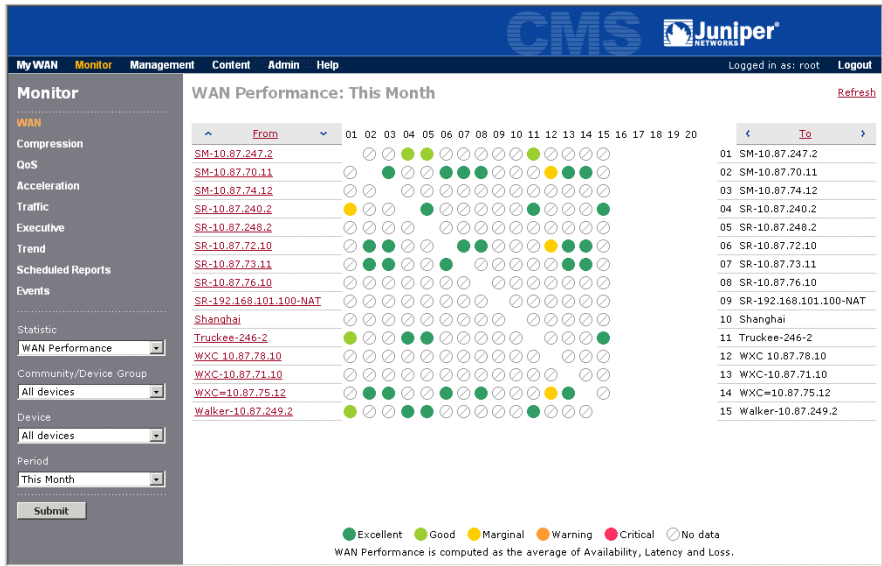
WAN Performance Statistics

The WAN Performance, Loss, Latency, and Availability reports show the WAN performance statistics and events, as measured from each device where WAN Performance Monitoring or Policy-Based Multi-Path is enabled for one or more remote devices.

To view the WAN performance reports:

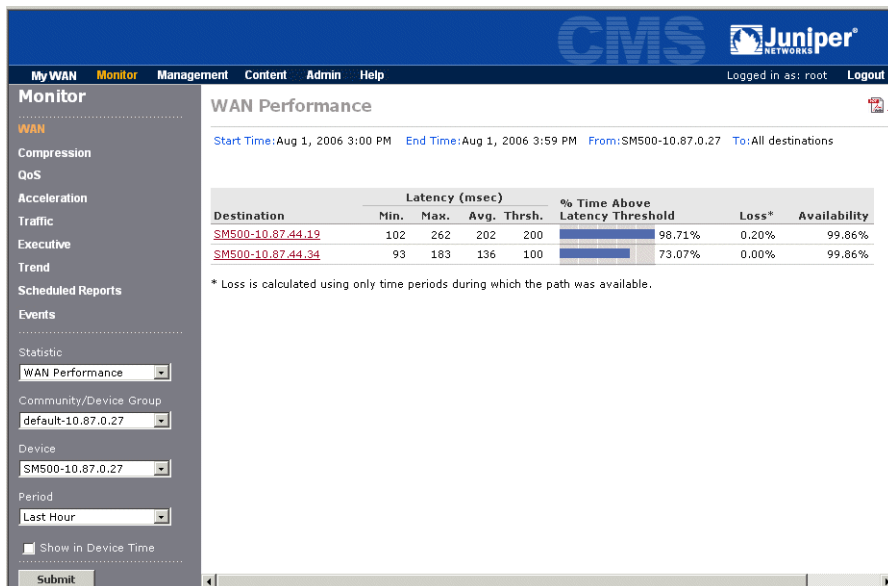
1. Click Monitor in the menu frame, and WAN in the navigation frame.
2. Select one of the following reports from the Statistic menu. Each report provides the same statistics when a specific device is selected. When All devices is selected, each report shows a different statistic as a matrix of color-coded cells to indicate the status between each pair of devices:
 - **WAN Performance.** Best, average, or worst values measured for loss, latency, and availability (the default is average). To change the performance ranges represented by each color, refer to “Setting WAN Performance Thresholds” on page 418.
 - **WAN Latency.** Percentage of the selected time period that the average latency exceeded the specified threshold.
 - **WAN Loss.** Percentage of probe packets that were lost.
 - **WAN Availability.** Percentage of the minutes in the selected time period for which at least one probe was acknowledged.
3. Change one or more of the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a specific device from the Device menu to view a table of performance statistics measured from the selected device to each remote device. The default is **All devices**, which shows a color-coded matrix for all monitored devices.
 - Select a time period from the Period menu. You can select the previous hour, day, week, month, or six months. The default is Last Hour.
4. If WAN Performance and All devices are selected, click Submit to open the WAN Performance page for the selected community or device group.

Figure 160: WAN Performance for All Devices in a Community/Device Group



From the report you can:

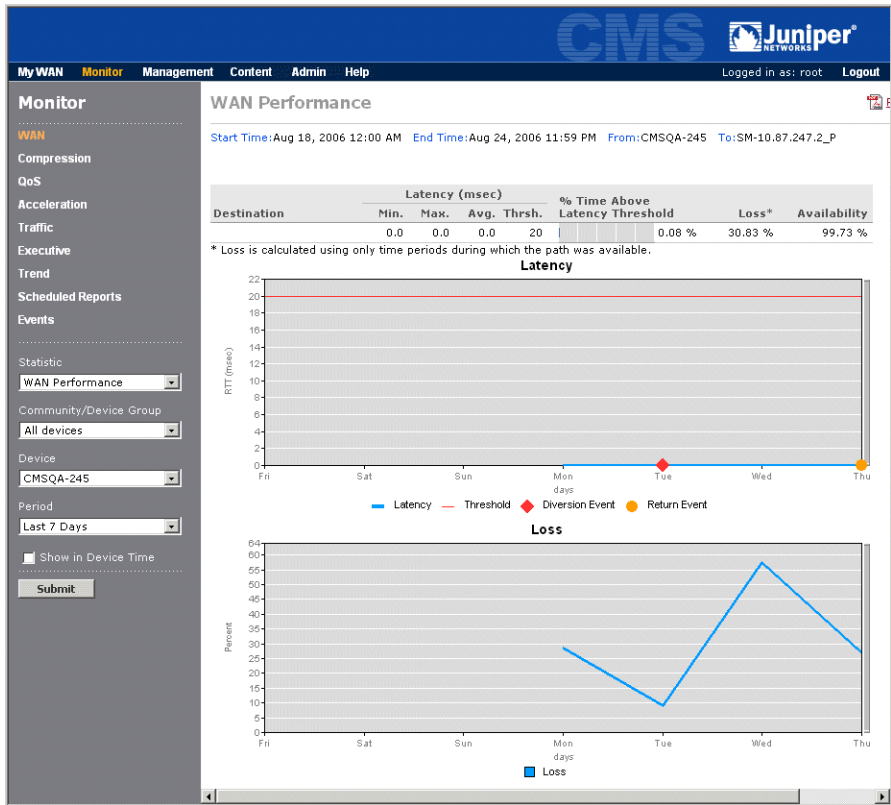
- View the WAN performance from each device in the From column to each device in the To column. The same devices are listed in both columns so you can see the performance in both directions. A cell is white if both WAN Performance Monitoring and Policy-Based Multi-Path are disabled between the devices.
 - Move the cursor over a cell to highlight the two devices and display the exact percentages in the browser's status bar for loss, time above the latency threshold, and availability (measured by the From device).
 - View the next group of devices by moving the cursor over the From or To column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.
5. To view the WAN performance statistics between a specific device and each of the other devices it is monitoring, click the device name in the From column, or select the device from the Device menu and click Submit.

Figure 161: WAN Performance Statistics for All Destinations

The following information is shown for each monitored device.

- **Device Name.** Name of the remote device. Devices that support Multi-Path have a “_Pri” or “_Sec” appended to the device name to indicate the primary or secondary path.
 - **Latency (msec).** Probes are used to measure the lowest, highest, and average round-trip times between the selected device and the remote device (in milliseconds). The latency threshold is also displayed.
 - **% Time Above Latency Threshold.** Percentage of the selected time period that the average latency exceeded the specified threshold.
 - **Loss.** Percentage of the probes that were lost.
 - **Availability.** Percentage of the minutes in the selected time period for which at least one probe was acknowledged.
6. To view the WAN performance graphs and events for a specific device, click the device name or select a colored cell on the matrix shown for All devices. The information on the performance graphs depends on whether the device is enabled for Multi-Path (Figure 162) or WAN performance monitoring (Figure 163).




Figure 162: Multi-Path WAN Performance Charts



For a Multi-Path device, the following information is shown on the Loss and Latency charts (Figure 162):

- The Latency chart shows the average round-trip time for the selected path, and indicates the configured latency threshold for the path.

- The following icons are used to indicate performance events. Move the cursor over the icon to view the number of events in the time period.

Icon	Description
	<p>Indicates that traffic was switched to the alternate path due to one of the following conditions:</p> <ul style="list-style-type: none">■ Loss or latency threshold exceeded. Eligible traffic is diverted only if the alternate path's 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).
	<p>Indicates that performance has returned to normal, and traffic was switched back to the preferred path (the tunnel must be up).</p>
	<p>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.</p>

- The Loss chart shows the percentage of the probes that were lost on the selected path. If the loss threshold is exceeded, a diversion to the alternate path is indicated on the Latency chart (as shown in Figure 162), provided the alternate path is not degraded.

Figure 163: Single-Path WAN Performance Charts



For WAN performance monitoring endpoints (Figure 163), the loss and latency are shown for the selected path, and the **E** icon indicates the loss or latency threshold was exceeded.



NOTE: If the remote device is unreachable, all paths will be down, the Latency chart will be blank (latency cannot be measured), and the Loss chart will show 100 % probe loss on all paths.

WAN Throughput Statistics

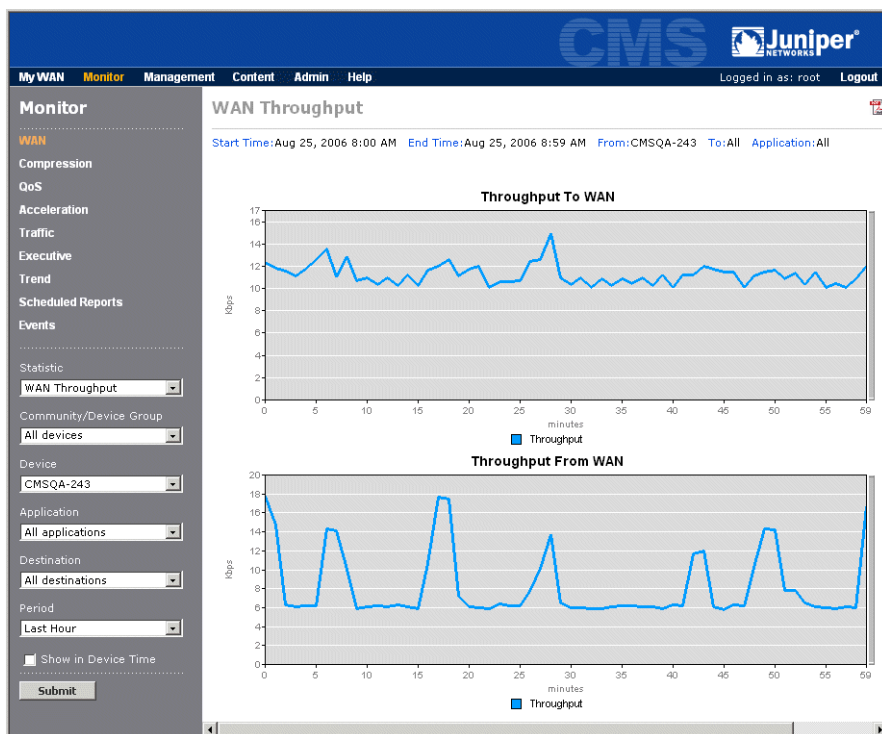
The WAN Throughput report shows the speed of the traffic to and from the WAN for a specific device. You can view a device's WAN throughput to all remote destinations, a specific remote WX device, or all non-WX destinations.

To view WAN throughput:

1. Click Monitor in the menu frame, and WAN in the navigation frame.
2. Select WAN Throughput from the Statistic menu.
3. Select a community or device group from the Community/Device Group menu., and select a device from the Device menu.

4. Change one or more of the following report parameters, and click Submit.
 - Select a specific monitored application from the Application menu. Select Others to view statistics for applications that are undefined or unmonitored. The default is All applications. To specify the monitored applications on a device, refer to “Monitoring Applications” on page 151.
 - Select a specific WX device from the Destination menu. Select Other Destinations to view statistics for all non-WX destinations.
 - Select a time period from the Period menu. You can select the previous hour, day, week, month, or six months. The default is Last Hour.

Figure 164: WAN Throughput Report



5. Review the following information on the two throughput graphs. Keep in mind that all values are for the selected application, destination, and time period.

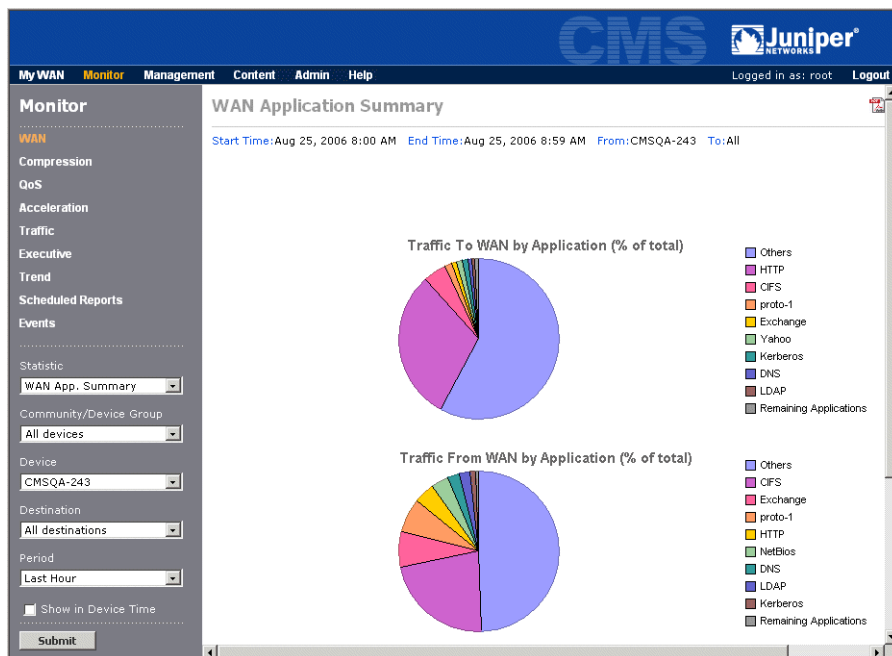
- The Throughput to WAN graph shows the average throughput of data sent to the WAN.
- The Throughput From WAN graph shows the average throughput of data received from the WAN. This graph is blank when the device is in Profile Mode.

WAN Application Summary

The WAN Application Summary shows the application traffic to and from the WAN for a specific device. You can view a device's WAN traffic to all remote destinations, a specific remote WX device, or all non-WX destinations. The traffic is shown for up to 40 monitored applications.

To view the WAN Application Summary:

1. Click Monitor in the menu frame, and WAN in the navigation frame.
2. Select WAN App. Summary from the Statistic menu.
3. Select a community or device group from the Community/Device Group menu., and select a device from the Device menu.
4. Change one or more of the following report parameters, and click Submit.
 - Select a specific WX device from the Destination menu. Select Other Destinations to view statistics for all non-WX destinations.
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 165: WAN Application Summary

5. Review the information on the following charts. Keep in mind that all values are for the selected destination and time period.
 - The two pie charts show the nine monitored applications that have the highest percentage of the total traffic sent to and from the WAN for the selected device. The **Remaining applications** category shows the traffic percentage for all other applications.
 - The application table shows the traffic in megabytes sent to and from the WAN for each monitored application. The applications are sorted in descending order by total traffic. The **Others** category indicates the traffic for applications that are undefined or unmonitored.

Compression Statistics

This section describes the compression reports:

- “Data Compression Statistics” in the next section
- “Application Summary Statistics” on page 328
- “Passthrough Statistics” on page 330
- “Packet Size Distribution Statistics” on page 332
- “Monitoring Tunnel Status” on page 333

Data Compression Statistics

The Compression reports let you view the percentage of compression for:

- Each pair of WX devices in the same community (matrix view).
- A selected device and each of the other devices in a community/device group.
- Each application for a selected pair of devices.
- A selected application from a specific device to each of the other devices in the same community.

The percentage of compression for the selected time period is based on the total number of bytes in and out of each device, as follows

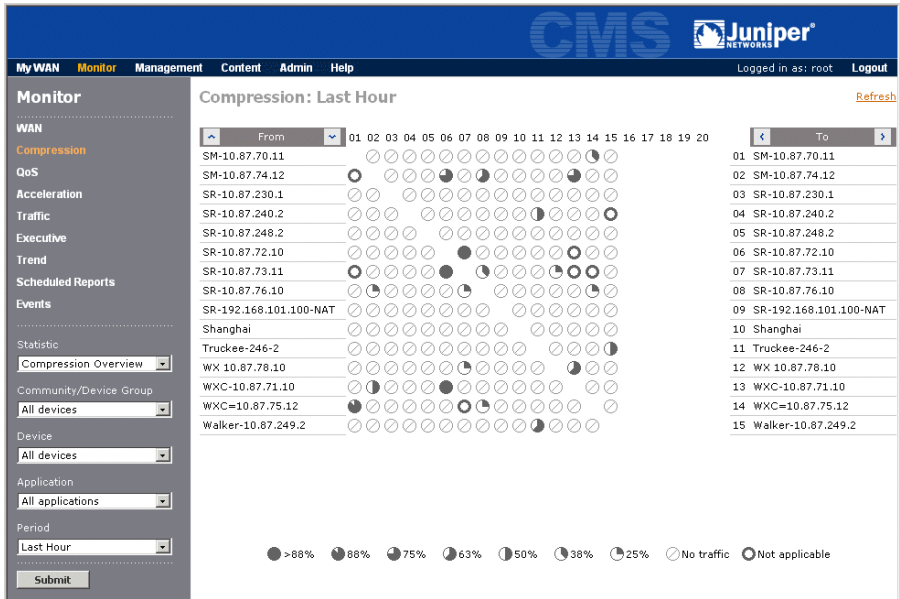
$$\% \text{ of Reduction} = \left(\frac{\text{Bytes In} - \text{Bytes Out}}{\text{Bytes In}} \right) \times 100$$

To view the Compression reports:


1. Click Monitor in the menu frame, and Compression in the navigation frame.
2. Select one of the following reports from the Statistic menu:
 - **Compression Overview.** Percentage of compression shown in a matrix between each pair of devices in the selected community or device group (when **All devices** is selected). Select a cell in the matrix to view graphs of compression by time. Select a device from the Device menu to view compression from the selected device to each remote device for one or all applications (in one or both directions).
 - **Compression by Time.** Percentage of compression by time from a selected device to one or all remote destinations, for one or all applications.
 - **Compression by Dest.** Same as the Compression Overview, except that a specific device must be selected (no matrix view).
 - **Compression by App.** Percentage of compression for all applications from a selected device to a remote destination.
3. Change one or more of the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a specific device from the Device menu to view compression statistics measured from the selected device to each remote device.
 - Select a specific monitored application from the Application menu. Select Others to view statistics for applications that are undefined or unmonitored. The default is All applications. To specify the monitored applications on a device, refer to “Monitoring Applications” on page 151.
 - Select a specific device from the Destination menu to view compression statistics measured to a specific device (on the by-time and by-application reports only).

- Select a time period from the Period menu. You can select the previous hour, day, week, month, or six months. The default is Last Hour.
4. If Compression Overview and All devices are selected, click Submit. to open the Compression page for all devices in the selected community/device group.

Figure 166: Percentage of Data Compression for a Community/Device Group

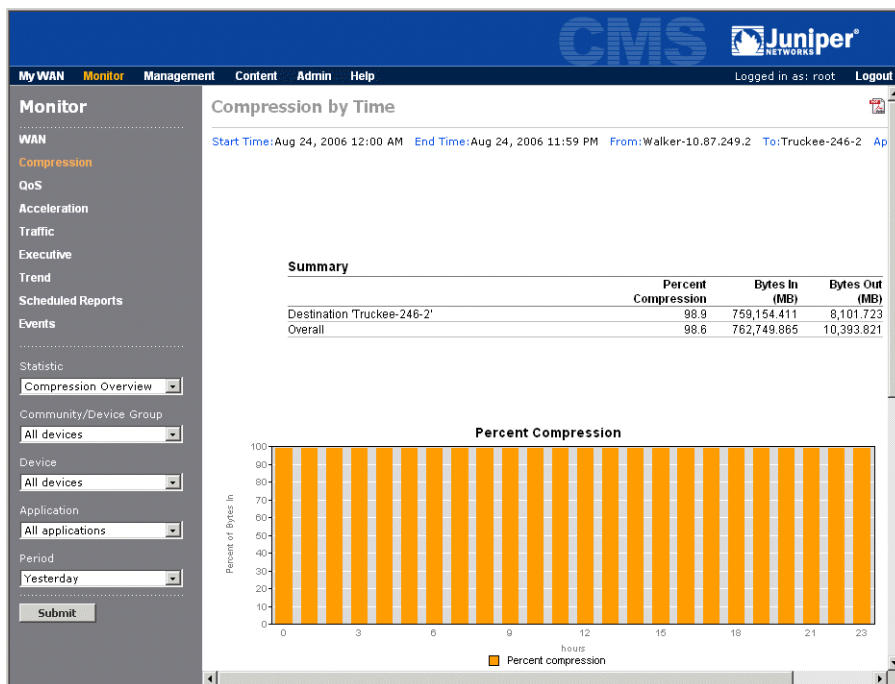


From the report page, you can:

- View the percentage of compression for traffic sent from each device in the From column to each device in the To column. The same devices are listed in both columns so you can see the compression in both directions. The icons indicate a percentage range.
- The  icon indicates that the device is down or unreachable, or there is no compression.
- Move the cursor over an icon to highlight the two devices and display the exact compression percentage in the browser's status bar, along with the number of bytes and packets in and out of the From device. Note that the compression percentage indicated by the icon is approximate.

- View the next group of devices by moving the cursor over the From or To column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.
5. Click the icon for a pair of devices to view graphs of compression by time for the traffic sent from a device in the From column to a device in the To column.

Figure 167: Percentage of Data Compression By Time

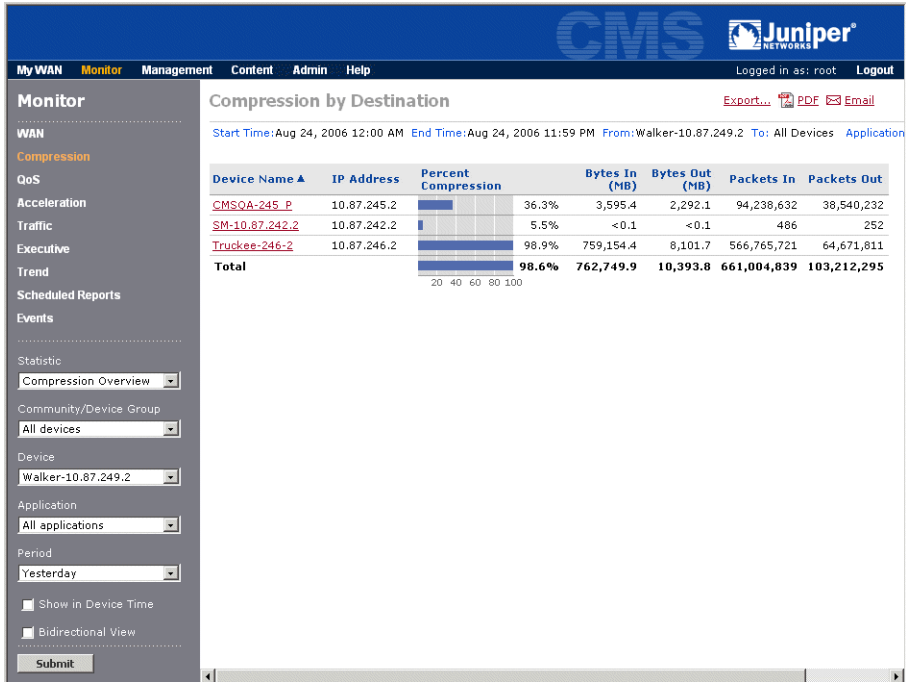


The report shows the percentage compression for the remote device, the overall compression for all remote devices, the number of bytes in and out of the compression engine, and graphs of the compression, bytes, and packets for the selected time period.

To view the same graphs from a selected device to all remote devices, select Compression by Time from the Statistic menu and select the device from the Device menu.

- 6. To view the percentage of compression between a specific device and each of the other devices in the same community, select the device name from the Device menu and click Submit. Alternatively, select Compression by Dest. from the Statistic menu and select the device from the Device menu.

Figure 168: Percentage of Data Compression for a Selected Device



From the report page, you can:

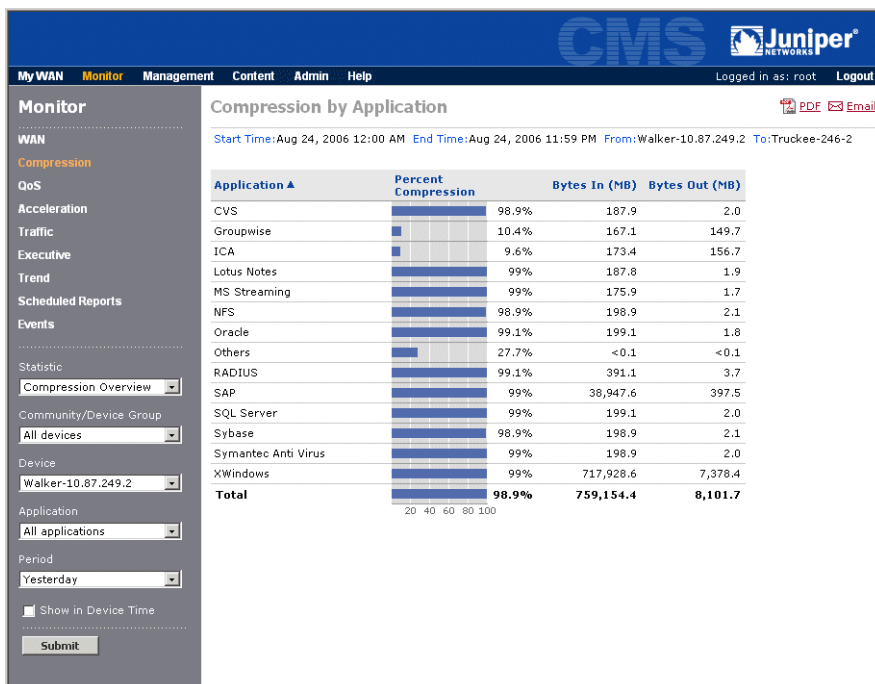
- View the outbound percentage of compression achieved by the selected device for each of the other devices in the same community. The number of bytes and packets in and out of the compression engine on the selected device is shown for each destination device.
- Click Export to view or save the displayed data in CSV format.



NOTE: If the selected device resides in multiple communities, the report includes compression statistics for devices in each community.

- Click a device name to view the percentage of compression for each compressed application in the traffic sent to the device. Alternatively, select Compression by App. from the Statistic menu and select the device from the Device menu.

Figure 169: Percentage of Data Compression By Application

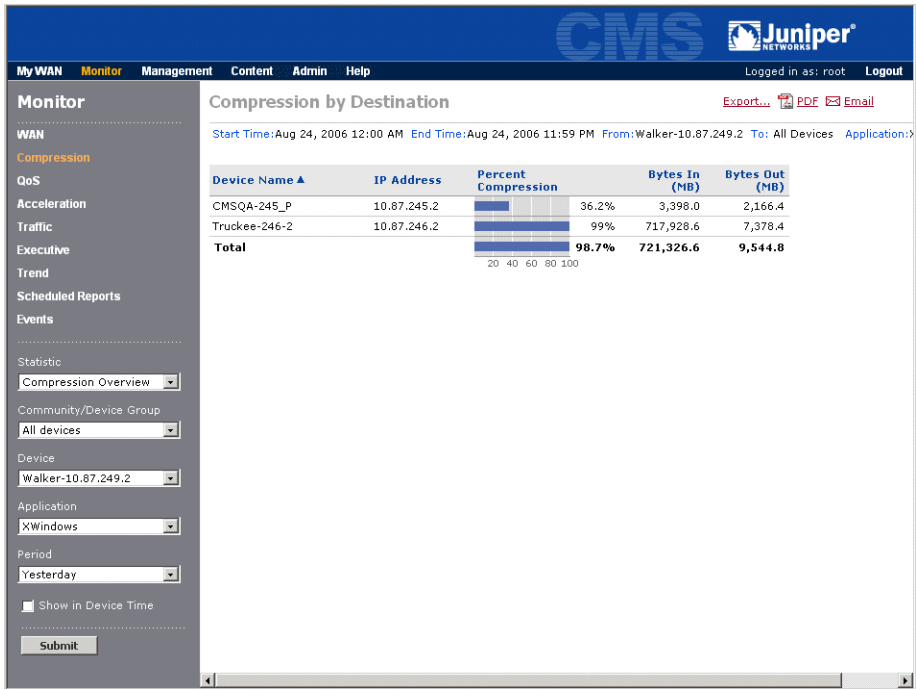


The report displays the percentage compression and number of bytes in and out of the device for each compressed application in the selected time period.

Note that the Others application is for applications that are undefined or unmonitored on the From device.

- To view the percentage of compression for a specific application from the selected device to each of the other devices in the same community, select a monitored application from the Application menu, and click Submit.

Figure 170: Percentage of Data Compression for a Selected Application



From the report page, you can:

- View the percentage of compression achieved for each destination device by the selected device and application. For the selected application, the number of bytes and packets in and out of the compression engine on the selected device is shown for each destination device.
- Click Export to view or save the displayed data in CSV format.

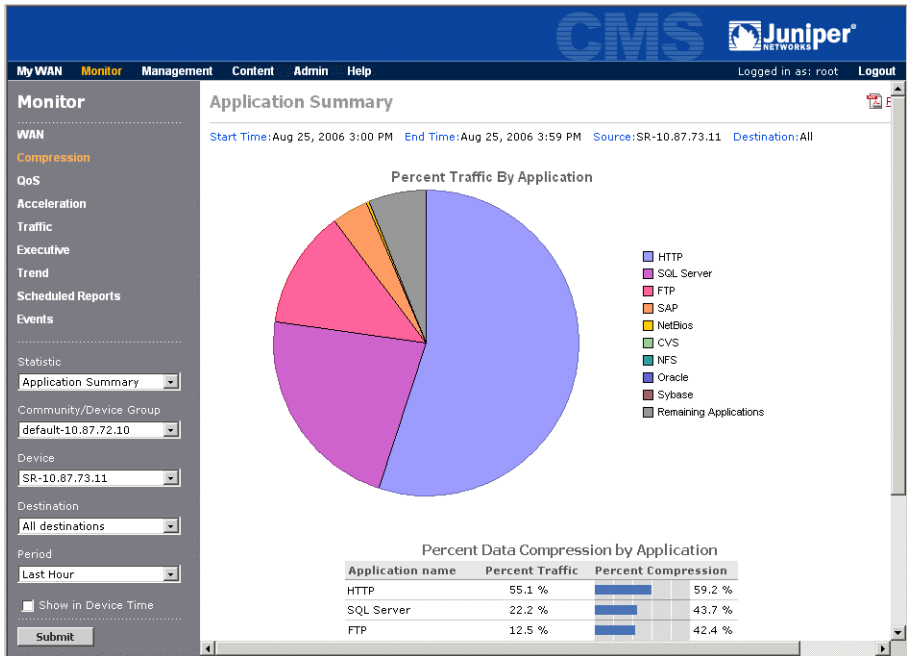
Application Summary Statistics

The Application Summary shows a pie chart of the nine monitored applications that have the highest percentage of the traffic into a selected device for one or all remote devices. A table is also included that shows the traffic statistics and percentage of compression for each monitored application (up to 40).

To view the Application Summary:

1. Click Monitor in the menu frame, and Compression in the navigation frame.
2. Select Application Summary from the Statistic menu.
3. Select a community or device group from the Community/Device Group menu., and select a device from the Device menu.
4. Change one or more of the following report parameters, and click Submit.
 - Select a specific device from the Destination menu.
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 171: Application Summary Statistics



5. Review the following information on the Application Summary.
- The pie chart shows the nine monitored applications with the highest percentage of the total traffic into the device for the selected destination. The **Remaining applications** category shows the traffic for all other applications (both defined and undefined).
 - The application table has the following columns.
 - **Application Name.** Names of the monitored applications, sorted in descending order by compression percentage. The Others category indicates the traffic for compressed applications that are undefined or unmonitored.
 - **Percent Traffic.** Percentage of the total traffic into the device's compression engine for each application.
 - **Percent Compression.** Percentage of compression achieved for each application. A dash is shown for applications that have no traffic or cannot be compressed (such as encrypted applications). Compression should be disabled for applications that consistently show little or no compression (refer to "Compressing Applications" on page 156).

Passthrough Statistics

Traffic that falls into one of several categories is passed through the devices with no attempt at 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.

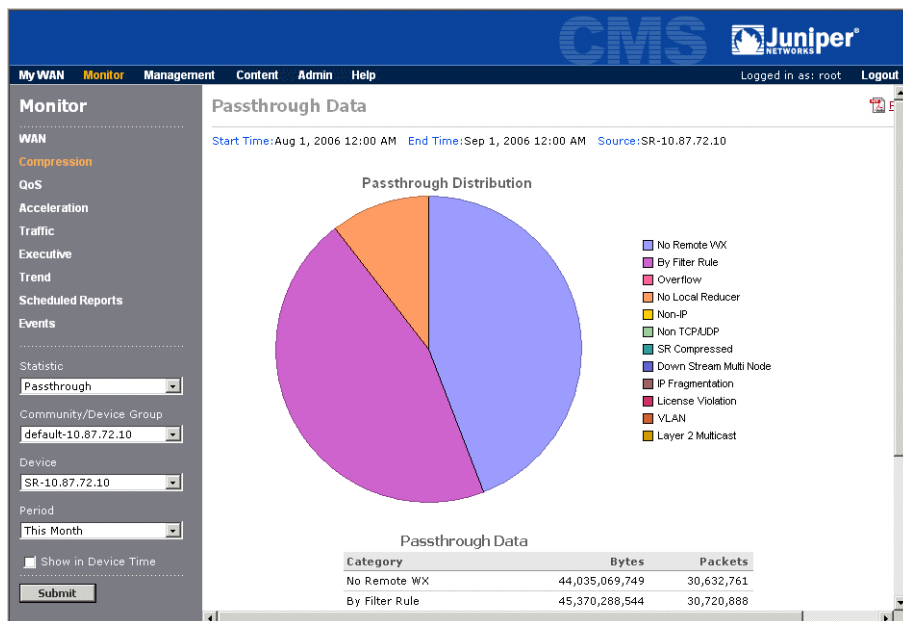


NOTE: An off-path device that uses RIP for packet interception has no passthrough statistics. All traffic is sent through the tunnel.

To view passthrough statistics:

1. Click Monitor in the menu frame, and Compression in the navigation frame.
2. Select Passthrough from the Statistic menu.
3. Select a community or device group from the Community/Device Group menu., and select a device from the Device menu.
4. Select a time period from the Period menu, and click Submit. You can select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 172: 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.
Overflow	Traffic volume exceeded the device capacity.
No Local Compressor	Compression is disabled on this device.
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 device.
IP Fragmentation	Always zero unless compression of IP fragments is disabled (refer to the “configure filter” CLI command in the operator’s guide).
License Violation	The licensed throughput speed was exceeded.
VLAN	Total VLAN traffic that was not 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.

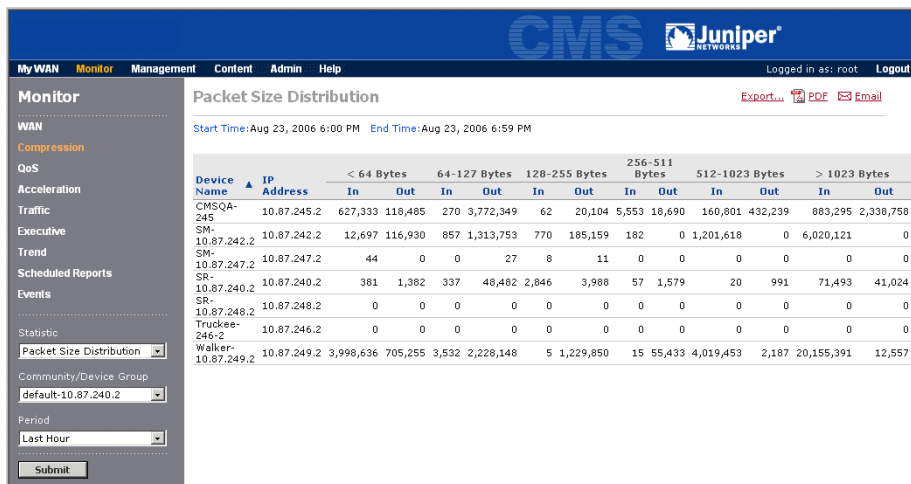
Packet Size Distribution Statistics

For each device in a selected community or device group, the Packet Size Distribution report shows the number of packets in and out of the compression engine for each of six packet-size ranges.

To view packet size distribution statistics:

1. Click Monitor in the menu frame, and Compression in the navigation frame.
2. Select Packet Size Distribution from the Statistic menu.
3. Select a community or device group from the Community/Device Group menu.
4. Select a time period from the Period menu and click Submit. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 173: Packet Size Distribution Statistics



Monitoring Tunnel Status

By default, each WX device attempts to form a pair of tunnels with each of the other devices in the same community. An outbound tunnel carries compressed data to another device; an inbound tunnel carries data compressed by another device. You can configure each device to specify which tunnels are formed.

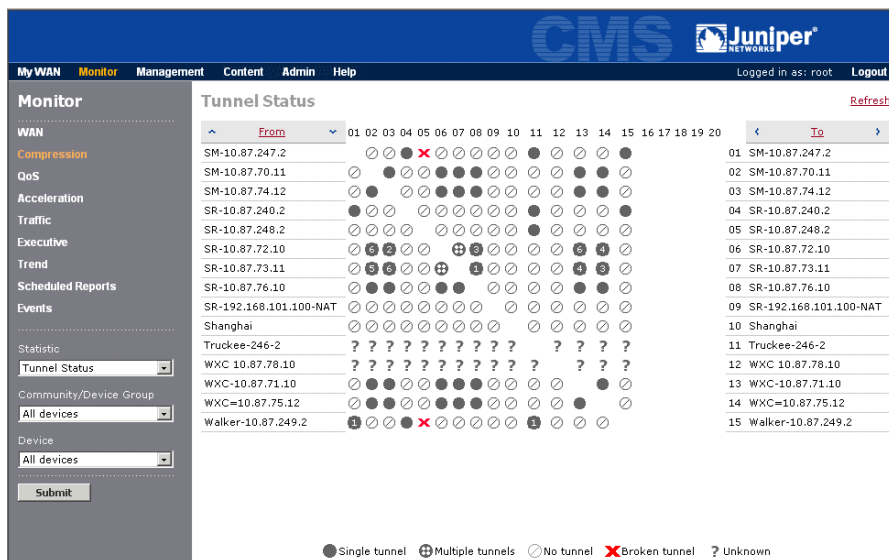
The Tunnel Status reports let you view the tunnel status for:

- The outbound tunnel for each pair of devices in the same community (matrix view).
- A selected device's outbound and inbound tunnels to and from each of the other devices in the same community (table view).

If devices in the same community are in different time zones, the CMS server time is shown in the Last Update field at the top of the Tunnel Status pages.

To view the Tunnel Status reports:

1. Click Monitor in the menu frame, and Compression in the navigation frame.
2. Select Tunnel Status from the Statistic menu.
3. Select a community or device group from the Community/Device Group menu.
4. To view a matrix showing the outbound tunnel status between each pair of devices in the same community, select All devices from the Device menu, and click Submit.

Figure 174: Monitoring Tunnel Status for a Community/Device Group

From the report page, you can:

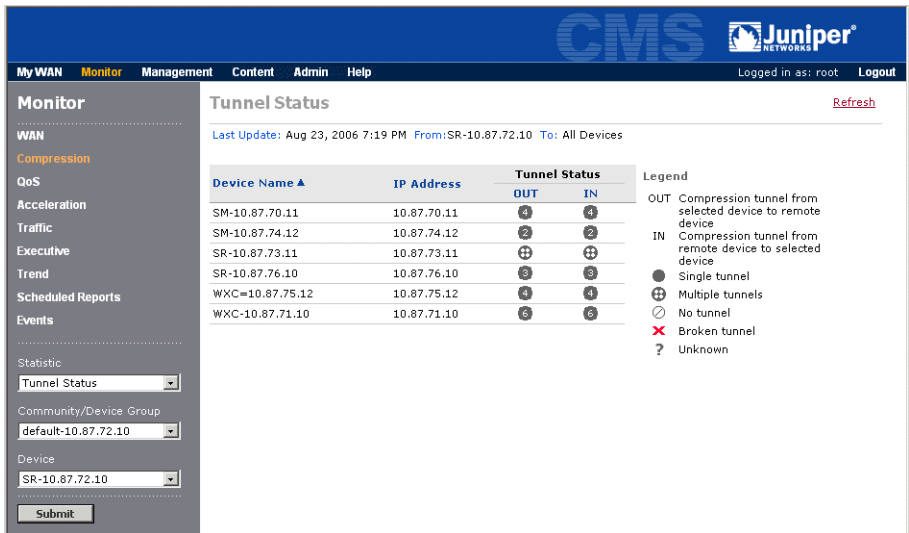
- View the outbound tunnel status from a device in the From column to a device in the To column. Move the cursor over an icon to highlight both devices. The status icons are described below.

Icon	Description
	Tunnel established — An outbound tunnel exists from a device in the From column to a device in the To column. For a WX 100 server running WXOS 5.4 or later, and one or more client devices, a number in the circle (1 to 6) indicates the number of the client hosting the tunnel.
	Multiple tunnels — Multiple tunnels exist (up to 6) between two WX 100 servers with the same number of client devices.
	No tunnel — No outbound tunnel exists due to a policy setting. For example, this icon is displayed if you manually disable compression from one device to another.
	Broken tunnel — No outbound tunnel exists due to an error or problem, such as low system resources. For more information, open the WXOS Web console for the “from” device and click Compression to view the Endpoints page. NOTE: If a device in a user-defined community is removed from the network, it will be shown with broken tunnels until it is deleted from the registration server.
	Temporarily unavailable — The tunnel is in a transitory state, or the device is down or unreachable.

- View the next group of devices by moving the cursor over the From or To column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.
 - To update the tunnel status from the devices, click Refresh.
5. To view a device’s outbound and inbound tunnels to and from each of the other devices in the same community, select the device name from the Device menu and click Submit.

The Tunnel Status page for the selected device opens (Figure 175).

Figure 175: Monitoring Tunnel Status for a Device



The tunnel status information shown here is the same as the status shown on the Endpoints page in the device’s WXOS Web console. Note the following:

- The OUT column indicates the status of the outbound tunnel from the selected device to each device in the table; the IN column indicates the status of the inbound tunnel on the selected device from each of the listed devices.
- An **✗** icon in the IN column indicates that the inbound tunnel has a problem or that compression to the selected device is disabled on the Endpoints page of the device listed in the table.



NOTE: If the selected device resides in multiple communities, the report includes the tunnel status for devices in each community.

QoS Statistics

This section describes the QoS reports.

- “Outbound QoS Statistics” in the next section
- “Inbound QoS Statistics” on page 340

Outbound QoS Statistics

If outbound QoS is enabled on a device, the QoS Outbound reports display the following statistics for the traffic into the Local (LAN) interface and out of the Remote (WAN) interface:

- Total number of bytes and packets in and out of a selected device for each destination. Includes the number of bytes and packets dropped.
- Byte and packet counts for each traffic class on a selected device for a specific destination. Includes the throughput for each class.
- Throughput in and out of a selected device for a specific traffic class and destination. Includes the rate of dropped packets.

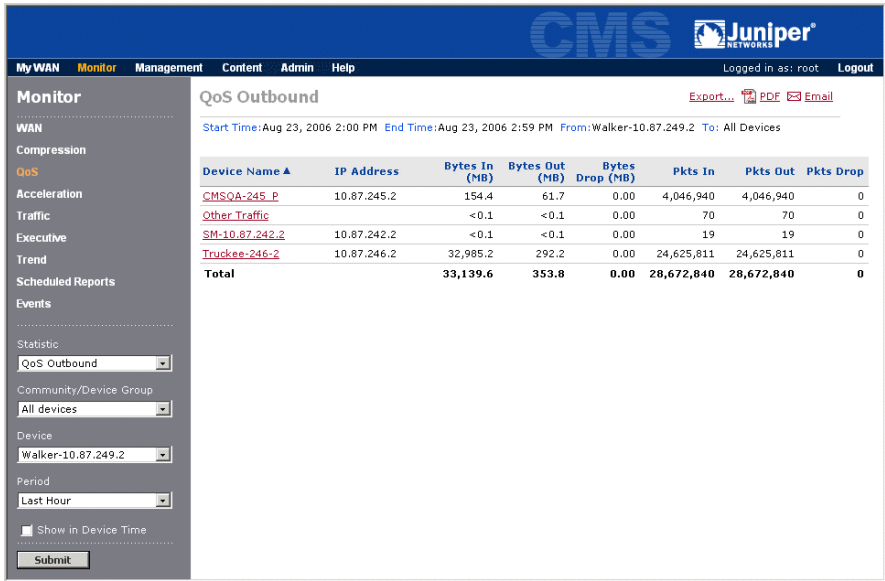


NOTE: Outbound QoS is not effective for an off-path device unless all outbound WAN traffic is routed through the device.

To view the QoS Outbound reports:

1. Click Monitor in the menu frame, and QoS in the navigation frame.
2. Select QoS Outbound from the Statistic menu.
3. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a device from the Device menu that has outbound QoS enabled. Devices using outbound QoS have a **QoS** on the Devices page (click Management in the menu frame to view the Devices page).
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 176: QoS Outbound Report for a Selected Device



From the QoS Outbound report page, you can:

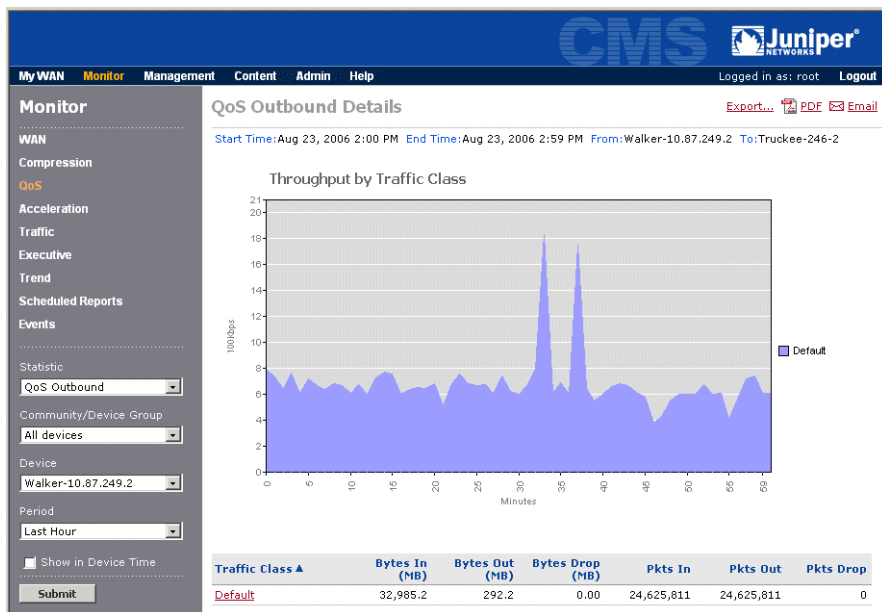
- View the total number of bytes and packets (both compressed and uncompressed) in and out of the selected device for each of the destination devices that are defined as QoS endpoints. The number of bytes and packets dropped by the device is also shown.

The Other traffic “device” does not have an IP address because it indicates all traffic that is not sent to a device that is designated as a QoS endpoint.



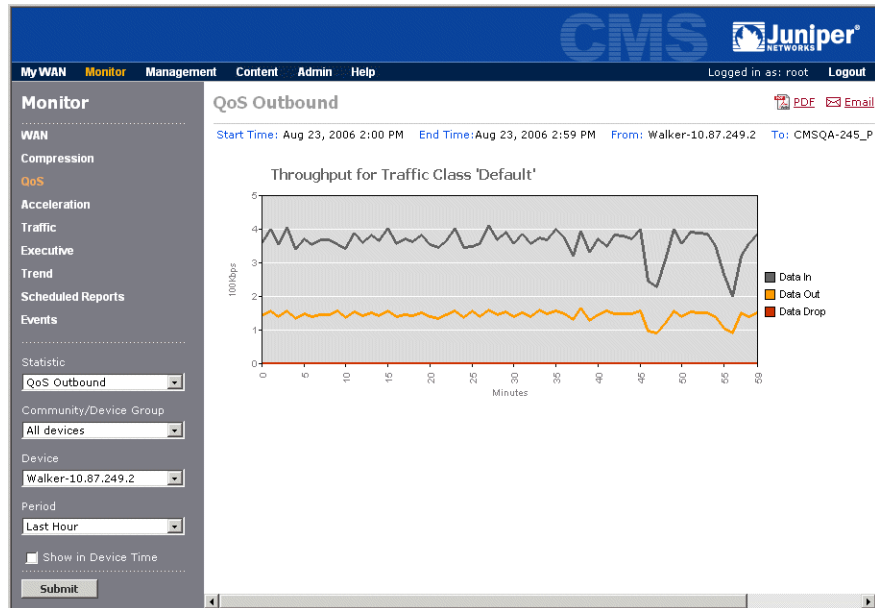
NOTE: If the selected device resides in multiple communities, the report includes the destination devices in each community.

- Click Export to view or save the displayed data in CSV format.
- Click a device name (destination) to view the throughput and byte and packet counts for each traffic class defined on the selected device (Figure 177) for the selected destination.

Figure 177: QoS Outbound Details by Traffic Class

From the QoS Outbound Details page, you can:

- View a graph of the throughput for each traffic class, and a table of the byte and packet counts for each traffic class, including the number of bytes and packets dropped by the device for this destination.
- Click Export to view or save the tabular data in CSV format.
- Click a traffic class name to view the throughput in and out of the device, and the rate of dropped traffic for the class (Figure 178).

Figure 178: QoS Outbound Throughput for a Selected Traffic Class

The Throughput graph shows the following:

- **Data In** (grey line). Average data throughput into the Local interface from the LAN side of the device.
- **Data Out** (orange line). Average throughput to the WAN side of the device. Indicates the compression achieved for the selected destination.
- **Data Dropped** (red line). Average rate that outbound data was dropped. Data is dropped when the traffic for the selected class exceeds the maximum allocated bandwidth or when the guaranteed bandwidth is exceeded while the circuit is fully utilized.

Note that brief bursts of traffic can cause data to be dropped, even when the average throughput is well below the maximum bandwidth.

Inbound QoS Statistics

If inbound QoS is enabled on a device, the QoS Inbound reports display the following statistics for the traffic into the Remote (WAN) interface and out of the Local (LAN) interface:

- Total number of bytes and packets in and out of a selected device. Includes the number of bytes and packets dropped.
- Byte and packet counts for the inbound traffic classes on a selected device. Includes the throughput for each class.
- Throughput in and out of a selected device for a specific traffic class. Includes the rate of dropped packets.

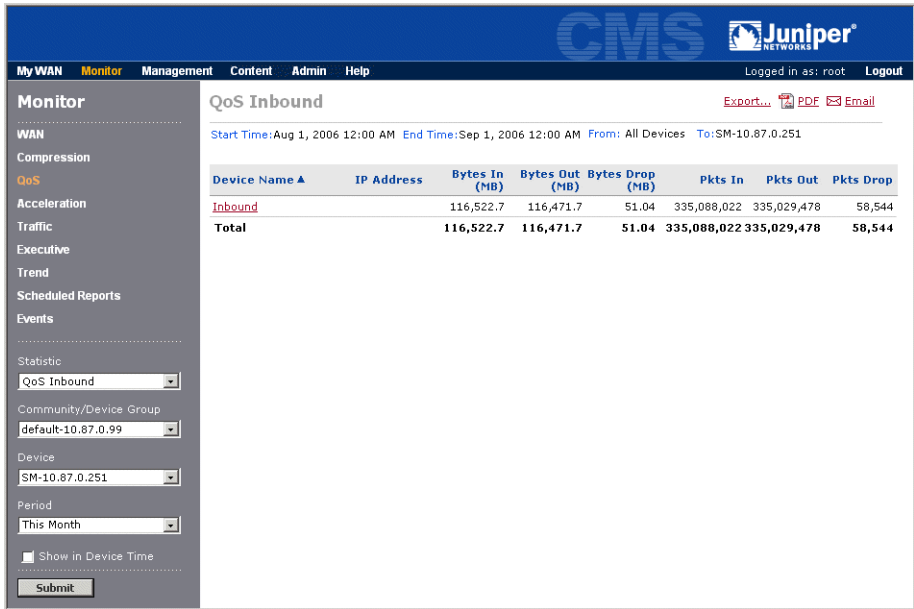


NOTE: QoS Inbound reports do not apply to off-path devices.

To view the QoS Inbound reports:

1. Click Monitor in the menu frame, and QoS in the navigation frame.
2. Select QoS Inbound from the Statistic menu.
3. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a device from the Device menu that has inbound QoS enabled.
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 179: QoS Inbound Report for a Selected Device



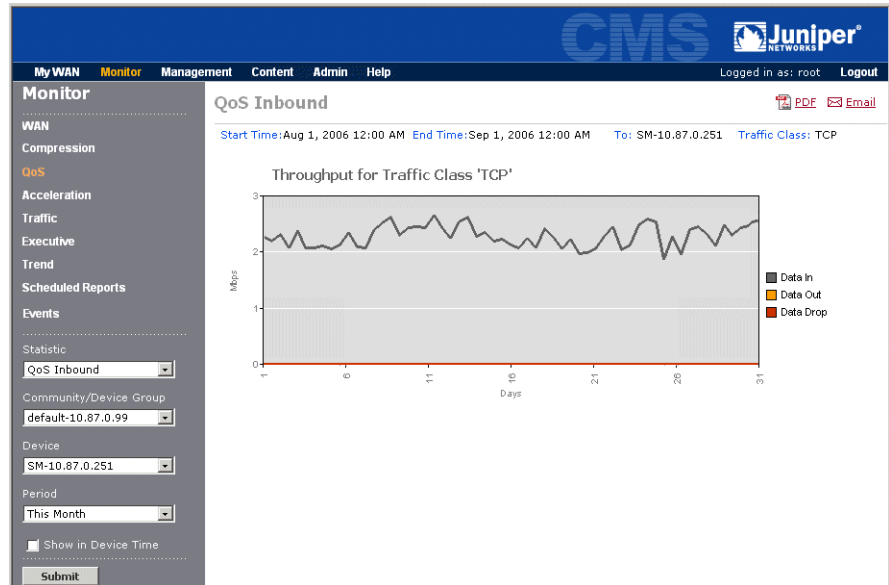
From the QoS Inbound report page, you can:

- View the total number of bytes and packets in and out of the selected device. The number of bytes and packets dropped by the device is also shown.
- Click Export to view or save the displayed data in CSV format.
- Click Inbound to view the throughput and byte and packet counts for each of the inbound traffic classes (Figure 180).

Figure 180: QoS Inbound Details by Traffic Class

From the QoS Inbound Details page, you can:

- View a graph of the throughput for each traffic class, and a table of the byte and packet counts for each traffic class, including the number of bytes and packets dropped by the device.
- Click Export to view or save the tabular data in CSV format.
- Click a traffic class name to view the throughput in and out of the device, and the rate of dropped traffic for the class (Figure 178).

Figure 181: QoS Inbound Throughput for a Selected Traffic Class

The Throughput graph shows the following:

- **Data In** (grey line). Average data throughput into the Remote interface from the WAN side of the device.
- **Data Out** (orange line). Average throughput to the LAN side of the device.
- **Data Dropped** (red line). Average rate that inbound data was dropped. Data is dropped when the traffic for the selected class exceeds the maximum allocated bandwidth.

Acceleration Statistics

This section describes the acceleration reports.

- “Acceleration Summary” in the next section
- “TCP Acceleration Statistics” on page 346
- “Fast Connection Setup Statistics” on page 348
- “CIFS and Exchange Acceleration Statistics” on page 351
- “HTTP Acceleration Statistics” on page 353

Acceleration Summary

The Acceleration Summary report consolidates compression and acceleration statistics for a pair of WX devices. Acceleration statistics are included for TCP Acceleration, CIFS, Exchange, and HTTP. Normally, compression and TCP Acceleration statistics must be measured on the source device, while CIFS, Exchange, and HTTP acceleration are measured on the destination device.

To view the Acceleration Summary:


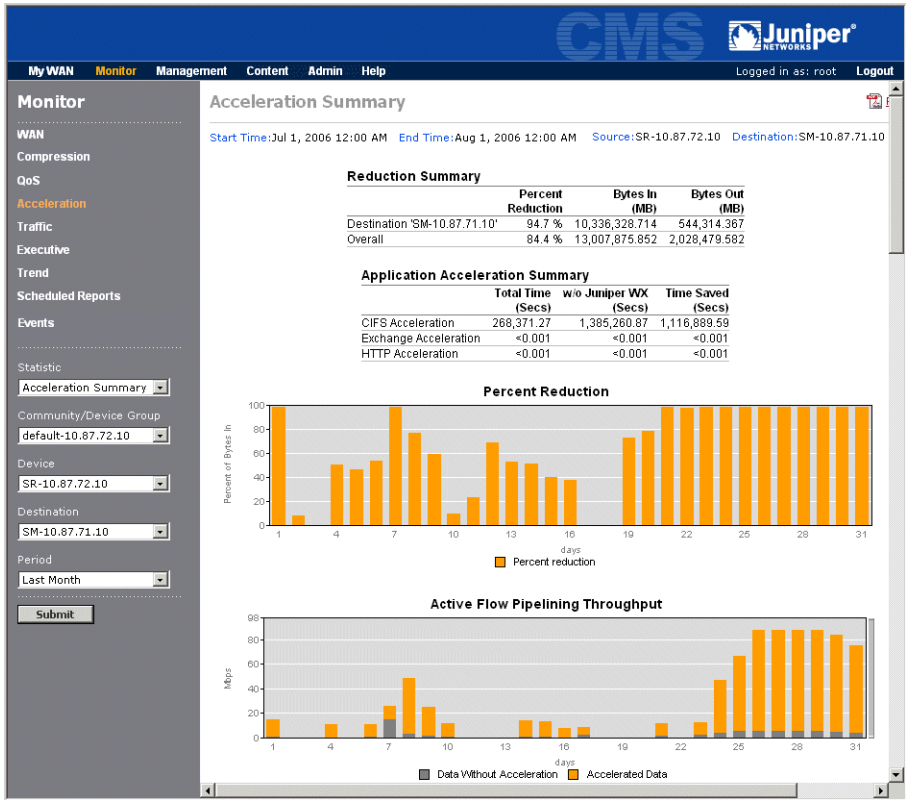
1. Click Monitor in the menu frame, and Acceleration in the navigation frame.
2. Select Acceleration Summary from the Statistic menu.
3. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a source device from the Device menu. Devices that have acceleration enabled have a  on the Devices page (click Management in the menu frame to view the Devices page).
 - Select a device from the Destination menu to view compression and acceleration statistics for traffic sent to the selected device.
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 182: Acceleration Summary



Review the following information.


- The Compression Summary shows the percentage compression for traffic sent to the destination device, the overall compression percentage for all remote devices, and the number of bytes in and out of the compression engine.
- The Application Acceleration Summary table shows the following statistics for CIFS, Exchange, and HTTP acceleration.
 - **Total Time.** Number of seconds required to complete the transactions that ended in the selected time period for all clients.
 - **w/o Juniper WX.** Number of seconds that would have been required if acceleration was disabled.
 - **Time Saved.** Number of seconds saved by acceleration.

- The Percent Compression bar chart shows the percentage compression over the selected time period for the destination device.
- The TCP Acceleration 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 last four bar graphs show the following for all CIFS, Exchange, and HTTP transactions over the selected time period. Exchange graphs are shown for bulk read/write transactions and for all transactions.
 - **Actual** (grey bars). Number of seconds required to complete the transactions that ended in the time period for all clients.
 - **Savings** (orange bars). Number of seconds saved by acceleration during the time period.

TCP Acceleration Statistics

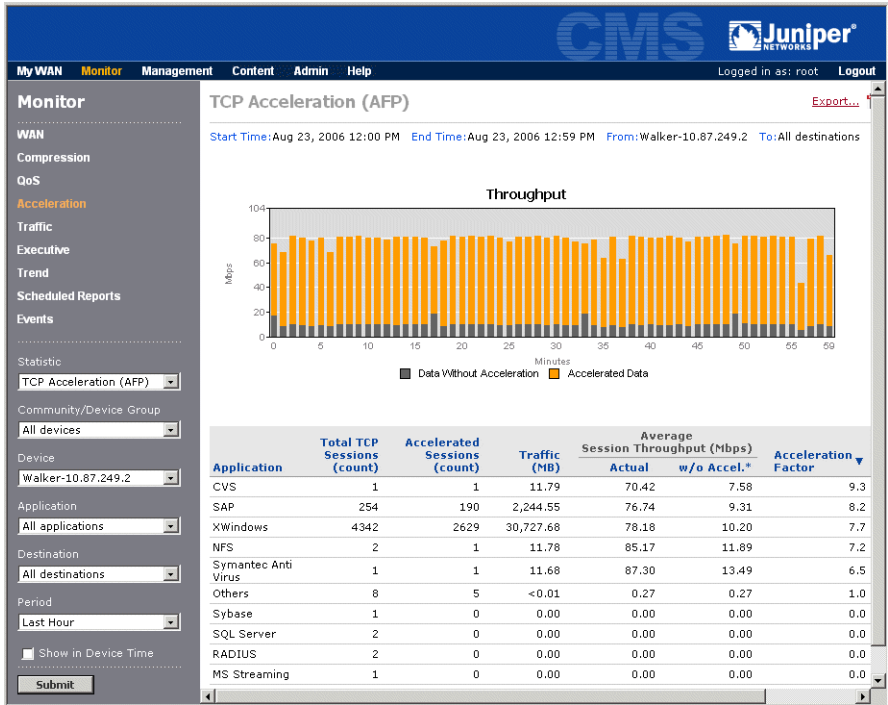
If TCP Acceleration is enabled for one or more endpoints and applications, the TCP Acceleration report shows the session statistics and the average throughput improvements due to TCP Acceleration.

To view TCP Acceleration statistics:

1. Click Monitor in the menu frame, and Acceleration in the navigation frame.
2. Select TCP Acceleration from the Statistic menu.
3. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a device from the Device menu that has acceleration enabled. Devices using acceleration have a  on the Devices page (click Management in the menu frame to view the Devices page).

- Select an application from the Application menu to view the acceleration statistics to each remote device. Select Others to view statistics for applications that are undefined or unmonitored. The default is All applications, which shows the average acceleration for all applications to all devices.
- Select the IP address of a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All destinations.
- Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 183: 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 is 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.
4. Click Export to view or save the tabular data in CSV format.

Fast Connection Setup Statistics

If Fast Connection Setup is enabled for one or more endpoints and applications, the Fast Connection Setup report shows the session statistics and the average percentage compression in session time due to Fast Connection Setup.

To view Fast Connection Setup statistics:

1. Click Monitor in the menu frame, and Acceleration in the navigation frame.
2. Select Fast Connection Setup from the Statistic menu.
3. Select the following report parameters, and click Submit.


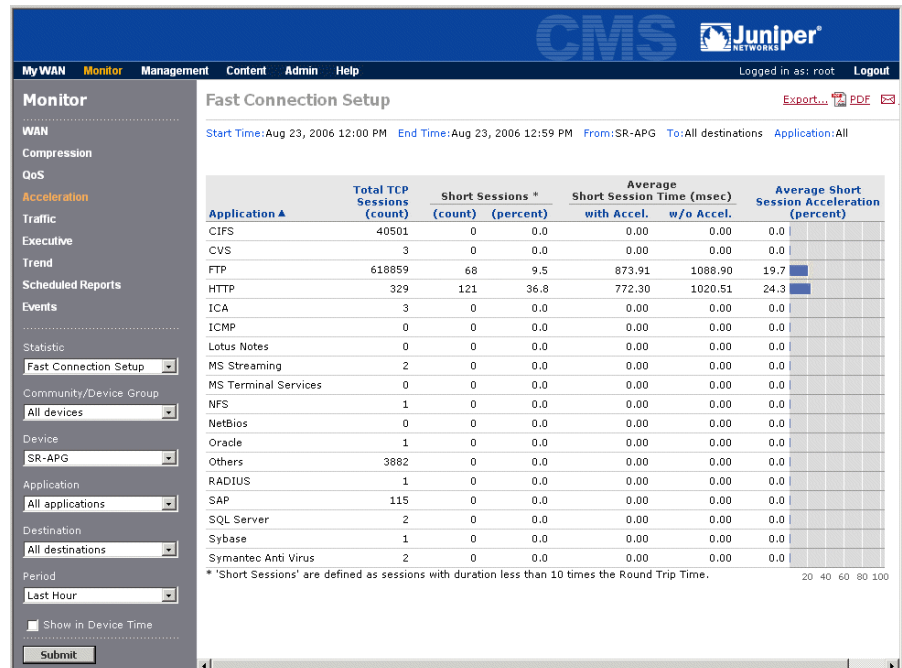
- Select a community or device group from the Community/Device Group menu.
- Select a device from the Device menu that has acceleration enabled. Devices using acceleration have a  on the Devices page (click Management in the menu frame to view the Devices page).
- Select an application from the Application menu to view the acceleration statistics to each remote device. Select Others to view statistics for applications that are undefined or unmonitored. The default is All applications, which shows the average acceleration for all applications to all devices.
- Select the IP address of a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All destinations.
- Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 184: Fast Connection Setup Statistics

4. Review the following information. Keep in mind that all values are for the selected application, destination, and time period.
 - **Application or Destination.** Name of the accelerated application(s) or, if you select a specific application, the IP addresses of each remote device.
 - **Total TCP Sessions.** Number of sessions that ended in the selected time period.
 - **Short Sessions.** Number of “short” TCP sessions accelerated, and the percentage of the total sessions. These columns show the relative number of sessions that benefit from Fast Connection Setup. Short sessions are those that last less than ten times the round-trip time (RTT). If a specific application traffic flow has five consecutive short sessions, subsequent identical traffic flows will be accelerated.
 - **Average Short Session Time (msec).** Average duration of the accelerated sessions (in milliseconds), versus what the average session time would have been if Fast Connection Setup was disabled.
 - **Average Short Session Acceleration (percent).** The average percentage 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.
5. Click Export to view or save the data in CSV format.

CIFS and Exchange Acceleration Statistics

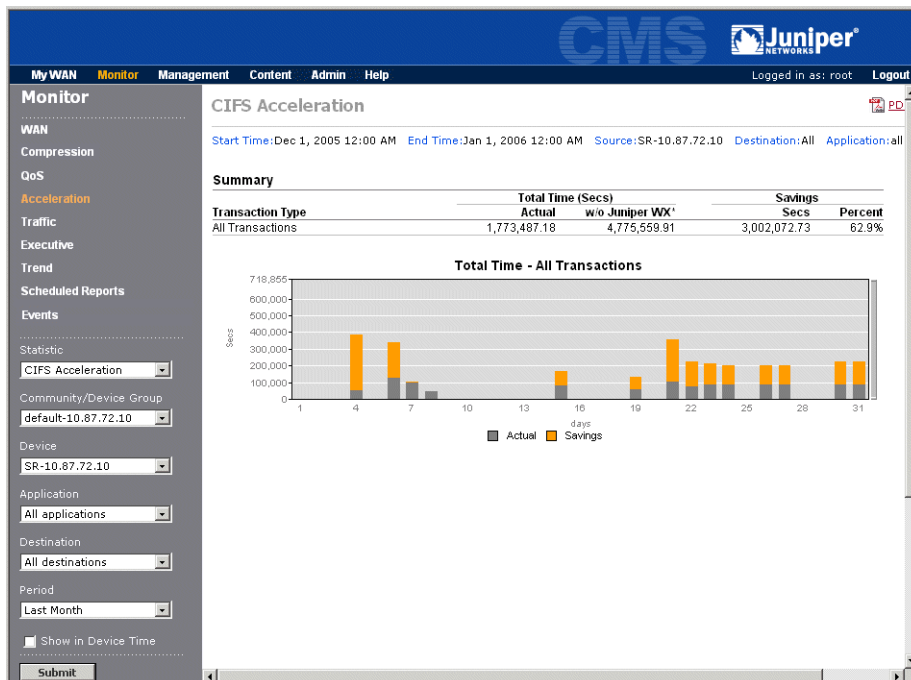
If CIFS or Exchange application acceleration is enabled for one or more application definitions on a device, the CIFS and Exchange acceleration reports shows the time saved due to CIFS and Exchange acceleration. TCP Acceleration must also be enabled.



NOTE: View CIFS and Exchange acceleration reports on the client-side device, not the server-side device. The acceleration statistics apply to the traffic in both directions. However, compression statistics should probably be viewed on the server-side device.

To view CIFS or Exchange acceleration statistics:

1. Click Monitor in the menu frame, and then click Acceleration in the left-hand navigation frame.
2. Select CIFS Acceleration or Exchange Acceleration from the Statistic menu.
3. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a specific device from the Device menu.
 - Select an application from the Application menu to view the acceleration statistics for a specific CIFS or Exchange application definition. The default is All applications.
 - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All destinations. If you select a specific destination, you can select the Show in Destination Time check box to view the report in the destination device's time (acceleration is measured from the destination device).
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 185: 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. The Exchange Acceleration report also shows bulk read/write 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 graph for all transactions shows the following (Exchange reports also have a graph for bulk read/write 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

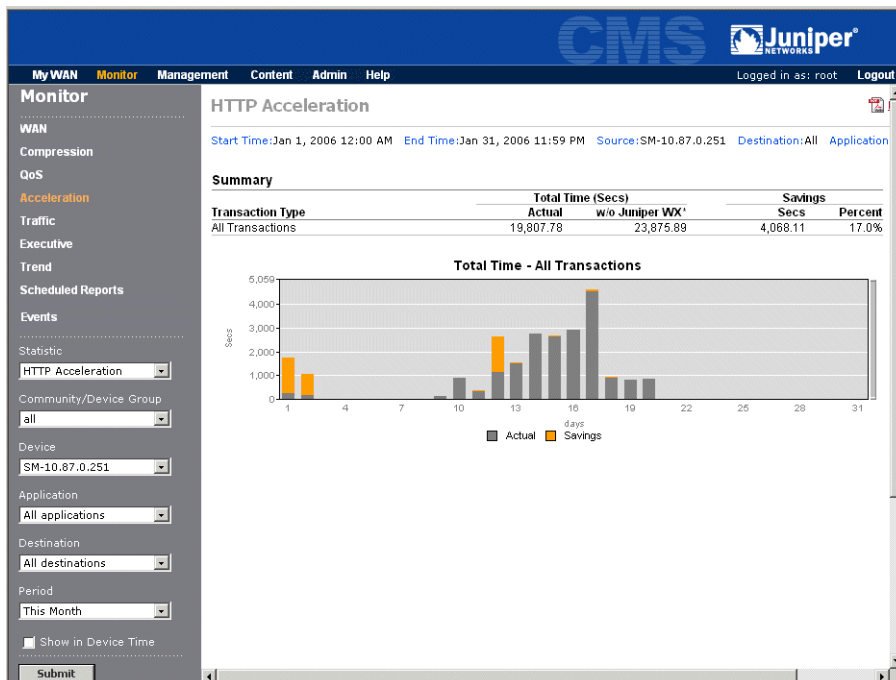
If HTTP acceleration is enabled for one or more application definitions, the HTTP Acceleration report shows the amount of time saved by HTTP acceleration. TCP Acceleration must also be enabled.



NOTE: View HTTP acceleration reports on the client-side device. The acceleration statistics apply to the traffic in both directions. However, 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.
3. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a specific device from the Device menu.
 - Select an application from the Application menu to view the acceleration statistics for a specific HTTP application definition. The default is All applications.
 - Select a specific device from the Destination menu to view statistics only for traffic sent to the selected device. The default is All destinations. If you select a specific destination, you can select the Show in Destination Time check box to view the report in the destination device's time (acceleration is measured from the destination device).
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 186: HTTP Acceleration Statistics

Review the following information. Keep in mind that all values are for the selected application, destination, and time period.

- The Summary table shows the following statistics for all transactions.
 - **Total Time.** Number of seconds required to complete the transactions that ended in the selected time period for all clients, and the number of seconds required if acceleration was disabled.
 - **Savings.** Amount of time saved by acceleration, shown in seconds and as a percentage of the time required if acceleration was disabled.
- The Total Time graph shows the following for all transactions:
 - **Actual** (grey bars). Number of seconds required to complete the transactions that ended in the time period for all clients.
 - **Savings** (orange bars). Number of seconds saved by acceleration during the time period.

Top Traffic Statistics

Each device collects statistics for its most active traffic flows, including the application name and protocol, source and destination addresses and ports, and the number of packets and bytes sent and received. The collected statistics can be sent to a Cisco NetFlow server and displayed in the Traffic report.

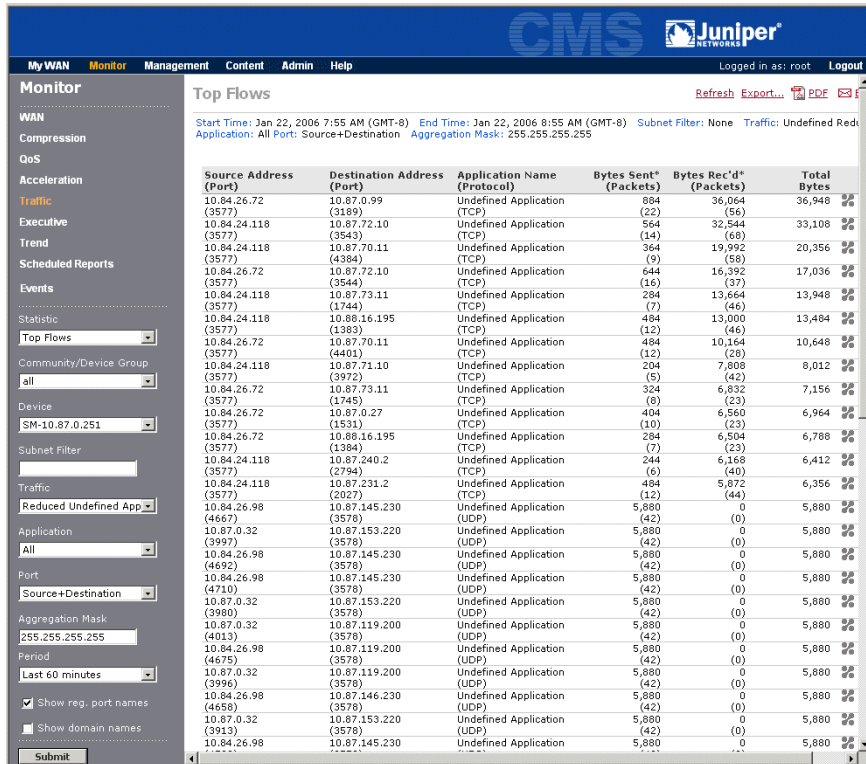
You can view the top traffic statistics for the past hour, the past 24 hours, or all available hours (the length of time depends on the traffic volume). The 65,000 most active flows are recorded. You can view the top 50 flows in the Web console, but the complete list can be exported to a file in CSV format.




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

To view the Traffic statistics for a device:

1. Click Monitor in the menu frame, and click Traffic in the left-hand navigation frame.
2. Select a community or device group from the Community/Device Group menu, select a device from the Device menu, and click Submit to view the top traffic flows for the past hour.

Figure 187: Top Traffic Statistics

Note that an  is shown next to the flows for undefined applications that are being compressed.

- 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	<p>If you select the top flows, sending addresses, or receiving addresses, you can enter a subnet to view just the traffic from that subnet. 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>
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-line devices or to in-line devices that use tunnel switching.
Application	Select an application to view the traffic for a specific application.
Port	<p>If you select the top flows, you can select a view of the port information.</p> <ul style="list-style-type: none"> ■ Ignore Port. Traffic is consolidated across all ports for each pair of source and destination addresses. ■ Source Only. Traffic is consolidated across the same source ports for each pair of source and destination addresses. ■ Destination Only. Traffic is consolidated across the same destination ports for each pair of source and destination addresses. ■ Source + Destination. Traffic is shown for each combination of source and destination port.
Aggregation mask	<p>If you select the top flows, sending addresses, or receiving addresses, you can enter a subnet mask to view all traffic from the same subnet as one consolidated entry. The default mask (“255.255.255.255”) shows a separate flow for each host. You can also use the “/n” format. For example, enter “/24” or “255.255.255.0” to consolidate all traffic flows with the same 24-bit network address.</p>
Period	<p>Select the time period (last 60 minutes, last 24 hours, or all). Note that if you select Last 60 minutes or Last 24 hours, only the traffic flows that started in the selected time period are shown.</p>
Show reg. port names	<p>If you select the top flows, click the check box to view the registered names for all ports in the collected data. Clear the check box to view the names only for port numbers up to 1024.</p>
Show domain names	<p>If you select the top flows, click the check box to view the domain names for each IP address. The DNS server used is the one specified on the Windows machine where CMS is installed. The IP address is displayed if its domain name cannot be resolved (the DNS queries may take a few seconds).</p>

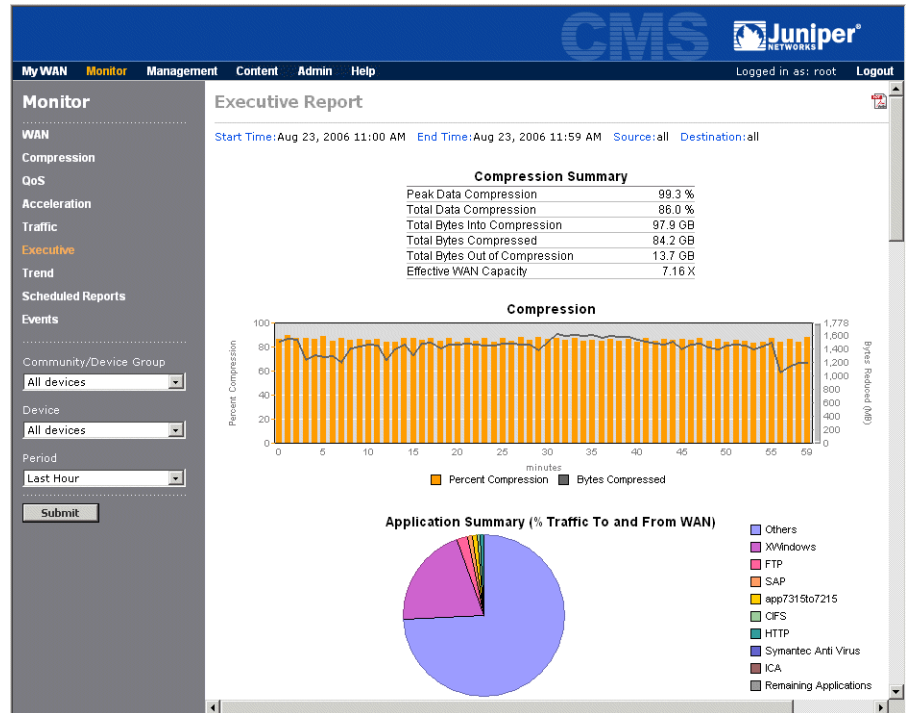
4. To export the traffic statistics to a file in CSV format, click Export. in the upper-right corner of the page.

Executive Summary

The Executive report summarizes compression results, traffic volume by application, and average WAN performance (latency and loss) for one or all remote devices.

To view the Executive statistics:

1. Click Monitor in the menu frame, and click Executive in the left-hand navigation frame.
2. Select the following report parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu.
 - Select a specific device from the Device menu. The default is All devices.
 - Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range.

Figure 188: Executive Summary

3. Review the following information.

- The Compression Summary table shows the following:
 - **Peak Data Compression.** Highest percentage of 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 compression engine.
 - **Total Bytes Compressed.** Number of bytes compressed.
 - **Total Bytes Out of Compression.** Number of bytes of traffic output after compression.

- **Effective WAN Capacity.** Factor increase in WAN capacity resulting from the total compression. For example, this value is 2.00 if total compression is 50 %.
- The Compression graph shows the average percentage of compression and the number of bytes compressed for the selected time period.
- The Application Summary pie chart shows the nine monitored applications with the highest percentage of the total traffic sent to and from the WAN for all destinations. The Remaining applications category shows the traffic for all other applications (both defined and undefined). Move the cursor over the legend to view the number of bytes for each application.
- The Volume by Application graph shows the traffic volume over the selected time period for the top nine monitored applications, plus the Remaining applications category.
- The Path Latency, Loss, and Availability distribution charts show the distribution of the average WAN latency, loss, and availability values measured for all destinations by devices that have WAN performance monitoring or Policy-Based Multipath enabled. To change the performance ranges represented by each color, refer to “Setting WAN Performance Thresholds” on page 418.

Trend Reports

The trend reports let you project future WAN traffic, throughput, loss, and latency for a WX device based on past performance for a specified date range. You can also view throughput trends by traffic class.

Trend lines are calculated using the least-squares, best-fit method. Times in the base period that have no data are shown as gaps in the performance curve (rather than zeros), and do not affect the trend line.

To view the trend reports:

1. Click Monitor in the menu frame, and click Trend in the left-hand navigation frame.
2. Select one of the following trend reports from the Statistic menu:

- **WAN Throughput.** Projected throughput to and from the WAN for a selected device to one or all remote WX destinations for one or all applications. You can also select “Other destinations” to project throughput to non-WX endpoints.
 - **QoS Class.** Projected throughput for a selected traffic class from a selected device to a remote WX destination.
 - **Application Volume.** Projected traffic volume to and from the WAN for a selected device to one or all remote WX destinations for one or all applications.
 - **WAN Latency.** Projected round-trip time (RTT) from a selected device to a remote WX destination.
 - **WAN Loss.** Projected percentage of probe loss from a selected device to a remote WX destination.
3. Select the following report parameters, and click Submit.
- Select a community or device group from the Community/Device Group menu.
 - Select a specific device from the Device menu.
 - For QoS Class trends, select a traffic class from the Class menu.
 - For WAN Throughput or Application Volume trends, select a monitored application from the Application menu. Select Others to view trends for applications that are undefined or unmonitored. The default is All applications.
 - For WAN Loss, WAN Latency, or QoS Class trends, select a specific device from the Destination menu. For WAN Throughput or Application Volume trends, the default is All destinations.
 - Select a base time period from the Base Period menu. Select the previous 30 to 180 days, or select Absolute and enter a start date.
 - Select a future time period from the Extrapolation Period menu. Select the next 30 to 180 days, or select Absolute and enter a future end date.

Figure 189: Trend Report for WAN Throughput

Events Reports

The Events reports list CMS “Device Task Failed” events and WX system and performance events for devices running WXOS 5.4 or later. System events for devices running WXOS 5.2 or 5.3 can be viewed only on the Devices page (refer to “Viewing Devices” on page 38). Acknowledging events removes them from the reports, but all past events can be viewed until they are purged.

Note the following:

- WXOS 5.4 devices must identify the CMS server as a Syslog server (refer to “Defining Syslog Servers” on page 124). CMS discards events from devices that are not in an imported community.
- To generate performance events, WXOS 5.4 devices must specify thresholds for the appropriate metrics. You can also enable or disable specific system events (refer to “Configuring Events” on page 280).
- To specify when events are purged, refer to “Configuring Device Polling” on page 405.

To view the Events report:

1. Click Monitor in the menu frame, and click Events in the left-hand navigation frame.
2. Select the report type from the Display Type menu:
 - **Matrix.** Displays a matrix of color-coded cells that indicate the highest severity performance event that occurred between each pair of devices in the selected community or device group. The device icons in the From column indicate the highest severity system event or CMS “Device Task Failed” event that occurred on each device. Select a device icon or a matrix cell to view the associated list of events.
 - **Console.** Displays a tabular list of up to 500 of the last WX and CMS events that occurred in the selected community or device group. A colored icon indicates the severity of each event. If you have more than 500 events (indicated at the top of the page), use the Period menu or other filters to view the older events.
3. Change the following report parameters, as needed, and click Submit.

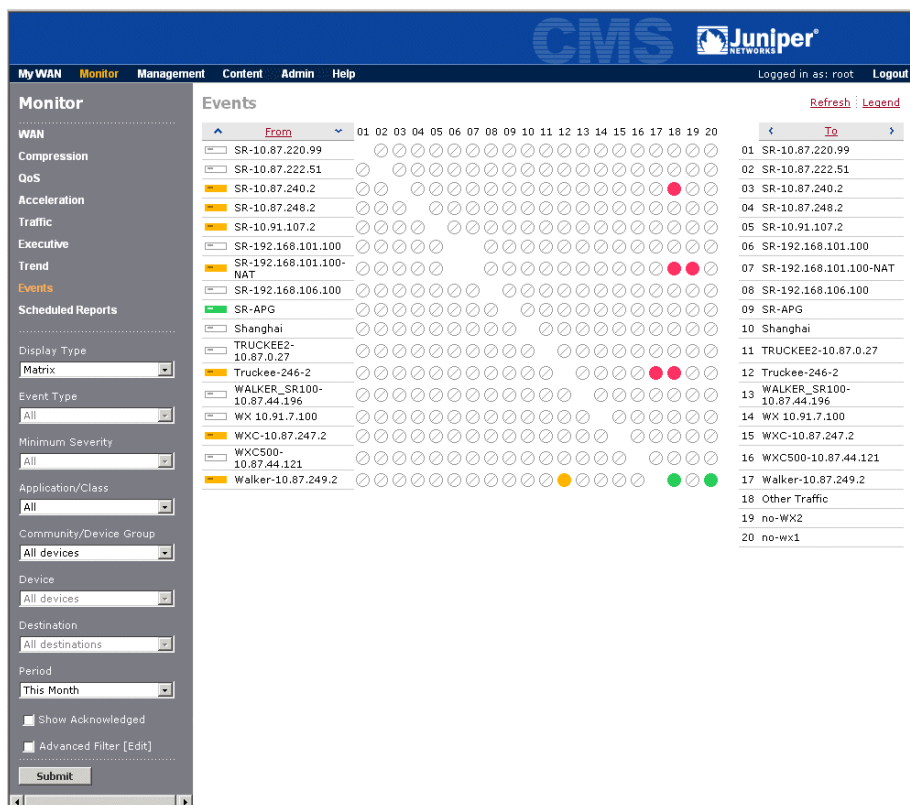
For a Matrix display, you can:

- 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.
- Select a community or device group from the Community/Device Group menu.
- Select a time period from the Period menu. Select the previous hour, day, week, month, or six months, or enter an absolute date range. The default is Last 60 Minutes.
- Select the Show Acknowledged check box to include acknowledged events.
- To view events for specific metrics, select the Advanced Filter check box, click Edit, select one or more metrics, and click OK. By default, events for all metrics are displayed.

For a Console display, you can also:

- Select the event type displayed (WX Performance, WX System, or WX CMS) from the Event Type menu. The default is All.

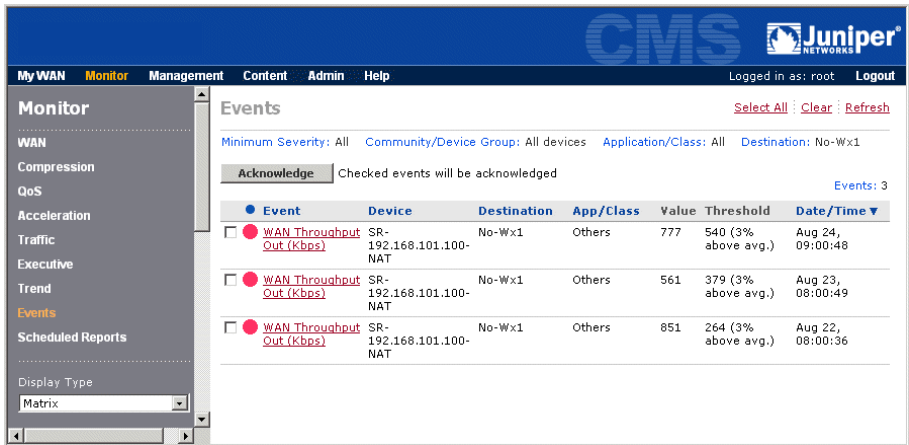
- Select the minimum severity level of the events displayed (OK, Warning, Major, or Critical), which correspond to the event color codes of green, yellow, orange, and red. The default is All.
 - Select a specific device from the Device menu. The default is All devices.
 - If you select WX Performance for the event type, you can select a device from the Destination menu to view performance events measured from for the selected destination. The default is All destinations.
4. If the Matrix display type is selected, click Submit to open the Events matrix. for the selected community or device group.

Figure 190: Events Matrix

From the Events matrix, you can:

- Click Legend in the upper-right corner to view the colors used for each severity level.
 - Identify the devices that have the highest severity events. The device icons in the From column indicate the highest severity of the system events and CMS “Device Task Failed” events that have occurred for each device. The matrix cells indicate the highest severity of the performance events that were measured by a device in the From column to a device in the To column.
 - Move the cursor over a colored device icon or matrix cell to view a description of the highest severity event.
 - View the next group of devices by moving the cursor over the From or To column headers and selecting a range of devices. You can also view the next or previous group of devices by clicking the arrows in the headers.
5. Click a colored cell in the Events matrix to view the list of performance events that were measured from the device in the From column to the device in the To column. The same event list is displayed if you select the Console display type, and then select the From device from the Device menu and the To device from the Destination menu.

Figure 191: Performance Events List



From the performance Events list, you can:

- Acknowledge events. Select the check box next to the appropriate events and click Acknowledge. Acknowledged events are removed from the list of events.
- Review the performance event information, including the event name, destination, application or traffic class, performance threshold, the value that violated the threshold, and the date and time the event occurred.
- Click the event name to view the evaluation time period (hourly or daily), whether the evaluation is limited to prime time days and hours, and a list of the related events.

Performance events are related if they are for the same metric, device, destination, application or traffic class, and are in the same time period.

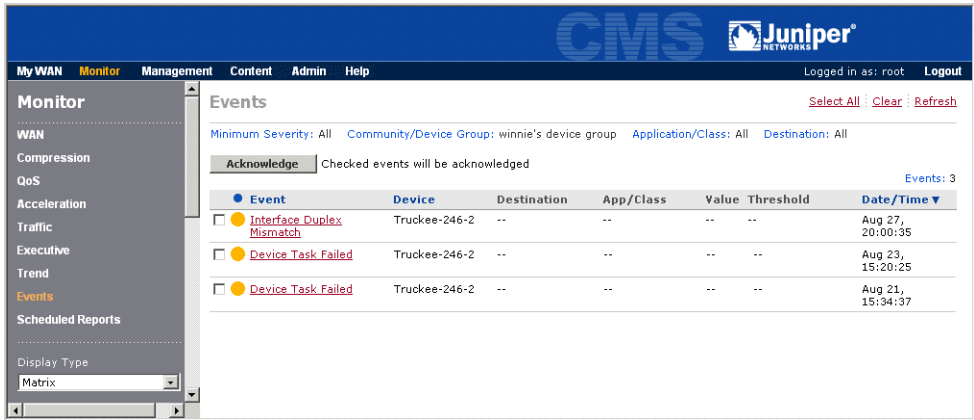
Figure 192: Performance Event Details

Event Details				
Metric	WAN Throughput Out (Kbps)			
Community	All devices			
Device	SR-192.168.101.100-NAT			
Destination	No-Wx1			
Application	Others			
Measured Value	777			
Threshold Value	540 (3% above avg.)			
Event Date/Time	Aug 24, 09:00:48			
Start Date/Time	Aug 23, 09:00:00			
Period	Daily			
Prime Time	Yes			
Severity	Critical			
				Total: 3
Other Instances	Value	Threshold	Event Date/Time ▼	Start Date/Time
	777	540 (3% above avg.)	Aug 24, 09:00:48	Aug 23, 09:00:00
	561	379 (3% above avg.)	Aug 23, 08:00:49	Aug 22, 08:00:00
	851	264 (3% above avg.)	Aug 22, 08:00:36	Aug 21, 08:00:00

6. Click a colored device icon in the From column of the Events matrix to view the list of WX system events and CMS “Device Task Failed” events that occurred on the selected device. Note that if you select the Console display type, and then select the device from the Device menu, you can view all event types, just CMS events, or just WX system events.

For a description of all the CMS and WX system events, refer to “System Events” on page 441.

Figure 193: System Events List

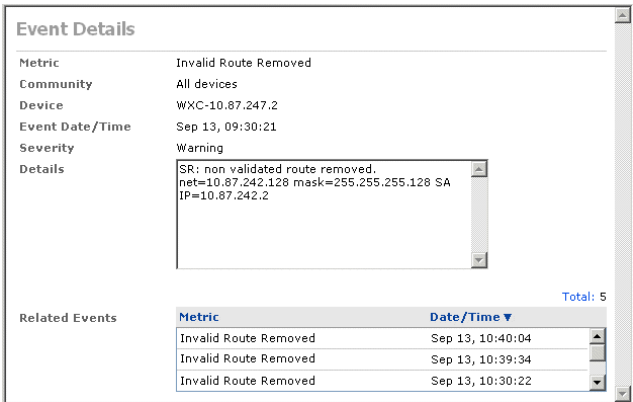


From the system Events list, you can:

- Acknowledge events. Select the check box next to the appropriate events and click Acknowledge. Acknowledged events are removed from the list of events.
- Click a WX system event name to view additional event details, and a list of related events (Figure 194). System events are related if they are for the same device, are in the same time period, and have the same or related metric. Metrics are related if they are inverses of each other, such as WAN Link Up and WAN Link Down.

Selecting a CMS “Device Task Failed” event opens the Schedule Details page where you can acknowledge or reschedule the failed CMS task (refer to “Managing Scheduled Tasks” on page 69).


Figure 194: System Event Details

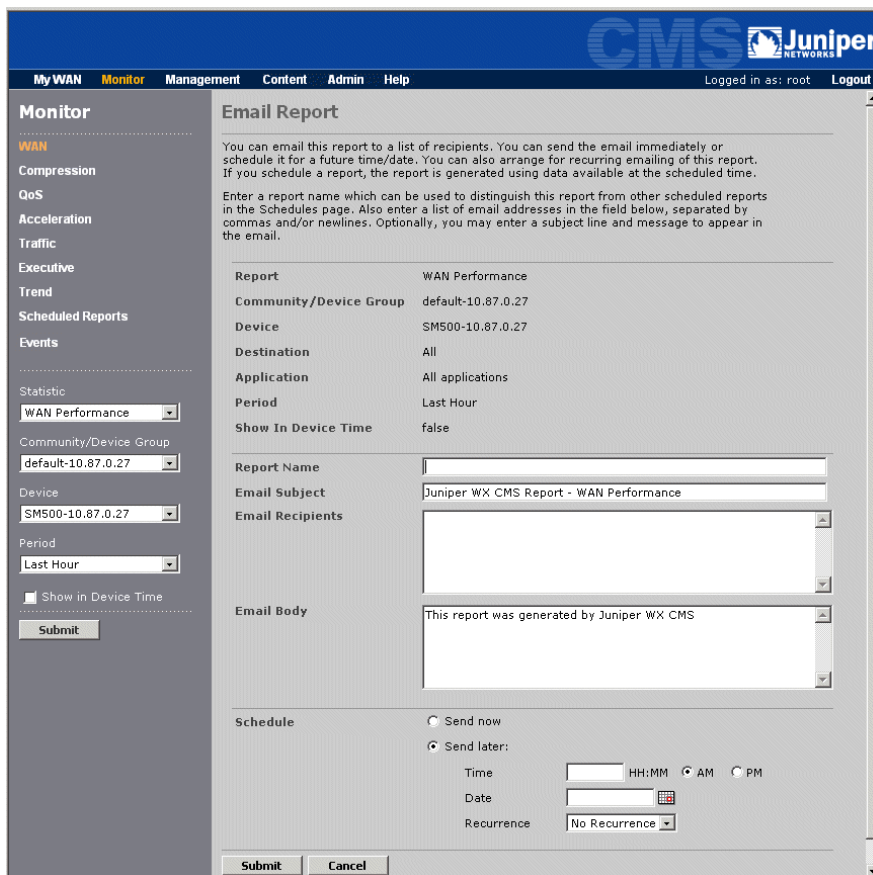


Scheduling Reports

Any monitoring report that can be displayed as a PDF file can be scheduled for email distribution on a one-time or recurring basis. This excludes the Events reports and the WAN Performance, Compression Overview, and Tunnel Status reports when All devices is selected. An email server must be defined in CMS (refer to “Defining an SMTP Server” on page 427). To view the scheduled reports, refer to “Managing Scheduled Reports” on page 369.

To email a report:

1. Click Monitor in the menu frame, and select the report you want to email.
2. Click  **Email** in the upper-right corner of the report to open the Email Report page.



Monitor

WAN

Compression

QoS

Acceleration

Traffic

Executive

Trend

Scheduled Reports

Events

Statistic

WAN Performance

Community/Device Group

default-10.87.0.27

Device

SM500-10.87.0.27

Period

Last Hour

Show in Device Time

Submit

Email Report

You can email this report to a list of recipients. You can send the email immediately or schedule it for a future time/date. You can also arrange for recurring emailing of this report. If you schedule a report, the report is generated using data available at the scheduled time.

Enter a report name which can be used to distinguish this report from other scheduled reports in the Schedules page. Also enter a list of email addresses in the field below, separated by commas and/or newlines. Optionally, you may enter a subject line and message to appear in the email.

Report	WAN Performance
Community/Device Group	default-10.87.0.27
Device	SM500-10.87.0.27
Destination	All
Application	All applications
Period	Last Hour
Show In Device Time	false

Report Name

Email Subject

Juniper WX CMS Report - WAN Performance

Email Recipients

Email Body

This report was generated by Juniper WX CMS

Schedule

☐ Send now

☒ Send later:

Time

HH:MM

AM

PM

Date


Recurrence

No Recurrence

Submit

Cancel

3. Specify the following information, and click Submit:

Report Name	Enter the name (up to 64 characters) to be shown on the Scheduled Reports page (refer to “Managing Scheduled Reports” on page 369).
Email Subject	Enter the text to appear in the subject line of the email. The default text indicates the name of the report.
Email Recipients	Enter the email addresses of the recipients, separated by commas. You can enter multiple lines of addresses.
Email Body	Enter the text to appear in the body of the email.
Schedule	<p>Select Send now or select Send later and enter a future time and date (in local CMS time):</p> <ul style="list-style-type: none"> ■ Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. ■ Enter a date in MM/DD/YYYY format or click  and select the month and date.

Managing Scheduled Reports

The Scheduled Reports page displays all reports scheduled for email distribution, including reports that are pending, successful, or failed. Scheduled reports can be canceled or changed by the user who scheduled the report, a CMS administrator, or the user group administrator. Only the schedule and email parameters can be changed (not the report). Canceled reports are deleted from the page.

If a scheduled report fails, move the cursor over the highlighted “Failed” status to view a description of the problem.

To manage scheduled reports:

1. Click Monitor in the menu frame, and then click Scheduled reports in the left-hand navigation frame.
2. Change one or more of the following parameters, and click Submit.
 - Select a community or device group from the Community/Device Group menu, or select All communities to view the reports scheduled for all devices.
 - Select a name from the Device menu to view scheduled reports only for the selected device. The default is All devices.

- Click the check box next to the status of the scheduled reports you want to view, such as Failed or Pending. Click Show Recurring Only to view scheduled reports that have the selected status AND are run periodically.

Figure 195: Scheduled Reports Page

The screenshot shows the 'Scheduled Reports' page in the CMS. The left sidebar has a 'Monitor' section with a 'Scheduled Reports' link. The main content area displays a table of reports. The table has columns: Name, Creation Time, User, Scheduled Time, Recipients, and Status. The first two reports, 'Top Sending Addresses' and 'Top Flows long report', are recurring (indicated by a circular arrow icon) and have a 'CANCEL' button next to their status. The other two reports, 'executive' and 'App Accl Summary', are one-time reports and have a 'Success' status.

Name	Creation Time	User	Scheduled Time	Recipients	Status
Top Sending Addresses	Dec 7, 2005 3:04 PM	root	Jan 18, 2006 6:00 PM	1	Success CANCEL
Reduction by time	Dec 7, 2005 2:55 PM	root	Jan 17, 2006 5:35 PM	1	Success
Top Flows long report	Jan 17, 2006 3:40 PM	root	Jan 17, 2006 4:00 PM	1	Success CANCEL
executive	Jan 13, 2006 4:56 PM	hc_dm	Jan 13, 2006 5:00 PM	1	Success
App Accl Summary	Jan 13, 2006 4:54 PM	hc_dm	Jan 13, 2006 4:54 PM	1	Success

The Scheduled Reports page provides the following information for each task:

- Name**—Name of the scheduled report. The icon indicates a recurring schedule. Move the cursor over the icon to view the frequency and the next run time.
- Creation Time**—Date and time the report was submitted.
- User**—ID of the user who submitted the report.
- Scheduled Time**—Date and time that the report is scheduled. For a recurring schedule that has run once, this is the scheduled time of the last run.
- Status**—The status of the scheduled report. For a recurring schedule that has run once, this is the status of the last run. If the status is “Failed”, move the cursor over the highlighted text to view a description of the problem.

3. To change the report schedule or email parameters, click the name of the scheduled report, make the necessary changes, and click Submit.
4. To cancel a pending or recurring scheduled report, click CANCEL. The scheduled report is deleted from the page.

Chapter 7

Content Management

This chapter describes how to use CMS to maximize user response times by preloading WXC compression dictionaries for files (content) that are large and/or frequently accessed over the Windows file system (CIFS).

- “Defining and Distributing Content” in the next section
- “Defining Distribution Groups” on page 378
- “Viewing the Status of Content Distribution Tasks” on page 380
- “Viewing the Schedule Log” on page 382
- “Accessing Network Drives” on page 383

Defining and Distributing Content

You can improve access times for large files by distributing the files in advance. The repeated patterns in the distributed files are added to the compression dictionaries of the specified WXC devices, so that compression occurs when the first user accesses the files through the Windows file system (CIFS).

The CMS server can send any number of files (of any size) to a source WXC, which then sends the files to one or more WXC destinations, such as from a central WXC server to one or more branches. For content that changes regularly, you can define a recurring schedule to automatically distribute the changed files.

Note the following:

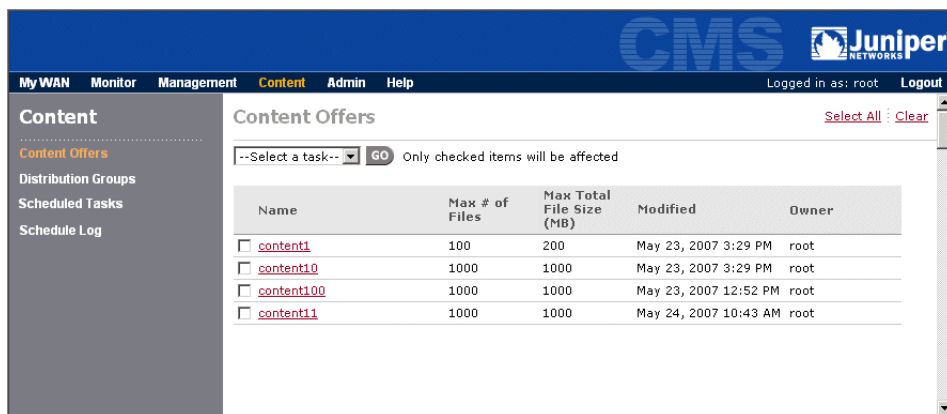
- You must define at least one distribution group to specify the source and destination WXC devices, such as WXC 500s, WXC 250s, and WXC 590s (refer to “Defining Distribution Groups” on page 378). Each destination group can specify up to 50 destinations.

- To distribute content located on remote network drives, you must change the default logon account of the JuniperCMS service (refer to “Accessing Network Drives” on page 383).
- TCP port 3578 is used to distribute content from the CMS server to the source WXC, and from the source WXC to the WXC destinations. Traffic sent to the destinations is included in the tunnel.
- On the source WXC, Juniper recommends creating an application definition for TCP traffic on port 3578. By default, content distribution traffic is managed by the Default QoS traffic class.
- On a WX 100 server, the dictionary updates are specific to the client that hosts the tunnel.
- Each user with a content management role can change or delete the content offers and distribution groups defined by all other users.

To define and distribute content:

1. Click Content in the menu frame, and then click Content Offers in the left-hand navigation frame.

Figure 196: Defining Content Offers



From the Content Offers page, you can:

- Define a new source of content for distribution (a content “offer”), as described in Step 2.
- Schedule a content offer for immediate or future distribution, as described in Step 3.

- Change a content offer. Click the content name, make any needed changes, and click Submit.
- Delete content offers. Select the check box next to one or more offers (or click Select All), select Delete from the task menu, and click Go. If you delete a content offer that has been scheduled, the distribution will fail. To cancel a scheduled distribution, refer to “Viewing the Status of Content Distribution Tasks” on page 380.

2. To define a source of content for distribution:

- a. Select New from the task menu, and click Go.

Figure 197: Adding a Content Offer

The screenshot shows the 'Content Offers > New' page in the Juniper CMS. The left sidebar contains links for Content, Content Offers, Distribution Groups, Scheduled Tasks, and Schedule Log. The main area contains a form for adding a new content offer. The form includes the following fields and options:

- Name:** A text input field.
- Path:** A text input field.
- Depth:** A dropdown menu set to '1' with the label 'directories'.
- Maximum # of Files:** A text input field set to '1000' with the label 'count'.
- Maximum Total File Size:** A text input field set to '2000' with the label 'MB'.
- Filter By:** A section with three checkboxes:
 - ☐ File Extension: Set to 'Accept' and 'DOC'.
 - ☐ Minimum File Size: Exclude files smaller than '10' MB.
 - ☐ Maximum File Size: Exclude files larger than '2000' MB.

At the bottom of the form are three buttons: Submit, Reset, and Cancel.

- b. Specify the following information, and click Submit:

Name Enter a name for the content (up to 32 characters).

Path Enter the name of a local or network file or directory, such as “c:\public” on the CMS server, “\\server\public”, or “\\10.20.30.1\public”. To access a network drive, you must change the logon account of the JuniperCMS service (refer to “Accessing Network Drives” on page 383).

Note the following:

- Path names cannot end in “\”, and paths such as “\\10.20.30.1” and “\\10.20.30.1*. *” are not supported.
- Path names cannot include non-English characters.
- Mounted drives are not supported.

Depth	Select the number of directory levels to be included in the distributed content (1 to 6, or all). A depth of 1 (the default) indicates just the specified directory.
Maximum # of Files	Enter the maximum number of files included in the content (default is 1000).
Maximum Total File Size	Enter the maximum number of megabytes included in the content (default is 2000 MB).
Filter By	<p>Optionally, select the check box for one or more of the following filters:</p> <ul style="list-style-type: none"> ■ File Extension. Select Accept or Reject and select a file type, or select Enter List and enter one or more file extensions separated by commas. The list is not case-sensitive. ■ Minimum File Size. To exclude relatively small files, enter the minimum file size in megabytes. ■ Maximum File Size. To exclude excessively large files, enter the maximum file size in megabytes.



NOTE: For the upper, lower, and total megabyte limits, files within 1 MB (+ or -) of the limit are sent. For example, if the file size in the Windows Explorer is 35,654KB, the file is sent if the upper limit is 35 MB.


- c. Click Submit to enter the changes, or click Cancel to discard them.

3. To distribute a content offer:
 - a. On the Content Offers page, select the check box next to a content offer, select Distribute from the task menu, and click Go.

Figure 198: Distributing a Content Offer

The screenshot shows the Juniper CMS interface for distributing a content offer. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content' (highlighted), 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar shows 'Content' with sub-items: 'Content Offers', 'Distribution Groups', 'Scheduled Tasks', and 'Schedule Log'. The main content area is titled 'Content Offers > Distribute'. Below the title is a description: 'Distribute the content offer to the selected distribution group.' The form contains the following fields: 'Content Offer' (text input with value 'content1'), 'Distribution Group' (dropdown menu with value 'dist1'), and a 'Schedule' section. In the 'Schedule' section, 'Distribute now' is selected with a radio button. Below it, 'Distribute later:' is selected with a radio button. This section includes 'Start Time' (text input with 'HH:MM' and 'AM/PM' radio buttons), 'Start Date' (calendar icon), 'Recurrence' (dropdown menu with 'No Recurrence'), and a checkbox for 'Skip if no change in content'. At the bottom are 'Submit', 'Reset', and 'Cancel' buttons.

- b. Specify the following information, and click Submit:

- | | |
|--------------------|---|
| Distribution Group | Select a distribution group to specify the source WXC and up to 50 WXC destinations. To define new distribution groups, refer to “Defining Distribution Groups” on page 378. |
| Schedule | <p>Select Distribute now, or select Distribute later and enter a future time and date (in local CMS time):</p> <ul style="list-style-type: none"> ■ Enter the time in HH:MM format (HH is 0-12 and MM is 0-59), and click AM or PM. Note that midnight can be 0:0 AM or 12:00 AM. ■ Enter a date in MM/DD/YYYY format or click  and select the month and date. ■ Optionally, select a recurring schedule (Daily, Weekly, or Monthly), and select the Skip if no change check box to avoid distributing the content unless one or more files have changed. |

- c. Click Submit to enter the changes, or click Cancel to discard them.

To view the status of the content distribution tasks, refer to “Viewing the Status of Content Distribution Tasks” on page 380.

Defining Distribution Groups

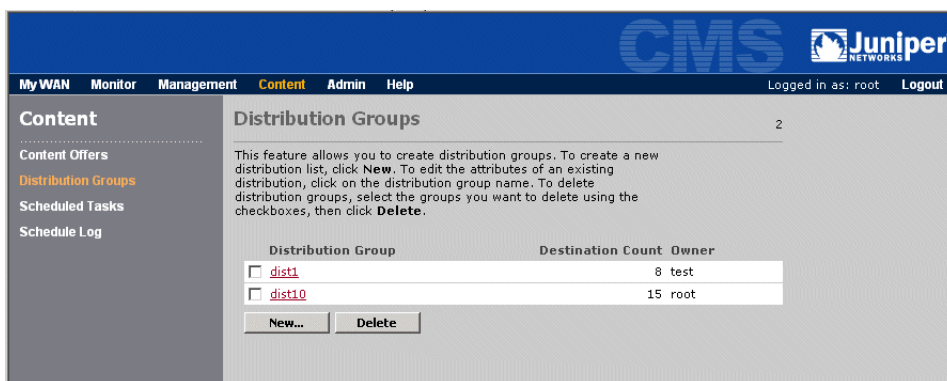
To improve access times for large files by distributing the files in advance, you must define at least one distribution group that specifies the source and destination WXC devices. Each destination group can specify up to 50 destinations.

After you define one or more distribution groups, you can schedule the content for distribution (refer to “Defining and Distributing Content” on page 373).

To define distribution groups:

1. Click Content in the menu frame, and then click Distribution Groups in the left-hand navigation frame.

Figure 199: Defining Distribution Groups



The Distribution Groups page lists the name and number of devices in each distribution group. The Owner column indicates the last user who changed the distribution group.

From the Distribution Groups page, you can:

- Define a new distribution group, as described in Step 2.
- Change a distribution group. Click the group name, make any needed changes, and click Submit.
- Delete distribution groups. Select the check box next to one or more groups, and click Delete. If you delete a group that is specified in a future scheduled distribution, the distribution will fail. To cancel a scheduled distribution, refer to “Viewing the Status of Content Distribution Tasks” on page 380.

2. To define a new distribution group:
 - a. Click New at the bottom of the page.

Figure 200: Adding a Distribution Group

- b. Specify the following information:

Distribution Group Name	Enter a name for the distribution group (up to 32 characters).
Community/Device Group	Select the community or device group of the source WXC device.
Source Device	Select the source WXC device. The device name and IP address are shown for each device in the selected community/device group. The IP address is enclosed in parentheses. If NAT is used, the IP address is the private address.

- c. Click Next to specify the destination WXC devices:
 - a. Select a community or device group from the Community/Device Group list.
 - b. Select the WXC destination devices, and click Add. To remove devices from the Destination Devices list, select the devices and click Remove.
 - c. Repeat Steps **a** and **b** for each community/device group (some devices may belong to multiple communities or groups).



NOTE: Do not specify more than 50 destinations. Any non-WXC devices in the group will cause errors during distribution.

- d. Click Submit to enter the changes, or click Back or Cancel to discard them.

To use a distribution group to distribute content, refer to “Defining and Distributing Content” on page 373.

Viewing the Status of Content Distribution Tasks

The Scheduled Tasks page lets you view the scheduled content distribution tasks, including tasks that are pending, running, sent, failed, or cancelled. To view or export the schedule log, refer to “Viewing the Schedule Log” on page 382.

Note that any user group administrator can cancel or reschedule any content distribution task. However, a rescheduled task will fail for any devices that are not in the administrator's user group(s).

To view or cancel content distribution tasks:

1. Click Content in the menu frame, and then click Scheduled Tasks in the left-hand navigation frame.

Figure 201: Viewing Content Distribution Status

Content Offer	Distribution Group	Destination Count	Creation Time	User	Scheduled Time	Status
network	winnie's group	12	Jun 25, 2007 2:32 PM	root	Jul 10, 2007 5:00 PM	Failed CANCEL
jenoffer1	jen_group1	2	Jul 9, 2007 5:43 PM	root	Jul 9, 2007 6:00 PM	Sent
jenoffer1	jen_group1	2	Jul 9, 2007 5:42 PM	root	Jul 9, 2007 6:00 PM	Cancelled
jenoffer1	jen_group1	2	Jun 26, 2007 12:50 PM	tadmin102	Jun 26, 2007 12:50 PM	Sent

The Schedules page includes the following information for each task:

- **Content Offer**—Name of the content offer that specifies the content to be distributed. A preceding the name indicates a recurring task. Move the cursor over the icon to view the time interval (Daily, Weekly, or Monthly) and the next execution time.
- **Distribution Group**—Name of the distribution group that specifies the source and destination WXC devices.

- **Destination Count**—Indicates the number of destination WXC devices specified by the distribution group.
 - **Creation Time** and **User**—Date and time that the task was created and the user who created it.
 - **Scheduled Time**—Date and time that the task is scheduled. For a recurring schedule that has run once, this is the scheduled time of the last run.
 - **Status**—The status of the task. For a recurring task that has run once, this is the status of the last run. Note the following:
 - A “Failed” status indicates the task has failed for at least one destination.
 - A “Sent” status is reported when files are sent successfully to the source WXC or if no files are sent due to the current filter settings (refer to “Defining and Distributing Content” on page 373).
 - Synchronization problems between the source and destination WXCs are not reported to CMS. If you suspect a problem, check for “PRESYNC_PROXY not initialized” messages in the system log of the source WXC.
2. To cancel a pending or recurring task, click CANCEL next to the task. The task status is changed to “Cancelled”.
 3. To reschedule a failed or pending task, or to view the details of a task, click the task status.

Figure 202: Viewing Task Details

Schedule Details

Close

Action

Distribute Content Offer

Content Offer

network

Source Device

WXC590-10.87.245.2(10.87.245.2)

Distribution Group

winnie's group

Scheduled time

Jul 12, 2007 5:00 PM

Recurrence

Daily

Creation time

Jun 25, 2007 2:32 PM

Skip if no change in content

No

User

root

--Select a task--

GO

Device Name	IP Address	Status	History	Completed	Detail
10.87.241.2	10.87.241.2	Failed	VIEW	Jul 12, 2007 5:00 PM	Command Error: Can not access Source Path.
10.87.243.2	10.87.243.2	Failed	VIEW	Jul 12, 2007 5:00 PM	Command Error: Can not access Source Path.
10.87.244.107	10.87.244.107	Failed	VIEW	Jul 12, 2007 5:00 PM	Command Error: Can not access Source Path.

The Schedule Details page displays additional task information, such as the source WXC device, and lists the status and completion time for each destination device. Note that a blank status indicates the files were sent to the source WXC, but does not guarantee that the destination received them. If the content distribution failed, the reason is shown in the Detail column. For a recurring task, click VIEW to view the details of each run for a specific device.

To reschedule a pending task or a task that failed for all devices, select Reschedule from the task menu and click Go. You can also cancel the task by selecting Cancel from the task menu.

Viewing the Schedule Log

You can view or export a log containing information about the completed content distribution tasks. You can save the file in CSV format on a local disk, and then import its contents into a spreadsheet program (such as Microsoft Excel).

Each log entry contains the following information:

- Task identification number and task name (“Distribute Content Offer”)
- Name of the content offer and distribution group
- Source and destination IP address (one entry per destination)
- Date and time the task was scheduled, and the task creation time
- Status of the completed task (blank, “Failed”, or “Cancelled”). A blank indicates the files were sent to the source device, but does not guarantee that the destination received them.
- Date and time that the task was completed
- User who scheduled the task
- Details for failed tasks

To view or export the schedule log:

1. Click Content in the menu frame, and then click Schedule Log in the left-hand navigation frame.
2. To save the file to a local disk, click Save and specify the file location.

Accessing Network Drives

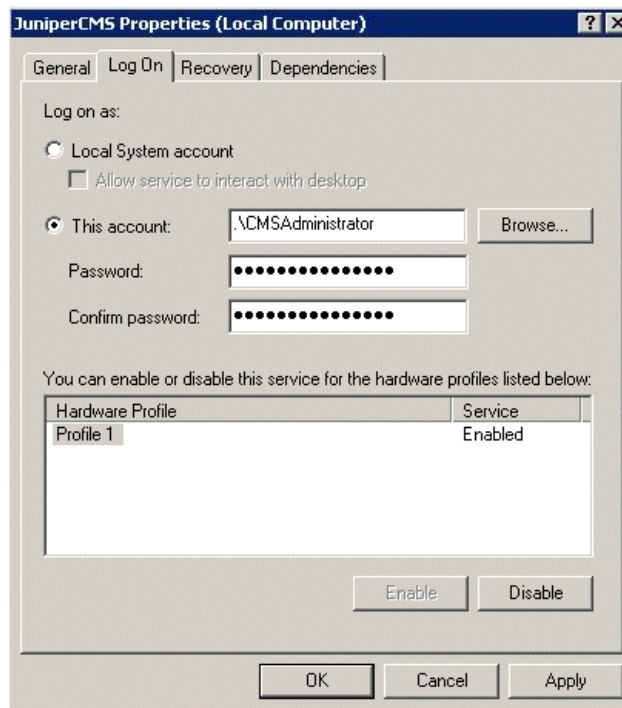
To distribute content stored on network drives, you must change the logon account of the JuniperCMS service to an account that has network privileges. The default Local System account has no network privileges.

To change the JuniperCMS logon account:

1. To create a new local or domain user account, do one of the following. Note that using the “root” account for the JuniperCMS service is not recommended for security reasons.
 - If the CMS server is a domain controller,
 - a. Select Start > Programs > Administrative Tools > Active Directory Users and Computers.
 - b. Navigate to Users in the left panel, right click on Users, and select New > User from the popup menu.
 - If the CMS server is NOT a domain controller,
 - a. Select Start > Programs > Administrative Tools > Computer Management.
 - b. Select System Tools > Local Users and Groups > Users in the left panel, right click on Users, and select New user from the popup menu.
2. Specify network privileges for the new account:
 - a. Go to User Manager for Domains.
 - b. Select Policies from the title bar menu, and select User Rights.
 - c. Select the option for Advanced User Rights.
 - d. In the drop-down list, verify that the following rights have been granted to the selected account:
 - Act as part of the operating system
 - Back up files and directories
 - Log on as a service
 - Restore files and directories

- e. Verify that the account is a member of the Users group.
3. Assign the new account to the JuniperCMS service:
 - a. Click Start > Run, enter “services.msc” and click OK.
 - b. In the Services window, right-click on JuniperCMS and click Stop.
 - c. Right-click on JuniperCMS, select Properties, and click the Log On tab.

Figure 203: Viewing Task Details



- d. Select This Account, enter the new account name, enter and confirm the account password, and then click OK.
- e. Restart the JuniperCMS service.

Chapter 8

CMS Setup and Administration

This chapter describes how to set up and administer CMS and covers the following topics:

- “Administering CMS Users” in the next section
- “Administering WX Devices” on page 397
- “Managing Event Forwarding Filters and Email Distribution Lists” on page 411
- “Administering CMS” on page 417

Administering CMS Users

The following sections describe the user-related administration tasks:

- “Partitioning Users by Customer” in the next section
- “Viewing My Account and Changing Passwords” on page 387
- “Defining CMS User Accounts” on page 388
- “Defining User Groups” on page 391
- “Defining User Roles” on page 393
- “Viewing Logged In Users” on page 396

Partitioning Users by Customer

Service providers can organize CMS users by customer, so that each user can view or change only the devices, configurations, and reports relevant to the associated customer. For each customer, the CMS administrator defines one or more user groups, and optionally, a user group administrator who can manage the users and devices for the specified user groups.

Use the following procedure to partition users for a customer:

1. Import the communities of WX devices from the registration server associated with the customer (refer to “Importing and Managing Communities” on page 397).
2. Add a user account that has the User Group Administrator role (refer to “Defining CMS User Accounts” on page 388).
3. Add a user group that specifies the new user account, the communities of devices the user can access, and whether the user has read or read/write access (refer to “Defining User Groups” on page 391). In general, the user group administrator should have read/write access to all customer devices.

A user group administrator can add more user accounts, user groups, and device groups, but is limited to the devices associated with the initial user group. User group administrators also can import communities from their registration servers and assign the communities to any of their user groups.

User group administrators can assign any user role to any of their users, except for the CMS admin user role (refer to “Defining User Roles” on page 393).

Viewing My Account and Changing Passwords

All CMS users can view their account settings. Users can also change their account password, depending on their account privileges (refer to “Defining CMS User Accounts” on page 388).

To view your account settings:

1. Click Admin in the menu frame, and then click My Account in the left-hand navigation frame.

Figure 204: Viewing My Account and Changing Passwords

2. To change your password:
 - a. Click Change Password.
 - b. In the Change Password page, type the current password, and then type the new password in the New Password and Verify New Password fields.
 - c. Click Submit to activate the new password.

Defining CMS User Accounts

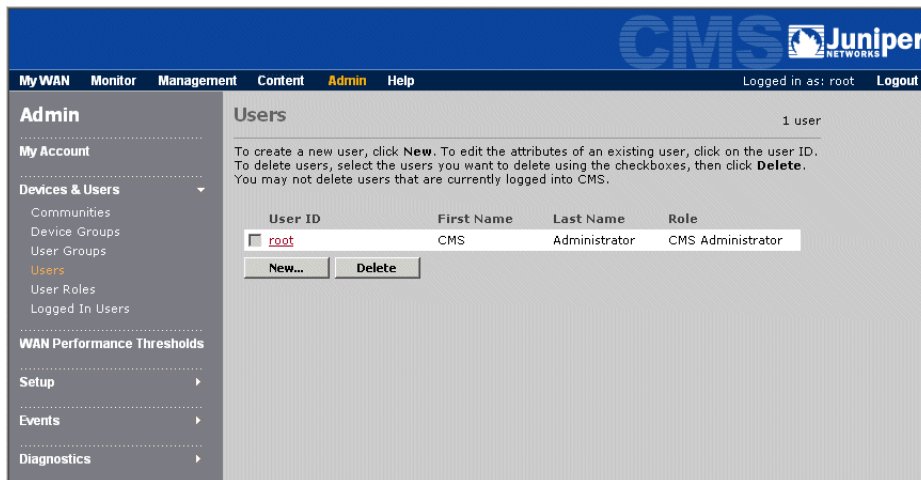
User accounts let you control access to CMS. The “role” assigned to each account determines the tasks the user can perform. The default “root” account has the CMS Administrator role (full access). Only a CMS administrator can change or add user roles, as described in “Defining User Roles” on page 393.

User groups determine the devices a user can access (read-only or read/write) and the configurations a user can view. A user cannot access any devices until the user’s account is assigned to at least one user group. To define user groups, refer to “Defining User Groups” on page 391. Only a CMS administrator or user group administrator can define user accounts and user groups.

To define CMS user accounts:

1. Click Admin in the menu frame, and then click Users in the left-hand navigation frame.

Figure 205: Administering CMS User Accounts



From the Users page, you can:

- Add new user accounts, as described in Step 2.
- Change a user account. Click the user ID, make any needed changes, and click Submit.

- Delete user accounts. Select the check box next to one or more accounts, and click Delete. The check boxes are grayed out for users who are logged in. You cannot delete the last CMS administrator account.
2. To add a new CMS user account, click New, specify the following information, and click Submit:

User ID	<p>Enter the account login name (up to 20 characters). In general, use only letters and numbers when defining user IDs. If necessary, you can use the following special characters:</p> <p>@ : # \$ & _ - / (.) ' ,</p>
First Name	Enter the user's first and last name (up to 32 characters each).
Last Name	Both names are required.
Role	<p>Select a role to specify the tasks the user can perform. The standard roles are:</p> <ul style="list-style-type: none"> ■ CMS Administrator. All tasks. Only a CMS administrator can assign this role to another user. ■ Content Administrator. Manage content distribution (refer to "Content Management" on page 373). ■ Device Administrator. Perform all device management tasks (refer to "Managing Devices" on page 37), edit configuration files, upload WXOS images, view monitoring reports, view and acknowledge events, manage content distribution, and schedule reports for email distribution. ■ Device Monitor. View monitoring reports, view events, and schedule reports for email distribution. ■ Device Operator. Manage and monitor devices, view and acknowledge events, and upload WXOS images. Cannot schedule reports. ■ User Group Administrator. All tasks for the user groups that the user belongs to. Excludes auto-deployment, license management, and the global CMS tasks described in "Administering CMS" on page 417 <p>The CMS Administrator role cannot be changed. To change the other roles or add new roles, refer to "Defining User Roles" on page 393.</p>
Authentication	<p>Select the check box to use a remote Active Directory or OpenLDAP server to authenticate the user. Note that CMS administrators and user group administrators must be authenticated locally.</p> <p>If you select this option, the user cannot be authenticated until you change the default authentication method (local) and identify the authentication server (refer to "Configuring AAA Settings for Remote Authentication" on page 421).</p>

Password	Enter the password twice (from 4 to 30 characters). Clear the check box to prevent the user from changing the password.
Member of	<p>Add the account to at least one user group by selecting the user group and clicking Add. Alternatively, you can add the account directly to the user group (refer to “Defining User Groups” on page 391).</p> <p>User groups provide read or read/write access to specific devices. Users cannot access any devices until their account is assigned to one or more user groups.</p> <p>NOTE: Since user groups allow access to devices by community or device group, access conflicts can occur (read vs. read/write) if a device belongs to multiple communities or device groups. These conflicts are resolved as follows:</p> <ul style="list-style-type: none"> ■ If the conflict occurs between two user groups that a user belongs to, the user receives read/write access. ■ If the conflict occurs within a single user group, the user receives read-only access.
Device Access URL	<p>For the devices of interest to the user, select the method used to generate URLs for the device name links shown on the Devices page. Clicking a link opens the WXOS Web console for the device.</p> <ul style="list-style-type: none"> ■ Use private IP address. If the devices are on the private side of a NAT device, the private IP address is used; otherwise the public IP address is used (the default). ■ Use public IP address. The public IP address is used for all devices. Use this option if the devices are on the private side of a NAT device, but the user's workstation is not in the private address space of the WX devices. ■ Use domain name. The URL is composed of the domain name and device name specified on the Addresses page of the WXOS Web console. The URL format is: https://<domain_name>/<device_name> The domain and device names can be downloaded from CMS in a Device Settings partial configuration (refer to “Configuring Device Addresses” on page 103). If a device does not report a domain name, only the device name appears in the URL. If both the domain and device names are not reported, the public IP address is used. The public IP address also is used if the reported device name contains periods. <p>For remotely authorized CMS users, the AD/LDAP server can be configured to return the type of URL access (refer to “Configuring AAA Settings for Remote Authentication” on page 421).</p>

Defining User Groups

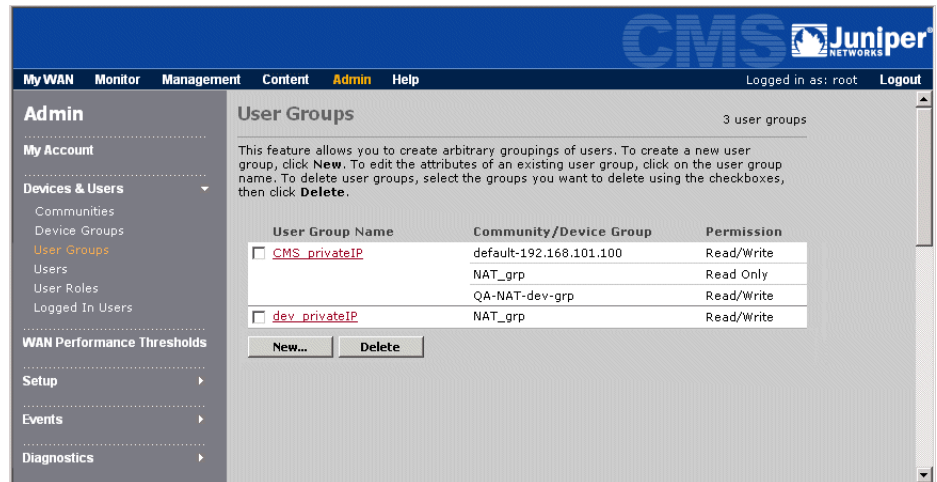
User groups specify the devices a user can access (read-only or read/write) and the configurations a user can view. A user cannot access any devices until the user's account is assigned to at least one user group.

Only a CMS administrator or user group administrator can define user groups. You can assign user accounts to user groups as described here, or you can select the user groups when you define the user account (refer to “Defining CMS User Accounts” on page 388).

To define user groups:

1. Click Admin in the menu frame, and then click User Groups in the left-hand navigation frame.

Figure 206: Defining User Groups



From the User Groups page, you can:

- Add new user groups, as described in Step 2.
- Change a user group. Click the group name, make any needed changes, and click Submit.
- Delete user groups. Select the check box next to one or more user groups, and click Delete. Deleted user groups are removed from all user accounts, and users in those groups may lose access to some or all of the devices specified in the deleted user groups.

- To add a new user group, click New.

Figure 207: Defining User Groups

Admin

My Account

Devices & Users

- Communities
- Device Groups
- User Groups
- Users
- User Roles
- Logged In Users

WAN Performance Thresholds

Setup

Events

Diagnostics

User Groups > Edit HELP

Members of the user group are permitted access to devices belonging to communities or device groups shown in the 'Devices' list.

Use the **Add/Remove** buttons to add or remove items from the 'Devices' or 'Members' lists. When you are done editing the lists, click **Submit**.

NOTE: Since device access permissions (Read-Only vs. Read/Write) are specified by community/device group, access conflicts can occur if a device belongs to multiple communities or device groups. These conflicts are resolved as follows:

If the conflict occurs within a single user group, the user receives Read-Only access.

If the conflict occurs between two user groups that a user belongs to, the user receives Read/Write access.

User Group Name

Devices	Community/Device Group	Permission
Add/Remove Devices...		

Members	User ID	User Name
Add/Remove Users...		

Submit **Cancel**

- Specify the following information, and click Submit:

- | | |
|-----------------|--|
| User Group Name | Enter the group name (up to 32 characters). |
| Devices | <p>To add communities and/or device groups to the user group:</p> <ol style="list-style-type: none"> Click Add/Remove Devices. Select one or more communities or device groups and click Add. When the list is complete, click Submit. <p>If you are a user group administrator, you can select only the communities or device groups associated with your user groups (the user groups that you belong to).</p> <ol style="list-style-type: none"> Select the Read Only or Read/Write permission for each of the selected communities and device groups, and click Submit. <p>To import communities from a registration server, refer to “Importing and Managing Communities” on page 397. To add more device groups, refer to “Managing Device Groups” on page 402.</p> |

Members

To add user accounts to the user group:

1. Click Add/Remove Users.
2. Select one or more accounts in the Users list and click Add. When the list is complete, click Submit.

If you are a user group administrator, you can select only the users that belong to your user groups.

NOTE: Since user groups allow access to devices by community or device group, access conflicts can occur (read vs. read/write) if a device belongs to multiple communities or device groups. These conflicts are resolved as follows:

- If the conflict occurs between two user groups that a user belongs to, the user receives read/write access.
- If the conflict occurs within a single user group, the user receives read-only access.

Defining User Roles

The role assigned to each account determines the tasks the user can perform. The devices a user can view and update depend on the user groups the user belongs to, and whether read/write access is enabled (refer to “Defining User Groups” on page 391).

CMS provides the following standard roles. All the standard roles can monitor devices, except the Content Administrator, and only the Content Administrator and Device Operator cannot schedule reports.

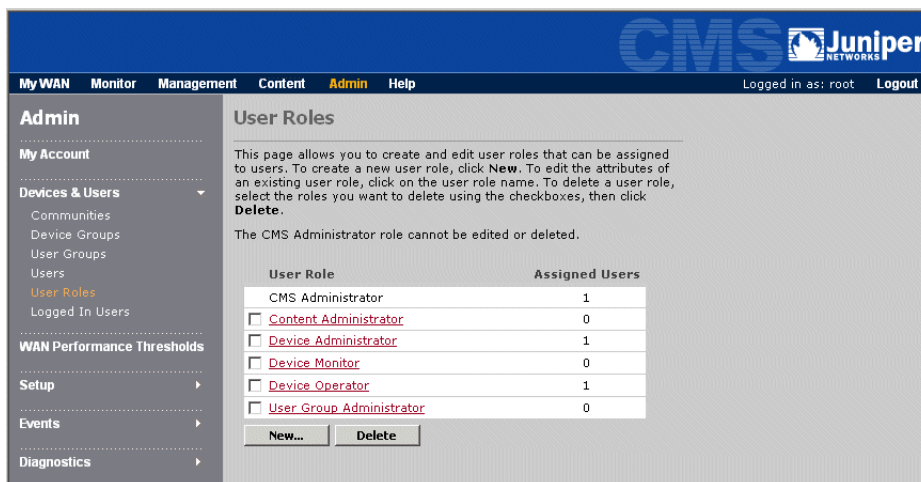
- **CMS Administrator.** Perform all CMS tasks (cannot be changed). Only a CMS administrator can change or add user roles, configure auto-deployment, manage licenses, or perform the tasks described in “Administering CMS” on page 417. The default “root” account has this role.
- **Content Administrator.** Manage content distribution (refer to “Content Management” on page 373).
- **Device Administrator.** Perform all device management tasks (refer to “Managing Devices” on page 37), edit configuration files, upload WXOS images, view monitoring reports, view and acknowledge events, manage content distribution, and schedule reports for email distribution.
- **Device Monitor.** View monitoring reports, view events, and schedule reports for email distribution.

- **Device Operator.** Manage and monitor devices, view and acknowledge events, and upload WXOS images. Cannot schedule reports.
- **User Group Administrator.** Perform all tasks for the user groups that the user belongs to (excludes auto-deployment, license management, and the tasks described in “Administering CMS” on page 417).

To define user roles:

1. Click Admin in the menu frame, and then click User Roles in the left-hand navigation frame.

Figure 208: Defining User Roles



From the User Roles page, you can:

- Add new user roles, as described in Step 2.
- Change a user role. Click the role name, make any needed changes, and click Submit. For changes that exclude previously allowed tasks, logged-in users are allowed to complete any task they have started.
- Delete user roles. Select the check box next to one or more roles, and click Delete. You cannot delete a role that is assigned to a user.

2. To add a new user role, click New, specify the following information, and click Submit:

User Role Name	<p>Enter the role name (up to 30 characters). In general, use only letters and numbers when defining names. If necessary, you can use the following special characters:</p> <p>: # \$ & _ - / () ' </p>
Monitor	<p>Select the Monitor check box and select the reports the user can view. You can also enable:</p> <ul style="list-style-type: none"> ■ View Scheduled Reports. View the reports scheduled for email distribution. ■ Manage Scheduled Reports. Cancel or change scheduled reports. ■ Acknowledge Events. If you enable the Events report, you can allow users to acknowledge events on the report. <p>If the WAN, Compression, QoS, Acceleration, or Events reports are enabled, the My WAN page is also enabled with access to the corresponding charts. If none of the monitoring options are checked, the Monitor and My WAN links are not shown in the Web console.</p>
Management	<p>Select the Management check box to allow users to view configurations, retrieve device files, reboot devices, apply a registration server password, and manage scheduled tasks.</p> <p>As appropriate, select additional options to allow users to edit, analyze, or backup/restore configurations, and load/rollback device images or configurations. If the Management box is not checked, the Management link is not shown in the Web console.</p>
Content	<p>Select the Content check box to allow users to view the schedule log for content distribution. You can also enable:</p> <ul style="list-style-type: none"> ■ Manage Content Offers. Define content offers and distribution schedules. ■ Manage Distribution Groups. Define the distribution groups that can be specified in content offers. ■ Schedule Content Distribution. View the status of scheduled distribution tasks. Users who can define content offers can also cancel or reschedule pending or failed tasks. <p>If the Content box is not checked, the Content link is not shown in the Web console.</p> <p>Note that a user who can define content offers and distribution groups can change or delete the content offers and distribution groups defined by any other user.</p>

Admin

Select the Admin check box to allow users to view their CMS account settings. You can also enable:

- **Manage User Groups.** Import communities, enable or disable device polling, view polling logs, and define user accounts, user groups, device groups, email distribution lists, and event forwarding filters. All functions are limited to the users and devices in the user groups that the user belongs to.

Users with this option are user group administrators, and must be authenticated locally.

- **Manage WXOS Boot Images.** Upload WXOS boot images to CMS.

If the Admin box is not checked, the Admin link is not shown in the Web console.

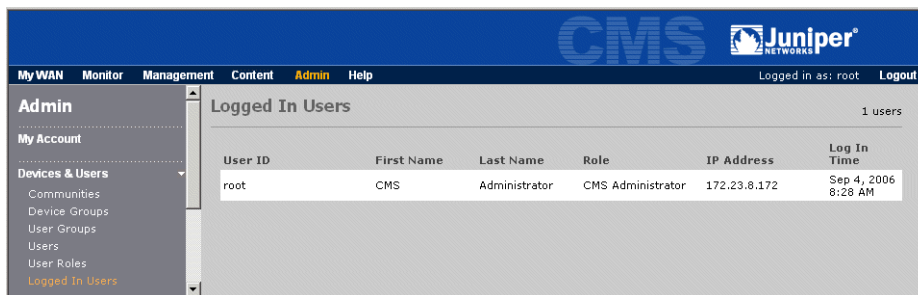
Viewing Logged In Users

CMS and user group administrators can view a list of the users who are currently logged in to CMS.

To view a list of the logged-in users:

1. Click Admin in the menu frame, and then click Logged In Users in the left-hand navigation frame.

Figure 209: Viewing Logged In Users



User ID	First Name	Last Name	Role	IP Address	Log In Time
root	CMS	Administrator	CMS Administrator	172.23.8.172	Sep 4, 2006 8:28 AM

2. Review the following information for each user:

- User ID and name.
- User role. For more information about user accounts and roles, refer to “Defining CMS User Accounts” on page 388.
- IP address of the client that the user logged in from.

- Date and time the user logged in.



NOTE: Users who close their Web browser without logging out are shown here until the session timeout expires.

Administering WX Devices

The following sections describe the device-related administration tasks:

- “Importing and Managing Communities” on page 397
- “Changing a Registration Server Address or Password” on page 401
- “Managing Device Groups” on page 402
- “Configuring Device Polling” on page 405
- “Configuring Data Retention” on page 407
- “Uploading a Boot Image” on page 408
- “Viewing the Polling Catch-Up and Failure Logs” on page 409

Importing and Managing Communities

A community is a group of WX devices that can compress and accelerate the traffic sent between them. Communities are defined on the devices that act as registration servers. When you install a WX device, you specify the IP address and password of a registration server, which is contacted periodically to identify the other devices in the same community.

To manage the devices in each community, CMS must import the communities defined on each registration server. Thereafter, the registration server is queried each day for changes to the imported communities. As new communities are added to a registration server, they must be imported into CMS. Only CMS administrators and user group administrators can import communities.

To import communities:

1. Click Admin in the menu frame, and then click Communities in the left-hand navigation frame.

Figure 210: Communities Page

The screenshot displays the Juniper CMS interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content', 'Admin', and 'Help'. The user is logged in as 'root'. The left sidebar shows the 'Admin' menu expanded, with 'Communities' selected. The main content area is titled 'Communities' and shows a list of 10 communities. Each row includes a checkbox, the community name, the registration server, the secondary registration server, and the number of devices. Some rows also have a 'NAT' icon. At the bottom of the list are 'Import...' and 'Delete' buttons.

Community Name	Registration Server	Secondary Registration Server	Devices
<input type="checkbox"/> backup-test	SR-10.87.240.2	SM-10.87.242.2	3
<input type="checkbox"/> cms-private	SR-192.168.101.100-NAT	192.168.243.2	3 NAT
<input type="checkbox"/> cms-test	SR-192.168.101.100-NAT	192.168.243.2	4 NAT
<input type="checkbox"/> default-10.87.207.11	SM-10.87.207.11		6
<input type="checkbox"/> default-10.87.229.10	SM-10.87.229.10		5
<input type="checkbox"/> default-10.87.235.2	SM-10.87.235.2		6
<input type="checkbox"/> default-10.87.240.2	SR-10.87.240.2	SM-10.87.242.2	7
<input type="checkbox"/> default-192.168.101.100	SR-192.168.101.100-NAT	192.168.243.2	3 NAT
<input type="checkbox"/> jnpr_test	SR-10.87.240.2	SM-10.87.242.2	6
<input type="checkbox"/> test	SR-10.87.240.2	SM-10.87.242.2	2

For each community, the Communities page lists the names of the primary and secondary registration servers and the number of devices in the community. From the Communities page, you can:

- Import communities, as described in Step 2.
- Change a registration server password or address, as described in “Changing a Registration Server Address or Password” on page 401.
- Change the public IP address for one or more WX devices that are on the private side of a NAT device. Click the number of devices next to the NAT icon, change the public addresses, and click Submit. You can also change device addresses from the Devices page.
- Delete communities from CMS. Select the check box next to one or more communities, and click Delete. Note the following:
 - All schedule information is also deleted.
 - If the devices in a deleted community belong to no other communities, they are removed from all device groups.



NOTE: To move all communities from one registration server to another, you can retain device group membership and all scheduled tasks by changing the server address rather than deleting and re-importing communities (refer to “Changing a Registration Server Address or Password” on page 401).

2. To import the communities from a registration server:
 - a. Click Import to open the Communities > Import page.

Figure 211: Importing Communities from a Registration Server

- b. Specify the following information, and click Next.

Public IP Address	Enter the IP address of the registration server.
Private IP Address	If the registration server is on the private side of a NAT device, enter the server's private IP address. If two registration servers have the same private address, their default communities have the same name. Change one address or move all devices out of the default community on both servers.
Password	Enter the registration server's password.

- c. Select the check box next to each community you want to import, and click Next. Communities that have already been imported do not have a check box.

- d. Select at least one user group that can access the imported devices, or click New User Group to define a new user group (refer to “Defining User Groups” on page 391), and then click Next.
- e. Click Finish. If the registration server is behind a NAT device, the Public IP Addresses page opens.

Figure 212: Entering Public IP Addresses

The screenshot shows the 'Public IP Addresses' page in the CMS. The page has a sidebar on the left with navigation links: Admin, My Account, Devices & Users, WAN Performance Thresholds, Setup, Events, and Diagnostics. The main content area has a title 'Public IP Addresses' and a description: 'The devices listed below are located behind a NAT router. In order to access the devices, you are required to enter their public IP addresses.' It also includes a warning: 'If any public IP address fields are highlighted in yellow below, the device reports a conflicting private IP address or registration server. Check that the public/private address mapping is correct.' and a note: 'You may not change the public IP address for a registration server here. Use the Registration Server page accessible from the Communities page instead.'

Showing 3 endpoints

Device Name	Private IP Address	Public IP Address
CMSQA-244	10.91.106.100	10.87.244.2
SR-192.168.101.100-NAT	192.168.101.100	10.87.241.2
CMSQA-243	192.168.243.2	10.87.243.2

Buttons: Submit, Cancel, Clear, Reset

- f. Enter the public IP address for each device listed, and click Submit. If the public address is missing or incorrect, you cannot select the device on the Devices page or Device Polling page (the check box is hidden). Duplicate or incorrect addresses are detected the first time the device is polled.

Changing a Registration Server Address or Password

If a registration server’s password or IP address(es) has changed, you must update CMS. After you change the password in CMS, you can download it to the devices in each community defined on the registration server (refer to “Applying a Registration Server Password” on page 65).

Only CMS administrators and user group administrators can update a registration server’s password or addresses.



NOTE: Changes made here affect only CMS, not the registration server. When you change a registration server’s IP address in CMS, all communities are migrated to the new address, and all scheduled tasks and device group memberships are preserved. You must transfer the registration server to the new address before you change CMS (refer to the *WX/WXC Operator’s Guide*).

To change a registration server’s password or IP address:

1. Click Admin in the menu frame, and then click Communities in the left-hand navigation frame.

Figure 213: Community Administration

Communities 10 communities

To import community information from a Registration Server, click on **Import**. To set the registration server password used by Juniper WX CMS to communicate with devices in a community, click on the registration server. To delete a community, select it using the checkbox and click **Delete**.

Community Name	Registration Server	Secondary Registration Server	Devices
<input type="checkbox"/> backup-test	SR-10.87.240.2	SM-10.87.242.2	3
<input type="checkbox"/> cms-private	SR-192.168.101.100-NAT	192.168.243.2	3 NAT
<input type="checkbox"/> cms-test	SR-192.168.101.100-NAT	192.168.243.2	4 NAT
<input type="checkbox"/> default-10.87.207.11	SM-10.87.207.11		6
<input type="checkbox"/> default-10.87.229.10	SM-10.87.229.10		5
<input type="checkbox"/> default-10.87.235.2	SM-10.87.235.2		6
<input type="checkbox"/> default-10.87.240.2	SR-10.87.240.2	SM-10.87.242.2	7
<input type="checkbox"/> default-192.168.101.100	SR-192.168.101.100-NAT	192.168.243.2	3 NAT
<input type="checkbox"/> jnpr_test	SR-10.87.240.2	SM-10.87.242.2	6
<input type="checkbox"/> test	SR-10.87.240.2	SM-10.87.242.2	2

Import... **Delete**

- Click the registration server name, and update the following information:.

Public IP Address	Enter the IP address of the registration server.
Private IP Address	If the registration server is on the private side of a NAT device, enter the server's private IP address.
Password Verify password	Enter the new password (must match the password defined on the registration server). If you make a mistake and click Submit, enter and verify the previous password first, and then enter the new password. Downloading the new password will fail if the previous password is incorrect.

- Click Submit to activate the changes, or click Reset to discard them.
- If you change the password, download the new password to the devices in all communities defined on the registration server (refer to “Applying a Registration Server Password” on page 65).

Managing Device Groups

Each CMS user can be granted read or write access to one or more communities or device groups. Device groups let you create arbitrary groups of devices that are independent of communities. A device group can include a single device, or multiple devices from one or more communities. The same device can appear in multiple device groups. Only CMS administrators and user group administrators can define device groups.

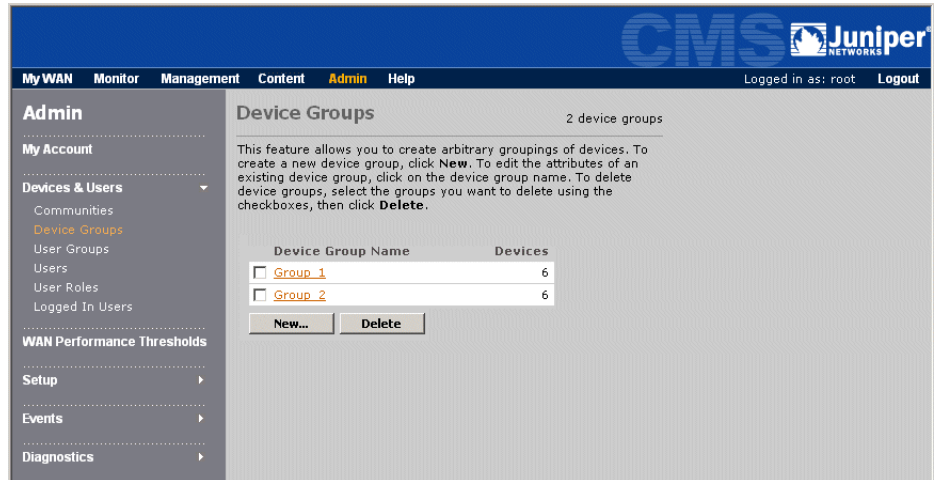


NOTE: When viewing reports by device group, performance statistics are relevant only for devices in the same community.

To define device groups:

1. Click Admin in the menu frame, and then click Device Groups in the left-hand navigation frame.

Figure 214: Device Groups



From the Device Groups page, you can:

- Add new device groups, as described in Step 2.
- Change a device group. Select a group name, click Add/Remove Devices to change the devices in the group, click Add/Remove User Groups to change the user groups who can access the in device group, and click Submit.
- Delete device groups. Select the check box next to one or more device groups, and click Delete. The deleted device groups are removed from all user groups, and users in those user groups may lose access to some or all of the devices in the deleted device groups.

2. To add a new device group:
 - a. Click **New** to open the Device Groups > New page.

The screenshot shows the 'Device Groups > New' page in the CMS Administrator's Guide. The page has a blue header with 'CMS Juniper NETWORKS' and a navigation bar with 'My WAN', 'Monitor', 'Management', 'Content', 'Admin', and 'Help'. The 'Admin' section is expanded, showing 'My Account', 'Devices & Users', 'WAN Performance Thresholds', 'Setup', 'Events', and 'Diagnostics'. The 'Devices & Users' section is further expanded, showing 'Communities', 'Device Groups', 'User Groups', 'Users', 'User Roles', and 'Logged In Users'. The 'Device Groups > New' page has a title bar 'Device Groups > New' and a step indicator 'Step 1: Select User Group(s)'. Below the title bar, there is a text input field for 'Device Group Name' and a list of 'User Groups' with checkboxes. The 'User Groups' list includes 'All Users', 'ad_cmsAdmin_grp', 'ad_devAdmin_grp', 'ad_devMonitor_grp', and 'ad_devOperator_grp'. A 'New User Group...' button is located below the list. At the bottom of the page, there are 'Back', 'Next', and 'Cancel' buttons.

- b. Enter a device group name (up to 32 characters).
 - c. Select at least one user group that can access the new device group. To add a new user group for yourself, click New User Group, enter a group name, and click OK. The new user group is added and selected.

By default, the selected user groups have read/write access to all devices in the group. To change the access to Read Only, or to view or change the members of a user group, refer to “Defining User Groups” on page 391.

Click Next.

- d. Select a community from the Community list. The device name and IP address are shown for each device in the selected community. The IP address is enclosed in parentheses. To import more communities, refer to “Importing and Managing Communities” on page 397.

- e. Select the devices you want to add to the device group, and click Add. To remove devices from the Members of Device Group list, select the devices and click Remove.
- f. Repeat Steps d and e for each community, as needed, and click Finish.

Configuring Device Polling

By default, performance data is collected every hour from WXOS 5.2 or later devices. CMS administrators and user group administrators can disable polling for specific devices.

If you have a large number of WX devices, a CMS administrator can reduce the polling interval to once a day to conserve system resources. The length of time that collected data is retained can also be changed (refer to “Configuring Data Retention” on page 407).

The following table provides a rough estimate of the daily growth of the database (actual results depend on the device configurations). Polling once a day eliminates the per minute data.

Type of Data	Daily Disk Space per Device	Max Days Retention
Per minute	10 MB	10
Hourly	180 KB	180
Daily	15 KB	365



NOTE: If you poll devices every hour, all devices should use an NTP server to ensure the accuracy of the hourly reports (refer to “Configuring NTP” on page 122). Also, statistics for the last hour cannot be viewed while hourly polling is in progress.

To configure device polling:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click Device Polling.

Figure 215: Configuring Device Polling

The screenshot shows the 'Device Polling' configuration page in the CMS. The sidebar on the left has 'Admin' selected, and 'Device Polling' is highlighted under the 'Setup' section. The main content area has a blue header with 'CMS' and 'Juniper' logos. Below the header is a navigation bar with 'My WAN', 'Monitor', 'Management', 'Content', 'Admin', and 'Help'. The 'Admin' section is expanded, showing 'My Account', 'Devices & Users', 'WAN Performance Thresholds', 'Setup', 'Events', and 'Diagnostics'. The 'Setup' section is expanded, showing 'AAA', 'CMS Access', 'Device Polling', 'Data Retention', 'FTP Server', 'License Key', 'Scheduler', 'Session Timeout', 'Web Server Port', and 'WX Boot Images'. The 'Device Polling' section is selected, showing a description: 'Checked endpoints will be polled at the specified interval. To configure the length of time that collected data is retained, go to the Data Retention page.' Below the description are three dropdown menus: 'Polling Interval' (set to '1 hour'), 'Polling Start Time' (set to '1:00 AM'), and 'Community/Device Group' (set to '-- All Devices --'). Below these is a table of endpoints with checkboxes for selection. The table has columns for 'Endpoints' and 'IP Address'. The endpoints listed are: CMSQA-10.87.244.107 (checked), CMSQA-243 (checked), SR-10.87.240.2 (checked), WX-10.87.244.2 (unchecked), WX-10.87.246.2 (checked), WX-10.87.248.2 (checked), and WX-10.87.249.2 (checked). At the bottom of the table are 'Select All' and 'Clear' buttons. A 'Submit' button is at the bottom of the page.

2. Specify the following information and click Submit:

Poll Checked Endpoints Select the check box next to each device that you want to poll. To select all devices, click Select All. To deselect all devices, select Clear. To view the devices in a single community or device group, select the community or device group from the Community/Device Group menu.

Polling Interval Select a polling interval:

- **1 hour.** Collects data for the previous hour (the default).
If a device does not respond to a poll, the next poll requests hourly data (rather than per-minute data) for the previous two hours. A “catch-up” poll can request up to the last 24 hours of data. To view the catch-up polls for each device, refer to “Viewing the Polling Catch-Up and Failure Logs” on page 409.
- **1 day.** Collects data for the previous day. Per minute data is not retained, and the Last Hour and Today time periods will be greyed out on reports.
- **Never.** Disables polling for all devices.

Polling Start Time If you select a one-day polling interval, select a start time (off-peak hours are recommended)

Configuring Data Retention

A CMS administrator can change the length of time that performance and other data is retained in the database. The growth of the database depends on the number of WX devices and the device polling interval (refer to “Configuring Device Polling” on page 405).

To configure data retention:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click Data Retention.

Figure 216: Configuring Data Retention

The screenshot shows the CMS web interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content', 'Admin' (selected), and 'Help'. The left-hand navigation menu has 'Admin' selected, with sub-items 'My Account', 'Devices & Users', 'WAN Performance Thresholds', and 'Setup'. Under 'Setup', 'Data Retention' is highlighted. The main content area is titled 'Data Retention' and contains the following configuration options:

Performance Data	
Minute Data	10 days (max: 10)
Hourly Data	30 days (max: 180)
Daily Data	180 days (max: 365)

Other Data	
Event Data and Content Schedule Log	16 days (max: 180)

Only CMS administrators may change the above settings.

Submit

2. Specify the following information and click Submit:

Performance Data Enter the number of days to retain the collected data in the database. The options are:

- **Minute data.** Up to 10 days (default is 7).
- **Hourly data.** Up to 180 days (default is 30).
- **Daily data.** Up to 365 days (default is 180).

Other Data Enter the number of days (up to 180) to retain event data and the schedule log for content distribution (default is 180).

Uploading a Boot Image

To load software upgrades on selected devices, the boot image must be uploaded to CMS from a local disk or an FTP server. To upload a boot image, users must have the Upload WX Boot Images privilege (refer to “Defining User Roles” on page 393).

To download a boot image from the CMS server to selected devices in a community, refer to “Loading Device Boot Images” on page 45. .

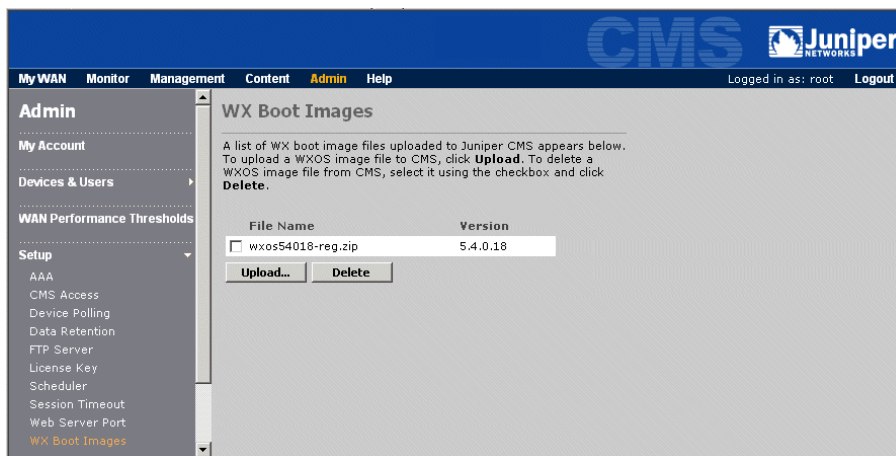


NOTE: You should retain the default name of the boot image to easily identify different releases and builds.

To upload a boot image to the CMS server:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click WX Boot Images.

Figure 217: Viewing Uploaded Boot Images



2. To remove an outdated image, select the check box next to the image and click Delete.
3. To upload a new image:
 - a. Click Upload.

Figure 218: Uploading a Boot Image

The screenshot shows the Juniper CMS web interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content', 'Admin' (highlighted), and 'Help'. The user is logged in as 'root'. The left sidebar shows the 'Admin' menu with options like 'My Account', 'Devices & Users', 'WAN Performance Thresholds', and 'Setup'. The 'Setup' menu is expanded, showing 'AAA', 'CMS Access', 'Device Polling', 'Data Retention', 'FTP Server', 'License Key', 'Scheduler', 'Session Timeout', 'Web Server Port', and 'WX Boot Images'. The main content area is titled 'Upload WX Boot Image' and contains the following form:

Upload WXOS image file to Juniper WX CMS from the following location.

☒ Local Disk File path/name:

☐ FTP Server IP Address:

File path/name:

User Name:

Password:

- b. Select Local Disk and click Browse to locate the boot image, or select FTP Server and specify the IP address of the FTP server, the pathname and filename of the boot image, and the user name and password. If the FTP server accepts anonymous user access, leave the user name and password blank.

The boot image must have either a “.bin” or “.zip” extension. CMS does not recognize files with other extensions.

- c. Click Submit to upload the boot image to the CMS server.

Viewing the Polling Catch-Up and Failure Logs

If an WXOS 5.2 (or later) device does not respond to an hourly poll for performance data, the next poll requests hourly data (rather than per-minute data) for the previous two hours. These catch-up polls can request up to the last 24 hours of data. If a device does not respond for an entire day, the performance data for that day will be lost.

Note that successful catch-up polls create discrepancies between the Last Hour and daily reports for a device. For example, the Today report will have data for the last hour, but the Last Hour report will be blank (no per-minute data). CMS administrators and user group administrators can verify the cause of these discrepancies by viewing the polling catch-up and failure logs.

To view the polling catch-up or failure log:

1. Click Admin in the menu frame, click Diagnostics in the left-hand navigation frame, and then click Polling Catch-Up Log or Polling Failure Log.

The polling catch-up log (Figure 219) shows the time of each catch-up poll for each device.

Figure 219: Viewing the Polling Catch-Up Log

CMS Juniper

My WAN Monitor Management Content **Admin** Help

Logged in as: root Logout

Admin

- My Account
- Devices & Users
- WAN Performance Thresholds
- Setup
- Events
- Diagnostics**
 - Polling Catch-Up Log**
 - Polling Failure Log
 - Show CMS Log
 - CMS Diagnostic File

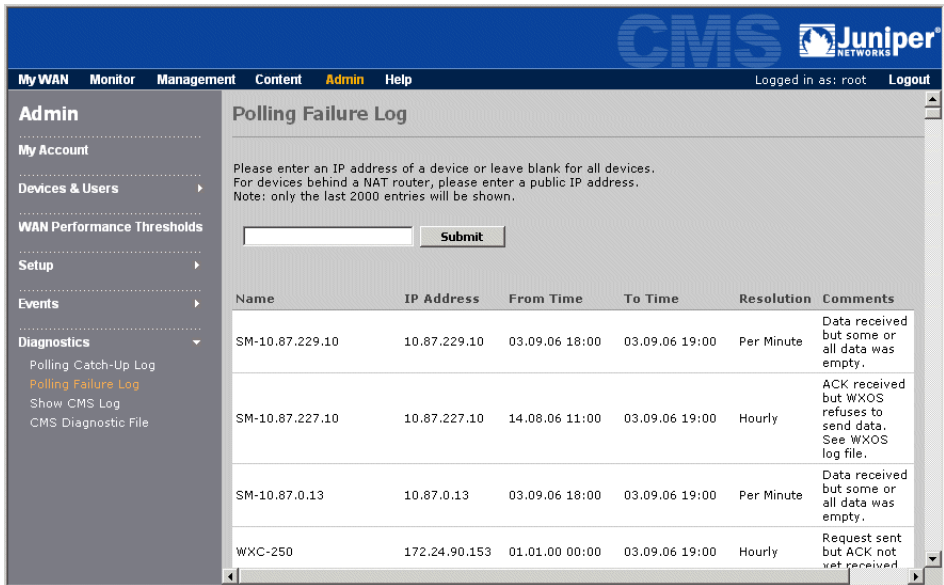
Polling Catch-Up Log

Please enter an IP address of a device or leave blank for all devices.
For devices behind a NAT router, please enter a public IP address.
Note: only the last 2000 entries will be shown.

Name	IP Address	From Time	To Time	Catch-up Type
SM-10.87.227.10	10.87.227.10	14.08.06 11:00	03.09.06 20:00	Hourly instead of per minute
WXC-250	172.24.90.153	--	03.09.06 20:00	Hourly instead of per minute
SM-10.87.227.10	10.87.227.10	14.08.06 11:00	03.09.06 19:00	Hourly instead of per minute
WXC-250	172.24.90.153	--	03.09.06 19:00	Hourly instead of per minute
SM-10.87.227.10	10.87.227.10	14.08.06 11:00	03.09.06 18:00	Hourly instead of per minute
WXC-250	172.24.90.153	--	03.09.06 18:00	Hourly instead of per minute
SM-10.87.227.10	10.87.227.10	14.08.06 11:00	03.09.06 17:00	Hourly instead of per minute
WXC-250	172.24.90.153	--	03.09.06 17:00	Hourly instead of per minute
SM-10.87.227.10	10.87.227.10	14.08.06 11:00	03.09.06 16:00	Hourly instead of per minute

The polling failure log (Figure 220) shows the time of each failed poll for each device, whether the poll was for per minute or hourly data (a normal or catch-up poll), and whether the device failed to respond or responded with no statistics.

Figure 220: Viewing the Polling Failure Log



- To view the log entries for a specific device, enter the IP address and click Submit. To view all entries again, delete the IP address and click Submit. The logs retain the last 2000 entries.

Managing Event Forwarding Filters and Email Distribution Lists

The following sections describe how to define event forwarding filters and email distribution lists for events and reports. These features are available only to users that have the CMS Administrator or User Group Administrator role:

- “Defining Email Distribution Lists” on page 412
- “Defining Event Forwarding Filters” on page 413

Defining Email Distribution Lists

CMS can send alerts to a list of email addresses for selected CMS and WX events. The subject and body of an email alert are defined as follows:

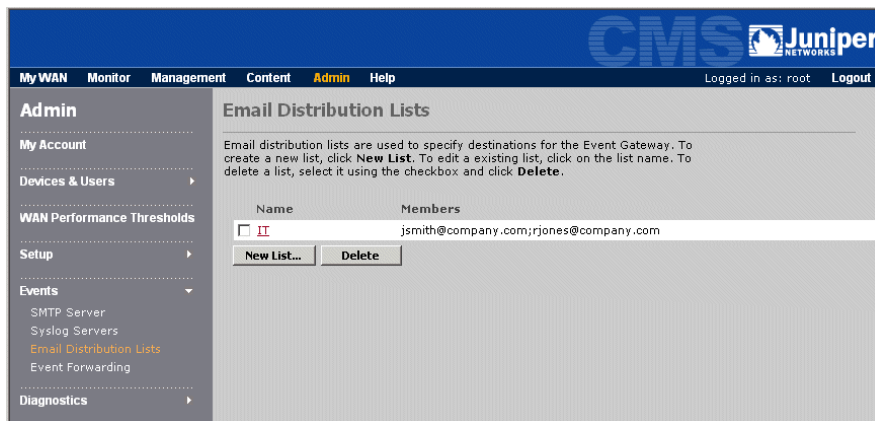
- **Subject.** “Juniper WX Event: < device_name > < metric name > ” or “Juniper WX CMS Event: < metric name > ”
- **Body.** One or more lines with “ < field > : < value > ”

To define filters that specify which events trigger email alerts, refer to “Defining Event Forwarding Filters” on page 413.

To define email distribution lists:

1. Click Admin in the menu frame, click Events in the left-hand navigation frame, and then click Email Distribution Lists.

Figure 221: Viewing Email Distribution Lists



CMS administrators can view all distribution lists, but user group administrators can view only the lists they created. From the Email Distribution Lists page, you can:

- Add email distribution lists, as described in Step 2.
- Change a distribution list. Click the list name, make any needed changes, and click Submit.
- Delete distribution lists. Select the check box next to one or more list names, and click Delete. You cannot delete a list that is used in an event forwarding filter. The lists created by a user group administrator are deleted automatically if the user's account is deleted.

2. To add a new email distribution list:
 - a. Select New List and specify the following information:

List Name	Enter the name of the list (up to 32 characters).
Email Addresses	Enter the email addresses (one per line) where you want to send event alerts.
 - b. Click Submit to activate the changes, or click Cancel to discard them.

Defining Event Forwarding Filters

CMS events and WX system and performance events can be selectively forwarded to Syslog servers, and/or to an SMTP server where email alerts are generated based on a specified distribution list. You can define any number of filters, and each filter can specify one Syslog server and one distribution list.

Note the following:

- To forward events to a Syslog server, verify that Syslog reporting is enabled in CMS (refer to “Enabling Syslog Reporting” on page 429).
- To generate email alerts:
 - Define at least one distribution list (refer to “Defining Email Distribution Lists” on page 412).
 - Verify that an SMTP server is defined and event email forwarding is enabled (refer to “Defining an SMTP Server” on page 427).
- For CMS to receive events, devices running WXOS 5.4 or later must identify the CMS server as a Syslog server (refer to “Defining Syslog Servers” on page 124). CMS discards events from devices that are not in an imported community.
- To generate performance events, WXOS 5.4 devices must specify thresholds for the appropriate metrics. You can also enable or disable specific system events (refer to “Configuring Events” on page 280).

To view the Events reports, refer to “Events Reports” on page 362. For a description of CMS and WX system events, refer to “System Events” on page 441.

To define event forwarding filters:

1. Click Admin in the menu frame, click Events in the left-hand navigation frame, and then click Event Forwarding.

Figure 222: Viewing Event Forwarding Filters



CMS administrators can view all event forwarding filters, but user group administrators can view only the filters they created. From the Event Forwarding Filters page, you can:

- Add event forwarding filters, as described in Step 2.
- Enable or disable a filter. Set or clear the check box next to the filter, and click Submit.
- Change a filter. Click EDIT next to the filter, make any needed changes, and click Submit.
- Delete a filter by clicking DELETE next to the filter. The filters created by a user group administrator are deleted automatically if their user's account is deleted.

- 2. To add an event forwarding filter:
 - a. Click New Filter.

Figure 223: Adding Event Forwarding Filters

The screenshot shows the Juniper CMS web interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content', 'Admin' (highlighted), and 'Help'. The user is logged in as 'root'. The left sidebar shows the 'Admin' menu with options like 'My Account', 'Devices & Users', 'WAN Performance Thresholds', 'Setup', 'Events' (with sub-items: SMTP Server, Syslog Servers, Email Distribution Lists, Event Forwarding), and 'Diagnostics'. The main content area is titled 'Event Forwarding Filter > Edit'. It contains a note about alert generation and a warning about the 'Any' selection for Community/Device Group. The configuration fields include: 'Event Severity' with checkboxes for OK, Warning, Major, and Critical; 'Event Type' and 'Event Metric' as dropdown menus (both set to 'Any'); 'Community/Device Group', 'Source', and 'Destination' as dropdown menus (all set to 'Any'); 'Created By' set to 'root'; 'Email Distribution List' as a dropdown menu (set to '-- None selected --'); and 'Syslog Server' with radio buttons for 'Select from list:' (set to '-- None selected --') and 'Enter IP Address:'. 'Submit' and 'Cancel' buttons are at the bottom.

- b. Specify the filter criteria. Only events that match all the criteria are forwarded.

Severity	<p>Select one or more severity levels of the events to be forwarded. The corresponding severity level used on Syslog events is shown in parentheses.</p> <ul style="list-style-type: none">■ OK (Notice)■ Warning (Information)■ Major (Error)■ Critical (Critical) <p>The severity levels of CMS and WX system events are shown in “System Events” on page 441 . Severity levels of WX performance events are user defined.</p>
Event Type	<p>To forward just one type of event, select the event type (WX CMS, WX Performance, or WX System). The default is any event type. To configure system and performance events on WXOS 5.4 devices, refer to “Configuring Events” on page 280</p>

Event Metric	<p>If you select an event type, you can select a single event metric to be forwarded. The default is any metric. For example, if you select the “WX CMS” event type, you can forward just the “Device Task Failed” events. CMS and WX system event metrics are described in “System Events” on page 441. WX performance metrics are described in “Configuring Events” on page 280.</p>
Community/Device Group	<p>To forward events from the devices in one community or device group, select the community or device group. The default is any device accessible to the user who creates the filter. If the selected community/device group is later deleted, “INVALID” is displayed in this field.</p>
Source	<p>If you select a community/device group, you can select a single device to forward events from. The default is any device in the selected community/device group. For CMS “Device Task Failed” events, the source device is the device where the task failed.</p> <p>Note that the public IP address of the device is used as the match criteria.</p>
Destination	<p>If you select a source device, you can select a single destination so that WX performance events from the source are forwarded only for performance measured to the specified destination. The default is any destination in the selected community/device group, including devices that are inaccessible to the user who creates the filter.</p> <p>Alternatively, you can select:</p> <ul style="list-style-type: none"> ■ All (Aggregated). Only events based on the overall performance of all remote endpoints are forwarded. ■ Non-WX Endpoint. Select a non-WX endpoint, such as “Other Traffic” to forward events only for the WAN Throughput, QoS Throughput, Bytes Dropped, and Packets Dropped WX performance metrics. Note that the “Other Traffic” endpoint includes only the traffic that is not sent to WX endpoints or other customized non-WX endpoints. <p>All non-WX endpoints are listed, not just those associated with the selected community/device group. Note that non-WX endpoints are not listed until CMS receives some events or statistics for the endpoint.</p> <p>Note that a destination device can refer to the secondary IP address; the source device always uses the primary IP address.</p>
Specify at least one forwarding method:	
Email Distribution List	<p>Select an email distribution list to forward an email alert to each address on the list when an event occurs that matches the filter criteria. To add email distribution lists, refer to “Defining Email Distribution Lists” on page 412.</p>
Syslog Server	<p>To forward events to a Syslog server, select a server from the list or enter a server IP address. To enable Syslog reporting and/or add Syslog servers to the list, refer to “Enabling Syslog Reporting” on page 429.</p>

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

Administering CMS

The following sections describe the features available only to users that have the CMS Administrator role:

- “Setting WAN Performance Thresholds” on page 418
- “Entering a Permanent License Key” on page 419
- “Controlling Client Device Access to CMS” on page 420
- “Configuring AAA Settings for Remote Authentication” on page 421
- “Defining the Session Timeout” on page 425
- “Defining an FTP Server” on page 426
- “Defining an SMTP Server” on page 427
- “Enabling Syslog Reporting” on page 429
- “Stopping and Starting the Scheduler” on page 430
- “Changing the CMS Server IP Address” on page 430
- “Changing the Web Server Port” on page 431
- “Viewing System Logs” on page 432
- “Generating a Diagnostic File” on page 433
- “Backing Up and Restoring the Database” on page 434
- “Moving CMS to Another Disk Drive” on page 437
- “Purging Temporary Java Files” on page 438
- “Changing the CMS Time Zone” on page 438

Setting WAN Performance Thresholds

The WAN Performance, Loss, Latency, and Availability charts use colored cells in a matrix to indicate the WAN status between each pair of devices in a community (refer to “WAN Statistics” on page 311 and “Executive Summary” on page 358). As needed, you can change the default performance ranges associated with each of the colors.

To change the WAN reporting thresholds:

1. Click Admin in the menu frame, and then click WAN Performance Thresholds in the left-hand navigation frame.

Figure 224: Setting WAN Reporting Thresholds

WAN Performance Thresholds

The table below allows you to modify the thresholds which determine the display of colored dots in the WAN performance reports.

	% Time Above Latency Threshold	% Loss	% Availability
Excellent	< 0.5	< 0.25	> 99.5
Good	0.5 to 1.0	0.25 to 0.5	99.0 to 99.5
Marginal	1.0 to 5.0	0.5 to 1.0	98.0 to 99.0
Warning	5.0 to 10.0	1.0 to 2.0	95.0 to 98.0
Critical	> 10.0	> 2.0	< 95.0

The WAN Performance report provides an aggregate health metric that combines latency, loss and availability. You can select the manner in which the aggregate measure is determined from the three parameters for each path in the WAN.

☐ Display the best of the parameters
☐ Display the worst of the parameters
☒ Display the average of the parameters

2. To change the performance ranges for a color, enter the new percentage values for time above the latency threshold, loss, and availability.
3. Select whether the WAN Performance report reflects the best, average, or worst values measured for loss, latency, and availability.
4. Click Submit to activate the changes, or click Reset to discard them. To restore the default values, click Set to Defaults.

Entering a Permanent License Key

CMS requires a permanent license key to operate beyond the 45-day evaluation period. The permanent license key determines the maximum number of devices that CMS manages. To obtain a permanent license, you need an authorization code and the IP address of the CMS server (refer to “CMS Licenses” on page 439).



NOTE: If you change the IP address of the CMS server, you must reboot the server and obtain a new license key before you can continue using CMS. To avoid this, install a Windows loopback adapter, and use the loopback IP address to obtain a CMS license. To install a loopback adapter, refer to <http://support.microsoft.com/kb/839013> or search support.microsoft.com.

To enter a permanent license key for CMS:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click License Key.

Figure 225: Entering a License Key

The screenshot shows the CMS web interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content', 'Admin' (highlighted), and 'Help'. The right side of the top bar shows 'CMS' and 'Juniper NETWORKS' logos, along with 'Logged in as: root' and a 'Logout' link. The left-hand navigation menu is expanded to 'Admin', which includes 'My Account', 'Devices & Users', 'WAN Performance Thresholds', 'Setup' (expanded), 'Events', and 'Diagnostics'. Under 'Setup', the following options are listed: AAA, CMS Access, Device Polling, Data Retention, FTP Server, License Key (selected), Scheduler, Session Timeout, Web Server Port, and WX Boot Images. The main content area is titled 'License Key'. It contains the following text: 'The maximum number of WX devices supported by Juniper WX CMS is determined by the license key. WX CMS will create an evaluation license on installation. However, the evaluation license will expire in 45 days from installation. A valid license key must be entered prior to the expiration date in order to ensure uninterrupted service.' Below this text is a table with the following information:

License key	GQZE-G6YK-NYBC-3KDI-JIAA-GA
Current license	1000 devices
Expires	Never

Below the table, it states: 'The license can be upgraded by entering a new license key below.' There is a text input field labeled 'License Key'. At the bottom of the form are two buttons: 'Submit' and 'Reset'.

2. Enter the permanent license key in the License Key field. Be sure to enter all characters, including hyphens (-), of the permanent license key.
3. Click Submit to activate the permanent license key. To restore the original license key, click Reset.

Controlling Client Device Access to CMS

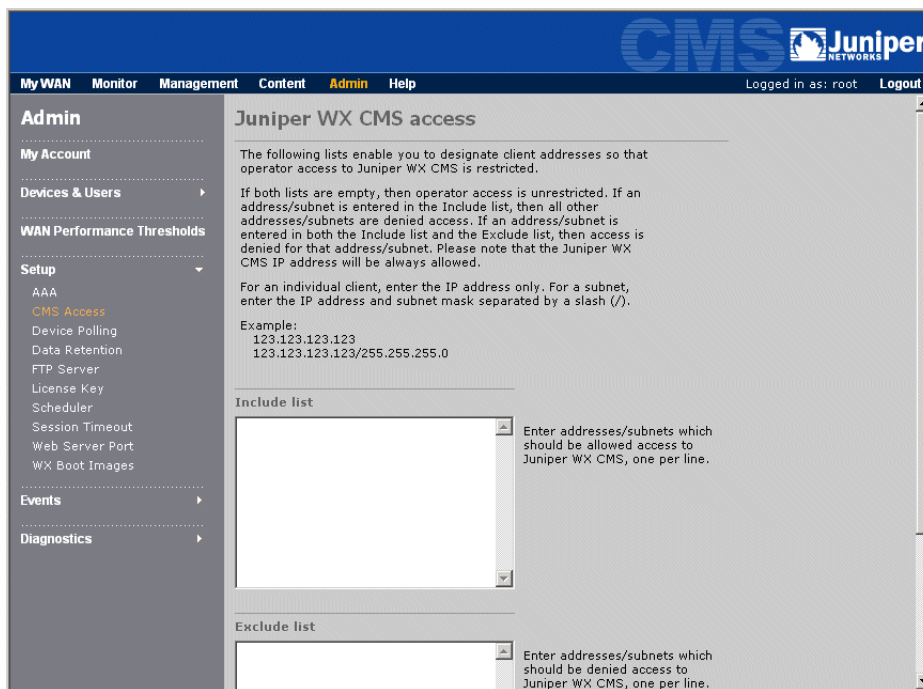
You can create an Include or Exclude list to allow or deny administrative access to CMS from specific IP addresses or subnets. For example, if you enter one address in the Include list, administrative users can log in only from the specified address. Alternatively, if you enter an address or subnet in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

By default, the Include and Exclude lists are empty, which means that administrative access is allowed from any address.

To restrict administrative access to CMS:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click CMS Access.

Figure 226: Controlling Client Device Access to CMS



2. To allow access to CMS only from specific IP addresses or subnets, enter the addresses or subnets in the Include list (one per line). The subnet format is:

<IP address>/<subnet mask>

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

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



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

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

Configuring AAA Settings for Remote Authentication

CMS administrators and user group administrators must be authenticated locally by CMS, but other users can be authenticated and authorized remotely by an Active Directory/OpenLDAP server.

To authenticate users remotely:


- Define the remote authentication server, as described here.
- Enable remote authentication for the appropriate user accounts (refer to “Defining CMS User Accounts” on page 388).
- Configure the authentication server to return the appropriate information. To provide authentication only, the remote server should return the following:

JuniperLocalUserName	<p>Name that matches the ID of a CMS user account, which specifies the user's role and associated user groups. Multiple users can be mapped to the same CMS account. The login name entered by the user is shown in the menu bar, while the first and last name in the CMS user account are shown on the My WAN page.</p> <p>If the returned name does not match a user in the CMS database, the user will see an error message indicating the account has not been configured properly. The log file will contain a more detailed error message.</p> <p>If this name is not returned, the login name must match a CMS user ID. Note that this name must be returned if the login name includes special characters that are disallowed by CMS.</p>
----------------------	--

JuniperDeviceAccessUrl	<p>Specifies one of the following keywords to indicate the method used to generate URLs for the device name links shown on the Devices page. Clicking a link opens the WXOS Web console for the device.</p> <ul style="list-style-type: none">■ usePrivateIpAddress. If the devices are on the private side of a NAT device, the private IP address is used; otherwise the public IP address is used (the default).■ usePublicIpAddress. The public IP address is used for all devices. Use this option if the devices are on the private side of a NAT device, but the user's workstation is not in the private address space of the WX devices.■ useDomainName. The URL is composed of the domain name and device name specified on the Addresses page of the WXOS Web console. The URL format is: https://<domain_name>/<device_name> The domain and device names can be downloaded from CMS in a Device Settings partial configuration (refer to "Configuring Device Addresses" on page 103). If a device does not report a domain name, only the device name appears in the URL. If both the domain and device names are not reported, the public IP address is used. The public IP address also is used if the reported device name contains periods. <p>If this attribute is not returned, the private IP address is used.</p>
------------------------	---

- To provide both authentication and authorization, the remote server must return the following:

JuniperUserRole	Name of a CMS user role. If the role is "CMS Administrator" or the role has the Manage User Groups activity enabled, the login is rejected.
JuniperUserGroups	Name of one or more CMS user groups. Multiple groups must be separated by commas.

 **NOTE:** When a device task is scheduled, the user's write access to the specified devices is verified twice—when the task is scheduled, and when the task is run. For a remotely authorized user, write access cannot be verified at run time.

To define the remote authentication server:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click AAA.

Figure 227: Defining a Remote Authentication Server

Admin

My Account

Devices & Users

WAN Performance Thresholds

Setup

AAA

CMS Access

Device Polling

Data Retention

FTP Server

License Key

Scheduler

Session Timeout

Web Server Port

WX Boot Images

Events

Diagnostics

AAA

This page allows you to choose between local and remote (Active Directory/LDAP) authentication and authorization methods. If local authentication is chosen below, then remote authentication and authorization are disabled; any user accounts in the local CMS database marked for remote authentication are effectively disabled. If 'Local and Active Directory' is chosen for authentication, then users in the local CMS database marked for Active Directory/LDAP authentication will be authenticated remotely.

If 'Local' is chosen for authorization below, then the role and associated user groups for all users are retrieved from the user records in the CMS database. If 'Follows Authentication' is chosen, then an entry for a user in the CMS database is not necessary. The Active Directory/LDAP server must be configured to provide the role and associated user groups for remotely authenticated users. Click the 'Help' button for further details.

Authentication: Local

Authorization: Local

Active Directory Server

Fully Qualified Domain Name:

IP Address:

Authentication Port: 389

Use SSL: ☐

Default port for SSL is 636. Default port for non-SSL is 389.

Base Distinguished Name:

Path in Active Directory from which to start searching for users.
Example: dc=companyname,dc=com

Submit Reset Import Certificate...

2. Specify the following information:

Authentication

Select the authentication method:

- **Local.** All users are authenticated locally (the default).
- **Local and Active Directory.** Non-administrative users whose accounts have remote authentication enabled are authenticated by the specified server. Remote authentication is also attempted if the login name does not match a CMS user ID.

Authorization

Select the authorization method:

- **Local.** All users are authorized locally (the default).
- **Follows Authentication.** For non-administrative users who are authenticated remotely, the remote server must be configured to return the user role and user groups.

Fully Qualified Domain Name	Enter the fully qualified domain name of the authentication server, such as "juniper.net".
IP Address	Enter the IP address of the authentication server.
Authentication Port	Enter the port number used by the authentication server. By default, an Active Directory server uses port 636 for Secure Socket Layer (SSL) connections, and port 389 for non-SSL connections.
Connection	Select the check box to use SSL to connect to the server.
Base Distinguished Name	<p>Enter the path on the server where the search for users begins. This name is entered automatically as you enter the domain name, but can be modified as needed.</p> <p>For example, if all users are under "cmslab.com/Users", the base DN could be:</p> <p>cn=Users,dc=cmslab,dc=com</p> <p>or</p> <p>dc=cmslab,dc=com</p>

3. Click Submit to activate the new settings. To restore the original settings, click Reset.
4. If you enable SSL connections, do the following:
 - a. Generate a certificate file on the authentication server, and copy the file to the local disk or an FTP server.
 - b. Click Import Certificate, specify the location of the certificate file, and click Submit.
 - c. Reboot the CMS server.

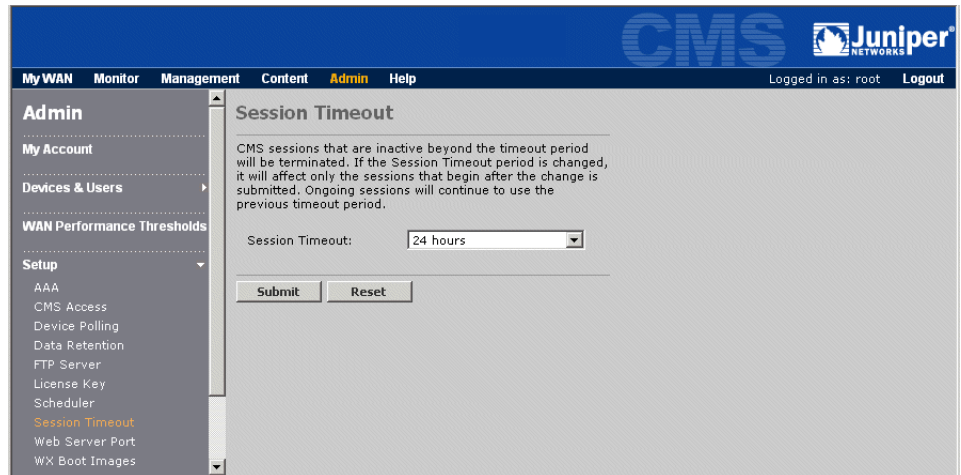
Defining the Session Timeout

The session timeout is the length of time a session can be idle before the session is closed (from 15 minutes to 24 hours). The default is 30 minutes.

To change the session timeout:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click Session Timeout.

Figure 228: Defining the Session Timeout



2. Select a timeout, and click Submit.

The new timeout affects only future sessions, not current sessions.

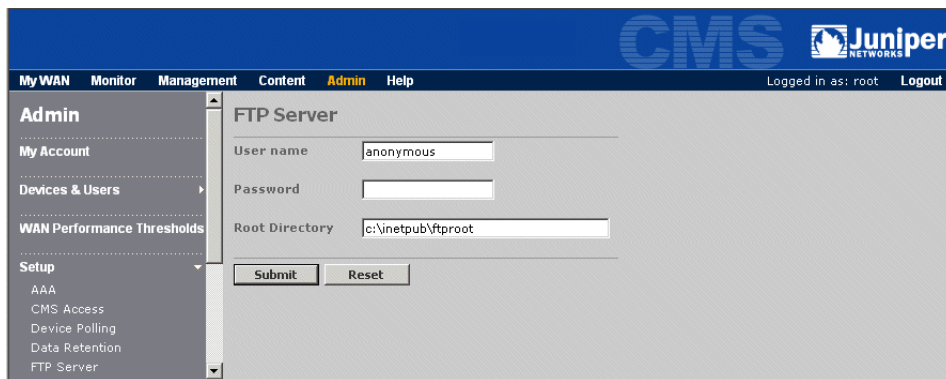
Defining an FTP Server

The Microsoft FTP server must be installed and running on the CMS server. During Quick Setup, you specified the FTP user name, password, and root directory. FTP is used by CMS to upload boot images, and by the devices to send data to CMS. You can change the FTP parameters at any time.

To define an FTP server:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click FTP Server.

Figure 229: Defining an FTP Server



The screenshot shows the CMS (Central Management System) web interface. At the top, there is a blue header with the 'CMS' logo and the 'Juniper NETWORKS' logo. Below the header is a navigation bar with tabs: 'My WAN', 'Monitor', 'Management', 'Content', 'Admin' (highlighted), and 'Help'. On the right side of the navigation bar, it says 'Logged in as: root' and 'Logout'. On the left side, there is a vertical menu under the 'Admin' tab with sections: 'My Account', 'Devices & Users', 'WAN Performance Thresholds', and 'Setup'. The 'Setup' section is expanded, showing sub-items: 'AAA', 'CMS Access', 'Device Polling', 'Data Retention', and 'FTP Server' (which is selected). The main content area is titled 'FTP Server' and contains three text input fields: 'User name' with the value 'anonymous', 'Password' (empty), and 'Root Directory' with the value 'c:\inetpub\ftproot'. Below these fields are two buttons: 'Submit' and 'Reset'.

2. Specify the FTP user name and password. If the FTP server allows anonymous user access, enter “anonymous” in the User name field and leave the Password field blank.
3. Verify that the FTP root directory is correct.
4. Click Submit to activate the new settings. To restore the original settings, click Reset.

Defining an SMTP Server

To email CMS reports as PDF files, or to generate email alerts for specific types of events, you must identify the local Simple Mail Transfer Protocol (SMTP) server that receives the emailed reports and alerts. You can change the SMTP parameters at any time.

If the SMTP server becomes unreachable, CMS attempts to reconnect every two minutes. If the internal queue of events to be sent via SMTP becomes full, subsequent events will not be forwarded. CMS attempts to send the queued events until successful or until SMTP event forwarding is disabled (see below).

To define an SMTP server:

1. Click Admin in the menu frame, click Events in the left-hand navigation frame, and then click SMTP Server.

Figure 230: Defining an SMTP Server

Admin

My Account

Devices & Users

WAN Performance Thresholds

Setup

Events

SMTP Server

Syslog Servers

Email Distribution Lists

Event Forwarding

Diagnostics

SMTP Server

In order to email CMS reports or to send email alerts from the Event Gateway, you must configure CMS to communicate with a local SMTP server.

SMTP Server Enter hostname or IP address

From Address Email address displayed in "From" field of email.

Display Name Name displayed in "From" field of email.

☐ Use SMTP Authentication

SMTP Username Enter any valid user account on the SMTP server

Password

Verify Password

☒ Enable Event Email Forwarding

2. Specify the following information:

Server	Enter the IP address or host name of a local SMTP server, such as "mail.juniper.net".
From Address	Enter the email address shown in the From field on emailed reports (such as "jsmith@juniper.com"). If you enable SMTP authentication, the address must match an account on the mail server; otherwise, any address can be used.
Display Name	Enter the name shown in the From field of emailed reports.

Use SMTP Authentication	<p>If the SMTP server requires authentication, select the check box and specify the following:</p> <ul style="list-style-type: none">■ SMTP Username. Enter any valid user account on the server.■ Password. Enter and verify the password defined on the server for the specified user. <p>SMTP servers typically require authentication to send email to external addresses, depending on the server configuration.</p>
Enable Event Email Forwarding	<p>Select the check box to enable the generation of email alerts. You can clear the check box to disable email alerts temporarily, as needed. When you disable event forwarding, any queued events waiting to be sent will be dropped.</p>

3. Click Submit to activate the new settings. To restore the original settings, click Reset.

Enabling Syslog Reporting

CMS can forward specific events to up to five Syslog servers. A Syslog server allows you to centrally log and analyze configuration events and system error messages, such as the failure of scheduled tasks. To define filters that specify the events to be forwarded to the Syslog servers, refer to “Defining Event Forwarding Filters” on page 413. For a description of Syslog messages, refer to “System Events” on page 441.

To enable Syslog reporting for CMS:

1. Click Admin in the menu frame, click Events in the left-hand navigation frame, and then click Syslog Servers.

Figure 231: Enabling CMS Syslog Reporting

The screenshot shows the CMS web interface. The top navigation bar includes 'My WAN', 'Monitor', 'Management', 'Content', 'Admin' (highlighted), and 'Help'. The left-hand navigation menu under 'Admin' includes 'My Account', 'Devices & Users', 'WAN Performance Thresholds', 'Setup', 'Events' (expanded), and 'Diagnostics'. Under 'Events', 'Syslog Servers' is selected. The main content area is titled 'Syslog Servers'. It contains a descriptive paragraph about forwarding CMS and WXOS events to a syslog server. Below this is a checkbox 'Enable Syslog Reporting' which is checked. Underneath the checkbox is a text area labeled 'Syslog Servers' with the IP address '10.91.5.50' entered. To the right of the text area is the instruction 'Enter IP addresses, one per line.' At the bottom of the form are 'Submit' and 'Reset' buttons.

2. Select the Enable syslog reporting check box to enable Syslog reporting, and then enter the IP addresses of up to five Syslog servers (one per line).
3. Click Submit to activate the changes, or click Reset to discard them.

Stopping and Starting the Scheduler

The CMS scheduler lets you schedule device management tasks for a future date and time (refer to “Managing CMS Schedules” on page 69). If your network is having problems, and you have critical tasks scheduled for execution, you might want to stop the scheduler until the problems are resolved.

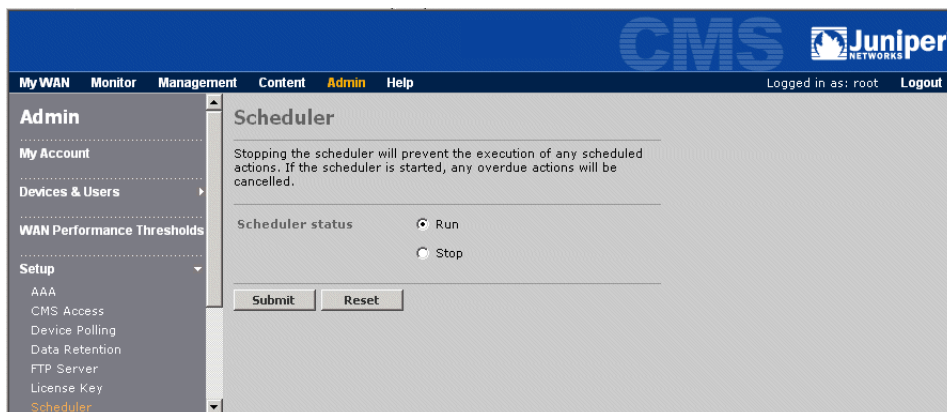


NOTE: You must reschedule any tasks that are scheduled to run while the scheduler is off.

To stop and start the scheduler:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click Scheduler.

Figure 232: Stopping and Starting the Scheduler



2. Click Stop or Run to stop or start the scheduler., and click Submit.

Changing the CMS Server IP Address

If you change the IP address of the CMS server, you must reboot the server and obtain a new license key before you can continue using CMS (refer to “Entering a Permanent License Key” on page 419).

If you change the IP addresses periodically, you can install a Windows loopback adapter, and use the loopback IP address to obtain a CMS license. To install a loopback adapter, refer to:

<http://support.microsoft.com/kb/839013> or search support.microsoft.com.

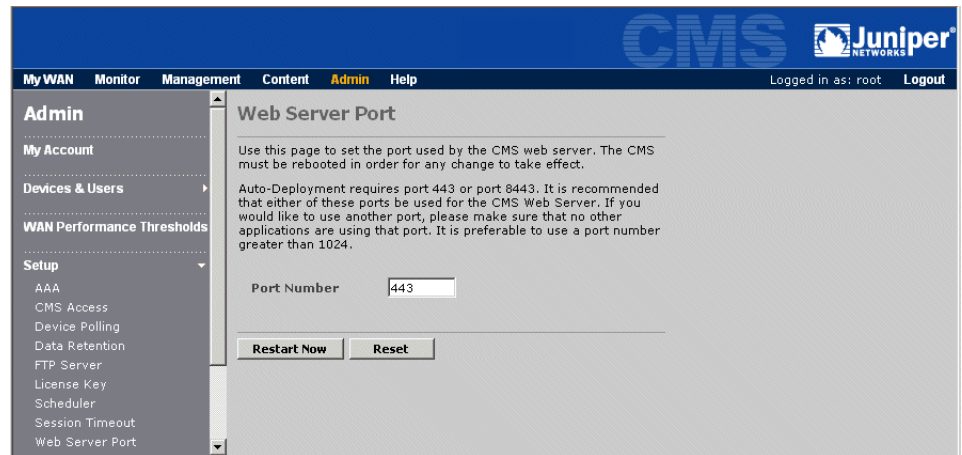
Changing the Web Server Port

By default, the CMS Web server uses port 443 (HTTPS). You can change this port if necessary. If you change the port, CMS must be restarted for the new port to take effect.

To change the CMS Web server port:

1. Click Admin in the menu frame, click Setup in the left-hand navigation frame, and then click Web Server Port.

Figure 233: Changing the Web Server Port Number



2. Enter the port number in the Port number field.

You can specify any unused port, but if 443 is not used, 8443 is recommended.

3. Click Restart Now to restart CMS and activate the new port number. To restore the original port number, click Reset.

After a restart, the Web console is redirected to the new port number in about 60 seconds.

Alternatively, you can also change the Web port number by changing the file at `<install_directory>/web server/conf/server.xml`. Search for:

```
<Connector port="443"
```

Change the port to the appropriate number (retain the quotation marks), and restart the JuniperCMS service.

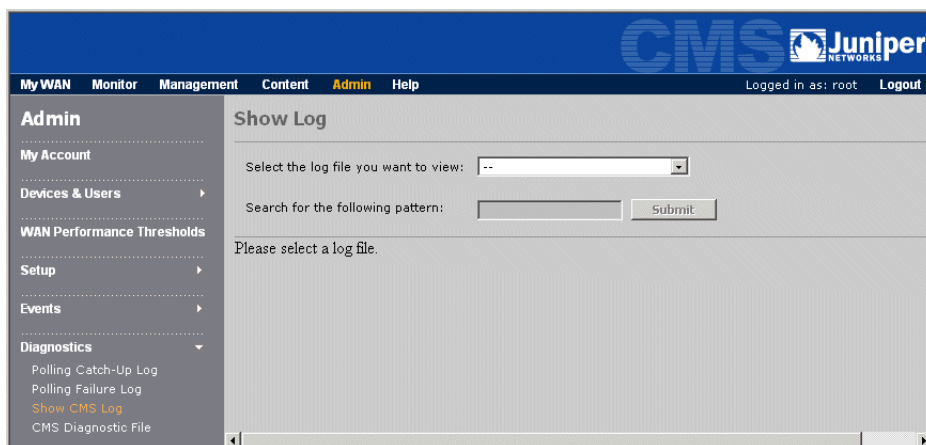
Viewing System Logs

The CMS system log files can be displayed in the Web console. If your network has dedicated Syslog servers, you can send CMS log messages to up to five Syslog servers, as described in “Enabling Syslog Reporting” on page 429.

To view the system logs:

1. Click Admin in the menu frame, click Diagnostics in the left-hand navigation frame, and then click Show Log.

Figure 234: Viewing the System Logs



2. Select the log you want to view.
 - **access.log.** Lists the time, user name, and client address of the most recent logins to CMS (also lists the times of failed logins). Each time access.log reaches 20 KB, it is saved as access.log.1, and existing logs are renumbered up to access.log.4 (older logs are discarded).
 - **activity.log.** Lists the various types of CMS tasks performed by each user since the last reboot.
 - **debug.log.** Lists the most recent informational messages, such as aggregation start and stop times, along with error messages, such as failures to communicate with a device. Each time debug.log reaches 5 MB, it is saved as debug.log.1, and existing logs are renumbered up to debug.log.10 (older logs are discarded).
3. To view all entries that have a specific text string, enter the text in the Pattern field (such as a partial IP address) and click Submit. To view all entries again, delete the text string and click Submit.

Generating a Diagnostic File

If you have problems with CMS, you can generate a diagnostic file to send to our support team. The diagnostic file contains current configuration, system information, and the most recent log files. By completing the form on the Diagnostic file page, your contact information is included with the file. After you generate and save the diagnostic file, email it to support@juniper.net.



NOTE: To generate a diagnostic file, the CMS Web server must be functioning.

To generate a diagnostic file:

1. Click Admin in the menu frame, click Diagnostics in the left-hand navigation frame, and then click CMS Diagnostic file.

2. Complete the form so that your contact information and a description of the problem is included with the diagnostic file.
3. Click Submit to generate the diagnostic file, and then click Save and specify a local file name and location.

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

Backing Up and Restoring the Database

The CMS database consists of a MySQL database and a set of independent files that define CMS schedules, device configurations, WXOS software images, and other configuration settings. You should periodically back up both sets of data. Note that the schedule log for content distribution is saved in the MySQL database, but the general device schedule log is not.

To back up the MySQL database, you can stop the server and back up the database manually, or use scripts to back up and restore the database automatically while the server is running. The non-MySQL configuration data must be backed up manually.

Manual MySQL Database Backups

To backup the MySQL database manually:

1. Stop the JuniperCMS service and the MySQL service:
 - a. Click Start > Run, enter “services.msc” and click OK.
 - b. In the Services window, right-click on JuniperCMS and click Stop, and then right-click on MySQL and click Stop.
2. Check the current size of the database in the following folder:

`<Install>\MySql\data`

 Where “< Install >” is the location where the database is installed. The default location is "C:\Program Files\Juniper Networks\CMS".
3. Copy the \MySql\data folder to a backup location that has sufficient disk space.
4. Restart the MySql and JuniperCMS services.

Restoring Manual MySQL Database Backups

To restore a database that was backed up manually:

1. Stop the JuniperCMS and MySql services.
2. Rename the <install>\MySql\data folder.
3. Copy the \data folder from the backup location to the <install>\MySql folder.

4. Restore the CMS configuration backup, if any (refer to “CMS Configuration Backups” on page 436).
5. If you are transferring the CMS database to a new server, you must remove the following files before starting the MySQL Windows service:

```
<install>/MySQL/data/ib_logfile0
<install>/MySQL/data/ib_logfile1
```

6. Restart the MySQL and JuniperCMS services.
7. Delete the old \data folder that was renamed in Step b.

Automatic MySQL Database Backups

Database backup and restore scripts are provided that let you back up the MySQL database while the server is running. Some data may not be backed up, so use this process only if you cannot stop the server.

To back up the MySQL database while the server is running:

1. Check the current size of the database in the following folder:

```
<Install>\MySQL\data
```

Where “< Install >” is the location where the database is installed. The default location is “C:\Program Files\Juniper Networks\CMS”.

2. Verify that the drive where you want to copy the database has at least twice the free disk space needed for the \data folder.
3. To execute the backup script, open a Command Prompt window and enter the following script name and the folder where you want to back up the database (or use the Windows “at” command to execute the script daily or weekly):

```
<Install>\MySQL\pnscripsts\InnoDBbackup.bat [<drive>:\]<backup_path>
```



NOTE: For the default installation location, enter “C:\Program Files\Juniper Networks\CMS” as “C:\Progra ~ 1\Junipe ~ 1\CMS”. Spaces are not supported in the path name.

If an error occurs during the backup, execute the backup script again.

Restoring Automatic MySQL Database Backups

To restore a database that was backed up automatically:

1. Stop the JuniperCMS service (the MySQL service must be running):
 - a. Click Start > Run, enter “services.msc” and click OK.
 - b. In the Services window, right-click on JuniperCMS and click Stop.
2. Open a Command Prompt window and enter the following script name and the folder where you previously backed up the database:

```
<Install>\MySQL\pnscripts\InnoDBrestore.bat [<drive>:\]<restore_path>
```



NOTE: For the default installation location, enter “C:\Program Files\Juniper Networks\CMS” as “C:\Progra ~ 1\Junipe ~ 1\CMS”. Spaces are not supported in the path name.

The current MySQL database is dropped, and a new database is created with the data imported from the backup database.

3. Restore the CMS configuration backup, if any (refer to “CMS Configuration Backups” on page 436).
4. Restart the JuniperCMS service.

CMS Configuration Backups

The CMS configuration data that is not stored in the MySQL database includes CMS device schedules, content distribution schedules, and device configurations, WXOS software images, customized My WAN settings for each user, and other miscellaneous settings. This data must be backed up manually.

To backup the non-MySQL configuration:

1. Stop the JuniperCMS service:
 - a. Click Start > Run, enter “services.msc” and click OK.
 - b. In the Services window, right-click on JuniperCMS and click Stop.
2. Check the current size of the following folder:

```
<Install>\CMS\data
```

Where “< Install >” is the location where CMS is installed. The default location is "C:\Program Files\Juniper Networks\CMS".

3. Copy the \CMS\data folder to a backup location that has sufficient disk space.
4. Restart the JuniperCMS service.
5. To restore a backed up configuration:
 - a. Stop the JuniperCMS service.
 - b. Rename the <install>\CMS\data folder.
 - c. Copy the \data folder from the backup location to the <install>\CMS folder.
 - d. Restart the JuniperCMS service.
 - e. Delete the old \data folder that was renamed in Step b.

Moving CMS to Another Disk Drive

To move CMS to another disk drive:

1. Copy the CMS database to a temporary folder on the new drive, as described in “Manual MySQL Database Backups” on page 434, but do not restart the CMS or MySQL services.
2. Uninstall CMS using the Microsoft Windows Add/Remove Programs function in the Control Panel. Elect NOT to remove the CMS database. This is not an option if you have a temporary CMS license.
3. Install CMS on the new drive (refer to “Installing CMS” on page 23).
4. Restore the CMS database that you saved in Step 1 (refer to “Manual MySQL Database Backups” on page 434).
5. When you are satisfied that CMS is running properly on the new drive, remove the old CMS installation directory.

Purging Temporary Java Files

In JRE version 1.5.0, Java cache files are accumulated in the Windows temporary folder on the CMS server. CMS will not start if the disk becomes full, so you should periodically delete the following files:

/WINNT/Temp/jar_cache.tmp*

Changing the CMS Time Zone

If you change the time zone of the CMS server, you must restart the CMS service for the change to take effect, as follows:

1. On the CMS server, select Start > Run, enter "services.msc" and click OK.
2. In the Services window, right-click on JuniperCMS and click Stop. Wait for the service to stop, and then right-click again and click Start.

Appendix A

CMS Licenses

The CMS license determines the maximum number of devices that CMS can manage. CMS has two types of licenses:

- Permanent license—If you have purchased CMS, you should receive an authorization code via email. Use the authorization code to obtain a permanent license key from https://www.juniper.net/generate_license. You will need the IP Address of your CMS server to obtain a permanent license key. You can then enter the permanent license key (refer to “Entering a Permanent License Key” on page 419).



NOTE: If you later change the IP address of the CMS server, you must reboot the server and obtain a new license key before you can continue using CMS. To avoid this, install a Windows loopback adapter, and use the loopback IP address to obtain a CMS license. To install a loopback adapter, refer to <http://support.microsoft.com/kb/839013> or search support.microsoft.com.

- Evaluation license—If you are evaluating CMS, you can use the evaluation license generated during the installation of the CMS software. The evaluation license is valid for 45 days and allows CMS to manage up to 5 devices. If you do not enter a permanent license key before the 45-day evaluation period expires, on the 46th day you will lose access to all CMS Web console pages, except the License Key page.

If you use an evaluation license during installation and subsequently purchase CMS and receive a permanent license, reinstallation is not necessary. Simply enter the permanent license key, as described in “Entering a Permanent License Key” on page 419.

If you want to manage a larger number of devices than is specified in your permanent license, you can purchase an upgrade license. Again, simply enter the new license key. Reinstallation is not necessary.

If you upgrade your CMS software to a newer version of software, the permanent license key last applied to CMS is retained and honored. To download the latest version of CMS, go to <https://www.juniper.net/customers/csc/software/>.

Appendix B

System Events

The following sections describe the SNMP traps and Syslog messages for the system events generated by CMS and the WX devices. All WX and CMS system events are displayed on the Events monitoring report, along with the WX performance events (refer to “Events Reports” on page 362).

- “Severity Levels” in the next section
- “CMS Syslog Messages” on page 442
- “WX System Events and SNMP Traps” on page 442
- “WX Syslog Messages” on page 447

Severity Levels

Some of the severity levels used by CMS event forwarding filters and the Events report filters are different from Syslog, as shown below. In the following sections, where the two terms are different, the Syslog term is shown first, followed by the CMS term in parentheses.

Syslog Severity	CMS Filter Severity
Notice	OK
Information	Warning
Error	Major
Critical	Critical

To define event forwarding filters, refer to “Defining Event Forwarding Filters” on page 413.

CMS Syslog Messages

Table 1 lists the Syslog messages generated by CMS.

Table 1: CMS System Events

Message	Warning: Exceeded licensed number of devices. Current licensed: < licensed# > and total devices: < actual# > .
Severity	Information (Warning)
Recommended Action	Call + 1-866-737-4248 to obtain a new license to support all of your devices.
Message	The schedule task < task name > for community (or device group) < dg name > failed.
Severity	Error (Major)
Recommended Action	To reschedule or acknowledge failed tasks, see “Managing Scheduled Tasks” on page 69
Message	The CMS was started.
Severity	Information (Warning)
Recommended Action	None

WX System Events and SNMP Traps

Table 2 describes each WX system event, and provides the SNMP trap name and its associated object ID (OID), if any. Table 3 describes the Syslog messages for each event.

Table 2: WX 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.	jnxWxEventCompressionBuffer Overflow 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.	jnxWxEventCompressionSession Closed 1.3.6.1.4.1.8239.2.2.1.3.2.0.3

Event Name	Severity	Description	SNMP Trap/OID
Compressor Session Opened	Notice (OK)	Compressor session to the device noted in jnxWxCommonEventDescr was opened.	jnxWxEventCompressionSession Opened 1.3.6.1.4.1.8239.2.2.1.3.2.0.5
DC Queue Above High Watermark	Information (Warning)	Decompression queue exceeded the high-watermark level.	jnxWxEventDCQAboveHiWatermark 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.	jnxWxEventDCQBelowHiWatermark 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.	jnxWxEventDecompressionSession Closed 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.	jnxWxEventDecompressionSession Opened 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.	jnxWxEventDiskFailure 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
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 noted by jnxWxCommonEventDescr.	jnxWxCommonEventInterface DuplexMismatch 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.	jnxWxCommonEventInterfaceSpeed Mismatch 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.	jnxWxCommonEventInterfaceSpeed Ok 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.	jnxWxCommonEventIpssecSecurity AssociationAdded 1.3.6.1.4.1.8239.2.1.3.2.0.15

Event Name	Severity	Description	SNMP Trap/OID
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
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
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.	jnxWxEventMultiPathStatusChange 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 289). If this system event is disabled, no SNMP traps or Syslog messages are generated for performance events.	jnxWxEventPerformanceThreshCrossed 1.3.6.1.4.1.8239.2.2.1.3.2.0.21

Event Name	Severity	Description	SNMP Trap/OID
Power Supply Failure	Error (Major)	One or more sources of power to the system has failed. A redundant power-supply has presumably taken over.	jnxWxCommonEventPowerSupplyFailure 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.	jnxWxCommonEventPowerSupplyOk 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).	jnxWxEventPrimaryDownBackupEngageFailed 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).	jnxWxEventPrimaryDownBackupEngaged 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.	jnxWxEventPrimaryRegServerUnreachable 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).	jnxWxEventPrimaryUpBackupDisengaged 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.	jnxWxEventRipAuthFailure 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.	jnxWxEventSecondaryRegServerUnreachable 1.3.6.1.4.1.8239.2.2.1.3.2.0.8
Secondary Registration IP Failed	Error (Major)	The IP address of the secondary registration server could not be set on the primary registration server.	None
Secondary Registration IP OK	Notice (OK)	The IP address of the secondary registration server was set on the primary 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 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 server, the link to a client device is down.	None
Stack Client Link Up	Notice (OK)	On a WX 100 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.	jnxWxCommonEventThruputLimit Exceeded 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

WX Syslog Messages

Table 3 lists the Syslog messages generated by WX devices. The critical- and error-level Syslog events for each device (if any) can be viewed on the Devices page (refer to “Viewing Device Events” on page 44).

Table 3: 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	<p>One of the following:</p> <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 top traffic data that can be exported to Cisco NetFlow servers, and the monitoring statistics that CMS can retrieve in CSV format from WX devices. For information about retrieving statistical data, see “Retrieving Device Files” on page 63.

This appendix covers the following sections:

- “NetFlow Version 5 Export” in the next section
- “Performance Statistics Export” on page 455
- “Top Traffic Export” on page 463

NetFlow Version 5 Export

Each device can export its top traffic data to a Cisco NetFlow server in Version 5 format (refer to “Generating NetFlow Records” on page 131).

Table 4 describes the NetFlow packet header.

Table 4: NetFlow Packet Header

Byte	Parameter	Description
0-1	Version	NetFlow export format version number (5).
2-3	Count	Number of flows exported in this packet (1 to 30).
4-7	Sysuptime	Number of milliseconds since the device was restarted.
8-11	Unix seconds	Number of seconds since 0000 1970 Coordinated Universal Time (UTC).
12-15	Unix nanoseconds	Residual nanoseconds since 0000 1970 UTC.
16-19	Flow number	Sequence counter of total flows seen.
20	Engine type	Not applicable.

Table 4: NetFlow Packet Header

Byte	Parameter	Description
21	Engine ID	Not applicable.
22-23	Sampling interval	Not applicable.

Table 5 describes each traffic flow entry in a NetFlow packet (up to 30 entries per packet).

Table 5: NetFlow Packet Entry

Byte	Parameter	Description
0-3	Srcaddr	Source IP address.
4-7	Dstaddr	Destination IP address.
8-11	Nexthop	Not applicable.
12-13	Input	SNMP index number of input interface.
14-15	Output	SNMP index number of output interface.
16-19	Packets	Number of packets in the flow.
20-23	Octets	Number of Layer 3 bytes in the flow.
24-27	First	SysUptime at start of flow.
28-31	Last	SysUptime when the last packet in the flow was received.
32-33	Source port	TCP/UDP source port number or equivalent.
34-35	Destination port	TCP/UDP destination port number or equivalent.
36	Pad1	Unused (zero).
37	TCP flags	Cumulative OR of TCP flags.
38	Protocol	IP protocol number (for example, TCP = 6; UDP = 17).
39	ToS	IP type of service.
40-41	Source system	Not applicable.
42-43	Destination system	Not applicable.
44	Source mask	Not applicable.
45	Destination mask	Not applicable.
46-47	Pad2	Unused (zero).

Performance Statistics Export

The following sections describe the performance data that can be extracted in CSV format from each device in the *AllData.csv* file.

- “General Device Information” in the next section
- “Data Section Information” on page 456
- “System Session Statistics” on page 456
- “Compression Session Statistics” on page 459
- “Application Session Statistics” on page 459
- “WAN Statistics” on page 460
- “Application Flow Acceleration Statistics” on page 460
- “Bandwidth Management Statistics” on page 461
- “WAN Performance Statistics” on page 462
- “Inbound Traffic By Port Statistics” on page 463

General Device Information

Table 6 describes the general device information.

Table 6: General Device Information

Parameter	Description
Device IP	IP address of the device.
Software version	Version of WXOS software that was running when the statistics were exported.
Serial number	Serial number of the device that exported the statistics.
License speed	Licensed speed of the device.
Monitor applications	Names of the applications being monitored.
Fast Connection applications	Names of the applications using Fast Connection Setup.
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 Profile mode.

Data Section Information

Table 7 describes the data section information that precedes the set of statistic tables for each time range.

Table 7: Data Section Information

Parameter	Description
< time > data section	Indicates the time range for the statistics tables that follow: <ul style="list-style-type: none"> ■ This hour ■ Last hour ■ Today ■ Yesterday ■ This week ■ Last week
ip =	IP address of the device.
device local time =	Local date and time of the export.
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 8 describes the system session statistics.

Table 8: 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).

Table 8: System Session Statistics

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 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 device in the same community.
Bytes PT SR-Hash	Number of bytes that passed through without compression because the data will be processed by another device.
Packets PT SR-Hash	Number of packets that passed through without compression because the data will be processed by another 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.
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 ¹ .

Table 8: System Session Statistics

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	Maximum compression rate for any five second interval within the selected time period. Peak percentage compression is calculated by the following formula: $10^5 \times \left(\frac{\text{Bytes In} - \text{Bytes Out}}{\text{Bytes In}} \right) = \text{Peak \% Reduction}$
Rsv H1 through Rsv H20	Reserved
Pkin1 to Pkin6	Six fields that show the number of packets in each of six packet-size ranges for traffic into the device, as follows: <ul style="list-style-type: none"> ■ Pkin1 Less than 64 bytes ■ Pkin2 64 to 127 ■ Pkin3 128 to 255 ■ Pkin4 256 to 511 ■ Pkin5 512 to 1023 ■ Pkin6 More than 1023 bytes
Pkout1 to Pkout6	Six fields that show the number of packets in each of six packet-size ranges for traffic out of the device, as follows: <ul style="list-style-type: none"> ■ Pkout1 Less than 64 bytes ■ Pkout2 64 to 127 ■ Pkout3 128 to 255 ■ Pkout4 256 to 511 ■ Pkout5 512 to 1023 ■ Pkout6 More than 1023 bytes

1. Data input is the number of IP bytes into the device from the Local port.

Compression Session Statistics

Table 9 describes the compression session statistics.

Table 9: 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 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 10 describes the application session statistics.

Table 10: Application Session Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
App Id	Application from which the data was received (e.g., FTP, HTTP, Lotus Notes).
Dst Ip	IP address of the device that receives compressed data from this device.
Bytes In	Number of bytes into the device that have been identified for compression, and addressed for the 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 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.

Table 10: Application Session Statistics

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.
FC Session Time	Number of milliseconds for all Fast Connection Setup sessions that ended during this time period.
Bytes Out NSC	Number of bytes out of this device after compression using NSC (WXC devices only), and addressed for the device listed with the destination IP address and application ID.

WAN Statistics

Table 11 describes the WAN statistics.

Table 11: WAN Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
App Id	Application for which the data was sent or received.
App Type	Type of application (0 = Default, 1 = HTTP, 3 = CIFS, 4 = Exchange).
Dst Ip	IP address of the remote device that sent or received data from this device.
Bytes From WAN	Number of bytes received from the WAN for the remote device and application.
Bytes To WAN	Number of bytes sent to the WAN for the remote device and application.

Application Flow Acceleration Statistics

Table 12 describes the Application Flow Acceleration statistics.

Table 12: Acceleration Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
App Id	Application for which the traffic was accelerated (HTTP, CIFS, or Exchange).
App Type	Type of application (0 = Default, 1 = HTTP, 3 = CIFS, 4 = Exchange).
Tran Id	Transaction ID number (0 = All, 1 = Bulk read/write).
Dst Ip	IP address of the remote device that received the accelerated traffic.
Time With Accel	Number of seconds required to complete the transaction.

Table 12: Acceleration Statistics

Time Without Accel	Estimated number of seconds required to complete the transaction with no acceleration.
--------------------	--

Bandwidth Management Statistics

Table 13 describes the bandwidth management statistics collected per application class for each tunnel.

Table 13: 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 device. Inbound bandwidth management: The total number of bytes into the Remote interface of the device by class.
Bytes Out	Outbound bandwidth management: The total number of application bytes out of outbound bandwidth management. Inbound bandwidth management: the total number of bytes out of inbound bandwidth management.
Bytes Dropped	Outbound bandwidth management: The total number of application bytes dropped by the bandwidth management feature. Inbound bandwidth management: The total number of bytes dropped by the bandwidth management feature.
Packets In	Outbound bandwidth management: The total number of application packets into the device. inbound bandwidth management: The total number of packets passed into the device by inbound bandwidth management.

Table 13: Bandwidth Management Statistics

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 14 describes the WAN performance statistics.

Table 14: WAN Performance Statistics

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Tunnel	IP address of a remote device.
Avg Latency	Average round-trip time to the remote device (in milliseconds). For hourly data, the median value is shown for each minute.
Latency Count	Number of minutes for which a latency value was measured.
Latency Above Thresh	Average percentage of minutes that the latency threshold was exceeded. For hourly data, the value is 0 or 1 for each minute (1 = above threshold).
Latency Above Thresh Count	Number of minutes for which the median latency exceeded the latency threshold.
Loss Pct	Average percentage of the probes that were lost.
Loss Count	Number of minutes for which a loss value was measured (excludes minutes for which none of the probes were returned).
Event Count	Number of times the loss or latency thresholds were exceeded or returned to normal.
Diversion Count	Number of times traffic was diverted to the alternate path (Multi-Path only).
Return Count	Number of times traffic was diverted back to the preferred path (Multi-Path only).
Last Down	Not used.
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 15 describes the Inbound traffic by port statistics.

Table 15: 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 16 describes the Traffic statistics retrieved in the *ip-flow.csv* file.

Table 16: Inbound Traffic By Port Data

Parameter	Description
Start Time	Start time for statistics generation.
End Time	End time for statistics generation.
Other Data	Number of bytes and packets sent and received for flows that exceeded the maximum retained by the device (16K for WX 15, 32K for WX 20, 65K for other models).
SrcIp	IP address of the flow source.
DstIp	IP address of the flow destination.
SrcPort	Source port number.
DstPort	Destination port number.
Proto	Traffic flow protocol (TCP, UDP, or protocol number).
Application	Traffic flow application name.
BytesSent	Number of bytes sent by the source.
PktsSent	Number of packets sent by the source.
BytesRcvdt	Number of bytes received by the source.
PktsRcvd	Number of packets received by the source.
TotalSendDelay	Cumulative delay between packets sent (in milliseconds).
TotalRcvDelay	Cumulative delay between packets received (in milliseconds).

Table 16: Inbound Traffic By Port Data

Type	Indicates the traffic type: <ul style="list-style-type: none">■ RA. Compressed application■ 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.

Appendix D

Common Application Port Numbers

Table 17 lists the common port numbers defined by the Internet Assigned Numbers Authority (<http://www.iana.org/assignments/port-numbers>).



NOTE: Port numbers 3577 and 3578 for TCP and UDP are reserved for data transmission.

Table 17: 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

Table 17: Common Application Port Numbers

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
https	443	TCP/UDP	https MCom
netnews	532	TCP/UDP	readnews
lotusnotes	1352	TCP/UDP	Lotus Notes
ms-sql-s	1433	TCP/UDP	Microsoft-SQL-Server
ms-sql-m	1434	TCP/UDP	Microsoft-SQL-Monitor
watcom-sql	1498	TCP/UDP	Watcom-SQL
orasrv	1525	TCP/UDP	Oracle
ccmail	3264	TCP/UDP	cc:mail/lotus

Glossary

access control list	List of IP addresses from which an administrator can log in to CMS.
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.
community	Two or more devices that can compress and decompress data for each other. Initially, all devices belong to the Default community. Each device contacts the registration server to identify the other devices in the same community.
compression subnets	Subnets that a device can advertise to the other devices in the community. The other devices can then compress traffic destined for those subnets.
content offers	Files used to preload WXC compression dictionaries so that traffic is optimized the first time the files are accessed.
decompression	Process used to restore compressed traffic to its original form.
endpoint	A device in a community or a virtual device used to apply outbound QoS policies to specific remote subnets.
filter	An operator defined IP address or TCP port number that determines valid addresses or applications for 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.

half-duplex	A mode of operation that allows only a single station to successfully transmit a frame at a given time.
hardware passthrough	Hardware-driven process by which all traffic is passed through the device at wire-speed. It is invoked automatically upon disruption.
HTTP	HyperText Transfer Protocol. The protocol most often used to transfer information from World Wide Web servers to browsers.
ICMP	Internet Control Message Protocol. An Internet Protocol used to communicate between devices on a network to manage errors and generate control messages.
Interior Gateway Protocol (IGP)	A group of protocols that provide routing information to the routers within an autonomous network.
Internet Protocol (IP)	The protocol that is used to route a data packet from its source to its destination over the Internet.
IP address	A numeric address, such as 10.10.124.22, assigned to every device on the network.
IP subnet mask	A numeric address, such as 255.255.0.0, used to define an IP subnet or to determine membership of an IP address in an IP subnet.
IP subnet	A group of IP addresses defined by the IP address and IP subnet mask pair, such as 10.10.0.0/255.255.0.0.
latency	The time necessary for a packet of data to travel from a source to a destination across a network.
log	A record of CMS activity. Logs are recorded for system information, performance, backup, and recovery.
MIB	Management Information Base. A database containing ongoing configuration information and statistics of a device in a network. MIBs are used with SNMP.
MTU	Maximum Transmission Unit. The largest size packet that can be transmitted by a device on a network.
OSPF	Open Shortest Path First. An interior gateway protocol that routes messages according to the least expensive path.
packet	A unit of data formatted for transmission on a network. Data is broken down into packets for sending over a packet switched network. Each packet has a header containing its source, destination, other control information, and a payload of data to be transmitted.

passthrough mode	A function of devices where all traffic passes through at wire-speed due to device disruption or overflow.
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.
registration server	The device that stores the network information for all devices in each community. Each device periodically contacts the registration server to identify the other devices in the same community.
response time	The time it takes for a host to respond to a user command.
RIP	See <i>Routing Information Protocol</i> .
round-trip time (RTT)	The time it takes to send a packet to a remote host and receive a response; used to measure delay on a network at a given time.
Routing Information Protocol (RIP)	An interior gateway protocol used in IP networks.
Simple Network Management Protocol (SNMP)	The Internet standard protocol for network management software.
Simple Network Time Protocol (SNTP)	A protocol that can synchronize clocks on local computers with radio or atomic clocks on the Internet.
software passthrough	Software-driven process by which a device transparently passes packets through the system without any processing.
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 in the data field of another packet.
User Datagram Protocol (UDP)	User Datagram Protocol. UDP is connectionless and does not guarantee reliable communication; the application itself must process any errors and check for reliable delivery. Defined in RFC 768.
warm reboot	A reboot of the device without powering off the unit.

Index

Numerics

3DES encryption for IPSec	278
802.1q VLAN support	120

A

AAA settings	
for CMS.....	421
for WX devices	132
acceleration	
CIFS	222
Exchange.....	225
Fast Connection Setup.....	218, 221
feature/topology setting	237
Forward Error Correction.....	218
HTTP.....	227
reports	344
TCP Acceleration.....	218, 220
access levels, CMS	388
access lists.....	420
access log	432
access URLs, device	390
acknowledging events.....	366, 395
acknowledging failed tasks.....	73
active FTP.....	147
activity log.....	432
advertising compression subnets	109
AES encryption for IPSec	278
aggregate WAN speed	
about.....	174
defining inbound QoS	205
defining outbound QoS	195
analyzing device configurations	52
Application Flow Acceleration	
about.....	210
CIFS and Exchange reports.....	351
CIFS traffic	222
Exchange traffic	225
feature/topology setting	237
HTTP reports	353
HTTP traffic.....	227
applications	
accelerating	
CIFS traffic	222

Exchange traffic.....	225
Fast Connection Setup	221
HTTP traffic.....	227
TCP Acceleration	220
assigning to traffic classes	149
common port numbers.....	465
compressing.....	156
managing.....	141
monitoring.....	151
monitoring percent compression	326
summary statistics	
all traffic.....	328
WAN traffic	319
visibility in tunnels	168
ARP, configuring.....	108
authentication	
for CMS	421
for SMTP servers.....	428
for WX devices.....	133
Authorization Codes	
importing.....	297
matching with devices	299
auto-deployment	
about	287
of configurations and software	288
of device licenses.....	296
status.....	294

B

backing up	
device configurations	60
the database	434
balancing, load	
across routers.....	128
across tunnels	160
bandwidth detection.....	187, 198
bandwidth management	
inbound.....	203
outbound	169
boot images	
downgrading	45
download and upload privileges	395
rolling back.....	48

upgrading	45	compression	
uploading to CMS	408	monitoring	321
browser support, CMS	19	statistics	359
bypass condition, multi-path	267	compression percentage, monitoring	321
C		compression subnets	
caching, for HTTP acceleration	212	configuring	109
cancelling		filtering source/destination	238
content distribution tasks	380	Configuration window	99
pending tasks	72, 381	configuration, initial	31
scheduled reports	371	configurations	
carving out addresses		about	75
and outbound QoS	180	analyzing	52
from off-path RIP advertisements	243	backing up	60
catch-up log, polling	409	changing	98
CIFS acceleration		changing a reference	97
about	210	CLI commands, appending	262
configuring	222	comparing	91
feature/topology setting	237	deleting	98
reports	351	displaying	92
circuit speeds		exporting	96
and router overhead	173	generating	
configuring	185, 197	creating	88
Citrix names, in application definitions	149	duplicating	87
classes, traffic		extracting from devices	54, 85
in the QoS Wizard	188	importing	95
inbound QoS	204	management recommendations	82
outbound QoS and Multi-Path	149	previewing before loading	57, 292
clean reboot for WXC devices	50	publishing	94
CLI commands, appending	262	restoring	61
client access, CMS	420	rolling back	58
CMS Web console		summaries by device	51
about	21	upgrading	98
browser support	19	verifying running configurations	51
logging in and out	20, 31	version tracking	82
user accounts	387, 388	viewing history	93
viewing logged in users	396	Configurations page	83
CMS Web server port		consistency checking	82
changing	431	content distribution	
default	26, 29	accessing network drives	383
communities		defining and distributing content	373
importing	397	distribution groups	378
migrating to another registration server ...		viewing and cancelling tasks	380
401		viewing and exporting the schedule log	382
Compressed traffic class	204	conventions, document	15

D

- data collection.....405
- data packets, Forward Error Correction.....219
- data retention407
- database
 - backups434
 - growth estimates405
- dead-time interval, RADIUS136
- debug log.....432
- decompressors
 - default162
 - preferred165
- dedicated WANs174
- default gateway, configuring103
- Default traffic class
 - inbound QoS.....204
 - outbound QoS and Multi-Path149
- deployment groups290
- deployment records292
- device access URLs390
- device groups, managing.....402
- device names104
- device time, viewing reports in.....305
- devices
 - about38
 - accessing the Web console68
 - adding to scheduled tasks.....73
 - cancelling pending tasks72, 381
 - data retention407
 - events
 - list of.....441
 - viewing.....44
 - exporting information68
 - failed tasks
 - icon on Devices page41
 - rescheduling and acknowledging73
 - groups of402
 - icons40
 - monitoring
 - inbound QoS340
 - outbound QoS336
 - percentage compression.....321
 - tunnel status333
 - WAN performance.....311
 - polling.....405
 - rebooting.....50
 - safe mode67
 - scheduling tasks for.....43
 - supported by CMS20
 - verifying running configurations51
 - viewing38
- Devices page.....38
- DHCP and auto-deployment.....287
- diagnostic files
 - CMS433
 - retrieving from devices63
- distributing content373
- distribution groups.....378
- distribution lists for event alerts.....412
- DNS servers
 - and auto-deployment.....287
 - configuring for the CMS Traffic report .357
- DNS servers, configuring for the device Traffic report.....105
- domain names
 - in device access URLs390, 422
 - in the Traffic report105, 357
- DSCP values, see "ToS/DSCP values"
- duplicating configurations87
- dynamic routes
 - OSPF and RIP126
 - router polling.....114

E

- email distribution lists for events.....412
- emailing reports368
 - defining an SMTP server.....427
 - managing scheduled reports.....369
- encryption, see "IPSec"
- endpoints
 - compression.....152
 - IPSec275
 - Multi-Path268
 - NSC.....154
 - outbound QoS.....195
 - packet flow acceleration214
 - WAN performance monitoring.....258
- events
 - acknowledging.....366, 395
 - configuring WX.....280
 - email distribution lists.....412
 - enabling CMS Syslog429
 - filters for forwarding413
 - list of CMS and WX441

retention policy	407	FTP server	
viewing Events reports	362	installing	26
viewing on the Devices page	44	on CMS server	63, 426
Exchange acceleration		G	
about	210	gateways, configuring	
configuring	225	default	103
feature/topology setting	237	in Multi-Path configurations	117
reports	351	global configurations	75
Executive report	358	groups	
exporting		CMS device	402
CMS configurations	96	CMS user	391
community and device information	68	guaranteed bandwidths	
content distribution logs	382	configuring	189, 194
device schedule logs	74	overriding	193
Fast Connection Setup data	350	H	
percentage compression, device..	325, 327	hardware requirements	23
QoS data	337, 338, 341	high-availability support	119
TCP Acceleration data	348	HMAC/SHA-1 authentication for IPsec	278
external routing for packet interception	244	HTTP acceleration	
extracting configurations	54, 85	about	212
F		configuring	227
failed tasks		feature/topology setting	237
icon on Devices page	41	reports	353
rescheduling and acknowledging	73	HTTPS	19
failure log, polling	409	Hub and Spoke topology	229, 230, 233
Fast Connection Setup		I	
configuring	218, 221	IANA port map	357
report	348	icons on Devices page	40
features, CMS	17	idle user timeout	137
features/topology, configuring	237	importing configurations	95
file locations, CMS and JRE	29	inbound QoS	
files		configuring	203
diagnostic, CMS	433	monitoring	340
retrieving from devices	63	installing CMS	23
filters		interface	
compression		link failure propagation	120
application	156	settings, configuring	119
source/destination	238	Intranet traffic class	204
content distribution	376	IP addresses	
event forwarding	413	primary device	103
flow statistics, retrieving	63	secondary, for Multi-Path	117
Forward Error Correction, configuring	218	IP compression tunnel mode, configuring ..	167
forwarding events	413		
front panel access, device	140		
FTP application type	147		

- IPSec
 - configuration procedure.....274
 - defining templates277
 - icon on Devices page.....41
- J**
- Java Runtime Environment (JRE)
 - location of files29
 - version.....26
- K**
- key lifetimes, IPSec278
- keys, RADIUS136
- L**
- latency threshold
 - Multi-Path.....269
 - WAN performance monitoring259
- Layer 2 multicast traffic331
- LDAP authentication server.....421
- license keys
 - CMS
 - about439
 - entering.....419
 - device
 - deployment procedure.....296
 - generating and applying.....299
 - importing Authorization Codes297
 - viewing status302
- License Server, accessing302
- lifetimes, IPSec key278
- link failure propagation.....120
- load balancing
 - across routers128
 - across tunnels.....160
- local domain name105
- local routes, adding static112
- local users, device.....137
- locations of CMS and JRE files29
- logging in and out20, 31
- login retries, SSH134
- logs
 - polling catch-up and failure409
 - system432
- logs, retrieving from devices63
- loopback IP address430
- M**
- MAC addresses108
- maximum bandwidths
 - inbound205
 - outbound
 - configuring189, 194
 - overriding193
- max-mem topology setting.....234
- MD5
 - for IPSec278
 - for OSPF127
- Mesh topology229, 230, 233
- meta packets
 - IP compression167
 - ToS/DSCP values in199
- metrics
 - in event forwarding filters416
 - in WX performance events282
- migrating communities401
- minimum WAN speed187, 198
- monitoring
 - applications151
 - inbound QoS340
 - outbound QoS336
 - percentage compression321
 - tunnel status.....333
 - WAN performance311
 - configuring258
- Monitoring pages305
- multi-flow emulation168
- Multi-Path, configuring
 - about263
 - addresses.....116
 - defining endpoints268
 - defining templates.....266
 - icon on Devices page41
 - router configuration270
- multiple tunnels on the WX 100260
- My Account page.....387
- My WAN page306
- N**
- NAT
 - changing registration server addresses 402
 - importing communities399

mapping public to private addresses	38, 400
NetFlow records, generating	131
network	
interfaces, configuring	119
settings, configuring	103
network drives, accessing	383
Not Applicable icon	323
NSC	
configuring applications	156
defining endpoints	154
icon on Devices page	41
NTP	
configuring	122
icon on Devices page	40

O

off-path deployments, configuring	242
operator access, device	138
OSPF	126
outbound QoS	
about	169
and packet flow acceleration	214
bandwidth detection	187, 198
configuration procedure	180
dedicated and oversubscribed WANs ..	174
defining endpoints	184, 185, 195
defining settings by endpoint	191, 275
defining templates	193
defining traffic classes	149, 188
excluding LAN/WAN addresses	111
monitoring	336
outbound speed	
about	174
defining	183, 195, 196
running the Setup Wizard	182
starting and stopping	202
ToS/DSCP values	199
virtual endpoints	199
outbound speed	183, 196
overflow, traffic volume	331
oversubscribed WANs	174
overviews	
CMS	19
configurations	75

P

Packet Flow Acceleration	
Fast Connection Setup	218, 221
Forward Error Correction	218
icon on Devices page	41
TCP Acceleration	218, 220
packet interception	
configuring	242
icons on Devices page	41
packet size distribution statistics	332
pages	
Auto-Deployment	289
Configurations	83
Devices	38
License Management	296
Monitoring	305
Scheduled Reports	369
Scheduled Tasks	69
partial configurations	75
partitioning users by customer	386
passthrough statistics	330
passwords	
CMS users	387
device	137
OSPF	127
registration server	
applying to devices	65
updating in CMS	402
RIP	127, 245
SMTP authentication	428
peak data compression	359
pending device tasks	
cancelling	72
pending tasks	
adding devices to	73
cancelling	371, 381
performance events, configuring	282
performance monitoring, WAN	
configuring	258
viewing reports	312
Point-to-Point topology	233
policy-based routing for packet interception ...	
244	
polling	
catch-up and failure logs	409
intervals	405

port numbers
 application465
 in application definitions148
 RADIUS server136
 port, CMS Web server
 changing.....431
 default26, 29
 preferred path.....267
 pre-fetch for HTTP acceleration.....212
 pre-installation tasks25
 prime time, defining241
 private IP addresses
 mapping to public addresses.....38, 400
 registration server399, 402
 privilege levels
 for CMS accounts388
 for device accounts.....137
 protocols, in application definitions148
 publishing configurations94

Q

QoS
 inbound203
 outbound, see "outbound QoS"
 quick setup31

R

RADIUS servers and server groups135
 read-only access388
 rebooting devices50
 recommended tasks.....35
 recovery packets, Forward Error Correction
 219
 recurring tasks, adding devices to73
 referenced configurations, changing.....97
 registration server
 changing addresses401
 designating130
 password
 applying to devices.....65
 updating in CMS.....402
 remote authentication
 for CMS
 defining the server.....421
 enabling by account389
 for WX devices132

remote circuit speeds
 and router overhead173
 configuring185, 197
 remote routes158
 reports
 about305
 Acceleration Summary344
 Application Summary
 all traffic328
 WAN traffic319
 CIFS and Exchange acceleration351
 compression.....321
 emailing.....368
 Events362
 Executive358
 Fast Connection Setup348
 HTTP acceleration353
 inbound QoS340
 managing scheduled reports.....369
 My WAN page306
 outbound QoS336
 packet size distribution.....332
 Passthrough Data330
 scheduling368
 TCP Acceleration346
 throughput, WAN traffic317
 traffic.....355
 trend.....360
 tunnel status.....333
 viewing in device time.....305
 WAN311
 rescheduling failed tasks73
 restoring
 backup databases434
 configurations61
 retransmissions, RADIUS136
 retries, SSH login.....134
 retrieving device files and statistics63
 RIP
 for dynamic routes126
 for packet interception243
 roles, defining user393
 rolling back
 boot images.....48
 configurations58
 root user account20
 route injection243

router balancing, route-based	128	NetFlow	131
router configuration		RADIUS	135
for Multi-Path	270	service providers, partitioning users for	386
for packet interception	247	Setup Wizard, outbound QoS	182
routes		severity levels	
adding static	112	in event forwarding filters	415
remote	158	of system events	441
router polling	114	of WX performance events	283
S		SMB signing, disabling	223
safe mode, devices	67	SMTP server for emailed reports	427
schedule logs		SNMP	123
content distribution	382	software requirements	23
device	74	source address, RADIUS	118
scheduled reports		source/destination filters	238
defining an SMTP server	427	Spoke topology	229, 230, 233
emailing	368	static routes, adding	112
managing	369	statistics	
scheduled tasks		acceleration	344
failed tasks	73	application	
selecting devices for	43	all traffic	328
viewing details	69	WAN traffic	319
scheduler, stopping and restarting	430	compression percentage	321
secondary IP address for Multi-Path	117	events	362
secret key, RADIUS	136	executive summary	358
security		inbound QoS	340
CMS		outbound QoS	336
access lists	420	packet size distribution	332
enabling accounts for remote		passthrough traffic	330
authentication	389	retrieving from devices	63
remote authentication	421	throughput, WAN traffic	317
user accounts	388	top traffic	355
device		trend reports	360
about	132	understanding retrieved data	453
defining local users	137	WAN	311
defining RADIUS servers and server		status	
groups	135	of applied licenses	302
front panel access	140	of auto-deployment	294
operator access	138	subnet mask, configuring	103
selecting authentication methods ...	133	subnets	
security features		advertising for compression	109
for CMS		defining whether encryption is required ...	279
user accounts	385	excluding from compression	238
server time, viewing reports in	305	excluding from outbound QoS	111
servers		filtering the Traffic report	357
DHCP	287	unadvertised subnets and outbound QoS..	180
DNS	105, 287, 357		

subnets, excluding from default	
decompressors.....	164
summary of device configurations.....	51
support	
browser	19
generating CMS diagnostic files	433
retrieving diagnostic files from devices	63
technical.....	15
Syslog	
enabling CMS.....	429
enabling on devices	124
list of events	441
retrieving from devices	63
system events	
enabling/disabling WX	285
list of CMS and WX.....	441
system log	432

T

tasks, managing scheduled	69
TCP Acceleration	
configuring.....	218, 220
feature/topology setting	237
report.....	346
TCP traffic class	204
technical support.....	15
templates	
IPSec.....	277
Multi-Path.....	266
outbound QoS	193
thresholds, loss and latency	
Multi-Path.....	269
WAN performance monitoring	259
thresholds, WX performance event	282
throughput statistics, WAN traffic.....	317
time settings	
CMS time zone	438
NTP server	122
WX time zone and daylight savings	107
timeout	
idle user.....	137
RADIUS server.....	136
topology settings	229
ToS/DSCP values	
defining by QoS traffic class.....	199
in application definitions	148
in Multi-Path configurations	265

traffic classes	
in the QoS Wizard	188
inbound QoS	204
outbound QoS and Multi-Path.....	149
traffic statistics	355
trend reports.....	360
tunnel mode	167
tunnels	
configuring endpoints	152
monitoring.....	333
types of applications.....	147

U

UDP	
and meta packets	167
in application definitions.....	148
unadvertised subnets and outbound QoS..	180
uninstalling CMS	31
upgrading	
configurations	98
the CMS software	27
URLs	
in application definitions.....	149
in device access links.....	390
user accounts	
CMS	
about.....	20
changing passwords.....	387
defining	388
defining roles for	393
device.....	137
user groups	391
users, logged in	396

V

validating remote routes	159
verifying running configurations.....	51
versions, configuration	82
virtual endpoints, outbound QoS	199
VLAN 802.1q support.....	120

W

WAN circuit speeds	
about	173
bandwidth detection	187, 198
configuring	197

WAN compression subnet	
for off-path devices	111
for VLAN environments.....	121
WAN performance monitoring	
configuring devices	258
configuring thresholds for reports.....	418
reports	311
WCCP for packet interception	243
Web console	
CMS	
about.....	21
logging in and out	20, 31
device, accessing.....	68
Web server port, CMS	
changing.....	431
default.....	26, 29
Weighted Fair Queuing	190, 203
Weighted Strict Priority	190, 203
Wizard, outbound QoS.....	182
WX 100 clients	
and multiple tunnels.....	260
maximum tunnels supported	232
on Tunnel Status report.....	334