FX Series

Administrator Guide

Version 6.2.2

IMPORTANT NOTE: The information contained in this document supersedes all previously published information regarding this product. Product specifications are subject to change without prior notice.

MN-FXSERIESADM6 Revision 6

Table of Contents

Table of Contents iii
Table of Figures vii

Using This Document ix

Document Organization ix
Contacting Product Support x
Key FX Series Appliance Information x
FX Series End User License Agreement xi
Patents and Trademarks xiii
Conventions and References xiv
Comtech EF Data Warranty Policy xv
Release Notes xvii
Version 6.2.2 Functionality Enhancements xvii
Version 6.1.1 Functionality Enhancements xvii
Version 6.1 Functionality Enhancements xviii
Version 6.0.3 Functionality Enhancements xix
Version 6.0.2 Functionality Enhancements xix
Version 6.0.1 Functionality Enhancements xx

1 Overview - FX Series 22

1.1 Stampede FX Series Product Line Update 22
1.2 Technologies that Optimize Satellite Bandwidth Acceleration 23
1.3 Single-Sided Solution 25
1.3.1 Load Balancing via WCCP 25
1.3.2 Source IP Preservation 25
1.3.3 Connection Management 26
1.3.4 ACM QoS 26
1.3.5 GZIP Compression 27
1.3.6 Image Reduction and Smoothing 27
1.3.7 Static Caching 28
1.3.8 TCP Optimization 28
1.4 Two-Sided Solution 29
1.4.1 Cache Differencing 29
1.4.2 Multiplexing of Large Data Objects 29
1.4.3 Partial Content Update Caching 30
1.4.4 Network Protocol Optimization 30
1.4.5 Dynamic Data Deduplication 30
1.4.6 Header Compression/Packet Aggregation 30
1.4.7 Multicator 31
1.5 FX Series Appliances 32
1.5.1 Theory of Operation 32
1.5.2 Reporting 32
1.5.3 Deployment Options 33
1.6 FX Series Remote Appliance 34
1.6.1 Theory of Operation 34
1.6.2 Reporting 35
1.6.3 Deployment Options 35
1.7 Mesh Networking with the FX Series 36

1.7.1 Theory of Operation 36
1.7.2 Mesh Capability with two FX Series appliances at each node 36
1.7.3 Mesh configuration with Redundancy 38
1.8 FX Series Appliances Data Sheet 39
1.8.1 Single Sided with the Application Delivery Controller (ADC) 39
1.8.2 Two Sided with the ADC and the Remote 39
1.8.3 Configuration Models 40
1.8.4 FX Series Hardware Specification 41
1.8.5 FX-4010 Physical Description 42
1.8.6 FX4010 DC Physical Description 43
1.8.7 FX-1005 Physical Description 45
1.8.8 FX Series FX-1005 Hardware Mounting Options 47
1.8.9 FX-1010 Physical Description 48

2 Initial Installation Information 50

2.1 Pre-Installation Information 50
2.1.1 Unpacking 50
2.1.2 User Interfaces 50
2.1.3 Documentation 50
2.2 How to Configure Appliance Management Address 51
2.2.1 All Installation Patterns 51
2.3 How to configure FX Series Installation Pattern (In-Line Mode) 52
2.3.1 Cable the Appliance 52
2.3.2 Configure the Appliance 52
2.4 How to configure FX Series Installation Pattern (Routed Mode) 54
2.4.1 Cable the Appliance 54
2.4.2 Configure the Appliance 54
2.5 How to configure FX Series Installation Pattern (WCCP Mode) 56
2.5.1 Cable the Appliance 56
2.5.2 Configure the Appliance 56
2.5.3 Configure WCCP Settings 57
2.6 How to Configure Two FX Series Appliances in a Mesh Configuration 59
2.6.1 Cable the Appliances 59
2.6.2 Configure the appliances 59
2.6.3 Mesh installation with Redundancy capability 59

3 FX Series Configuration 60

3.1 Standard Configuration Overview 60
3.2 Management Settings 62
3.2.1 How to Configure Network Interfaces 62
3.2.2 How to Configure Host/DNS Settings 64
3.2.3 How to Configure SNMP Settings 65
3.2.4 How to Configure Web Admin Settings 66
3.3 Traffic Interface Settings 67
3.3.1 How to Configure In-Path Interfaces 67
3.3.2 How to Configure LAN Interfaces 69
3.3.3 How to Configure Port Definitions 70
3.3.4 How to Configure WCCP 72
3.4 Quality of Service 79
3.4.1 Theory of Operations for QoS and Traffic Shaping 79
3.4.2 QoS Configuration Hierarchy Screen 81

3.4.3 How to Configure QoS Links 84
3.4.4 How to Configure QoS Groups 88
3.4.5 How to Configure QoS Group Filters 89
3.4.6 How to Configure QOS Queues 91
3.4.7 How to Configure QoS Queue Filters 93
3.5 FX Series Multicator 95
3.5.1 Multicator Settings 96
3.5.2 How to set the Multicator General Configuration 98
3.5.3 How to set the Multicator Controller Configuration 98
3.5.4 How to set the Multicator Transmitter Configuration 98
3.5.5 How to set the Multicator Receiver Configuration 99
3.6 Redundancy 100
3.6.1 Redundancy Configuration Settings 100
3.6.2 How to Configure Key-Exchange 101
3.6.3 How to Configure 1:1 Redundancy with Fail Over 102
3.6.4 How to Synchronize Configurations in a WCCP Cluster 103

4 FX Series ADC General Settings 104

4.1.1 How to Configure FX ADC in 'Configuration-Only' mode: 105
4.1.2 Object Retrieval Logging 106
4.1.3 Traffic Interception 106
4.1.4 Active Flows 107
4.1.5 System Time 107
4.1.6 Software Updates 107
4.1.7 Other 108

5 FX Series Remote General Settings 109

5.1 System Time 109
5.2 Traffic Interception 110
5.2.1 How to Configure FX Remote in 'Configuration-Only' mode: 110
5.3 Active Flows 112
5.4 Other 112

6 FX Series Status 113

6.1 FX Series ADC Status 113
6.1.1 FX Series ADC WANOP Monitor 114
6.1.2 FX Series ADC Current Statistics 117
6.2 QOS Monitors 123
6.2.1 QoS Link Monitor 124
6.2.2 QOS Queue Monitor 126
6.3 FX Series Remote Status 129
6.3.1 FX Series Remote WANOP Monitor 129
6.3.2 FX Series Remote Current Status Reports 131

7 FX Series Optimization Settings 133

7.1 Application Policies Overview 133
7.1.1 FX Series Optimization Summary 133
7.1.2 Single-Sided Optimizations: 134
7.1.3 Two-sided Optimizations 134
7.1.4 Authorization Realms 134

7.1.5 Web Application Policies 136
7.1.6 Authorization realm 137
7.1.7 Enable Acceleration 137
7.1.8 Allow Access 137
7.1.9 Caching 137
7.1.10 Content Validation 138
7.1.11 Image Optimization 138
7.1.12 Back-End Server Interface Options 139
7.1.13 When Application Policies Take Effect: 139
7.1.14 Web Application Firewall 139
7.2 How to Configure Basic Web Application Policies 140
7.2.1 How to Set the Policy Applicability 140
7.2.2 How to Set Specific Users Access 141
7.2.3 How to Restrict Acceleration for Specific Sites, or Users 141
7.2.4 How to Set Specific Optimization Techniques 142
7.3 Layer 5 Application Policies 143
7.3.1 How to Configure Certified Applications 143
7.3.2 How to Configure Other Applications 143
7.3.3 How to Configure Layer 5 Optimizations 144
7.3.4 Layer 5 Protocols 145
7.3.5 ToS handling method 145
7.3.6 Layer 5 Acceleration - Discussion 146

8 FX Series Operations Features 147

8.1 Basic Operations Functions 148
8.1.1 How to Backup/Restore Configuration Files 148
8.1.2 How to Initiate Disaster Recovery Procedure 149
8.1.3 How to Change Passwords 149
8.1.4 How to Manage Licenses / Fast Codes 150
8.1.5 How to Shutdown/Restart 151
8.2 How to do Network Trouble Shooting with Packet Capture 152
8.3 How to Update Software 154
8.3.1 Software Update Discussion 154
8.3.2 How to Download and Apply Image from ADC (FX Remote Only): 155

9 FX Series Support 156

9.1.1 Support Contact Information 156
9.1.2 SNMP MIBS 156
9.1.3 Product Information and Support Links 157

10 Appendix 158

10.1 Sample Acceleration Status Reports 158
10.2 FX Series Console Management Functions 159
10.3 How to Update FX Series Appliance Software at 5.78.0 or earlier 161
10.3.1 Base Platform Image (BPI) Upgrade Process 161
10.3.2 Upgrade Kit and Prep 161
10.3.3 The Upgrade Process 161

Table of Figures

Figure 1-1 FX Series Multicator Theory of Operation 31
Figure 1-2 FX Series Basic Mesh Connectivity Diagram 36
Figure 1-3 FX Series Hub Spoke Mesh Connectivity Diagram 37
Figure 1-4 FX Series Mesh with Redundancy Connectivity Diagram 38
Figure 1-5 FX Series Appliances Data Sheet 39
Figure 1-6 FX Series Hardware Specifications 41
Figure 1-7 FX Series FX-4010 Back Panel 42
Figure 1-8 FX Series FX-1005 Front Panel 45
Figure 1-9 FX Series FX-1005 Rear Panel 46
Figure 1-10 FX Series FX-1010 Front Panel 48
Figure 1-11 FX Series FX-1010 Rear Panel 49
Figure 2-1 FX Series Mesh Connection Diagram 59
Figure 3-1 FX Series Main Configuration Screen 60
Figure 3-2 FX Series Basic Network Interfaces Screen 62
Figure 3-3 FX Series Host/DNS Settings Screen 64
Figure 3-4 FX Series SNMP Edit Screen 65
Figure 3-5 FX Series Web Management Interface Screen 66
Figure 3-6 FX Series In-Path Interfaces Screen 67
Figure 3-7 FX Series LAN Interfaces Screen 69
Figure 3-8 FX Series Port Definitions Screen 70
Figure 3-9 FX Series WCCP Definitions Screen 73
Figure 3-10 FX Series QoS Hierarchy Screen 81
Figure 3-11 FX Series QoS Links Screen 84
Figure 3-12 F Series QoS Link Edit Screen 84
Figure 3-13 FX Series ACM QOS Status by Modem Report 87
Figure 3-14 FX Series QoS Groups 88
Figure 3-15 FX Series QoS Group Edit Screen 88
Figure 3-16 FX Series QoS Group Filters 89
Figure 3-17 FX Series QoS Group Filters Edit Screen 90
Figure 3-18 FX Series QoS Queues 91
Figure 3-19 FX Series QoS Queue Edit Screen 91
Figure 3-20 FX Series QoS Queue Filter Edit Screen 93
Figure 3-21 FX Series Multicator General/Controller Edit Screen 96
Figure 3-22 FX Series Multicator Transmitter/Receiver Edit Screen 97
Figure 3-23 FX Series Redundancy Edit Screen 100
Figure 4-1 FX Series ADC General Edit Screen 104
Figure 5-1 FX Series Remote General Edit Screen 109
Figure 6-1 FX Series Status Menu 113
Figure 6-2 FX Series ADC WANOP Monitor Screen 114
Figure 6-3 FX Series ADC Current Status Menu 117
Figure 6-4 FX Series QOS Monitor by Link 124
Figure 6-5 FX Series QoS Monitor by Queue 126
Figure 6-6 FX Series Remote Real-Time Monitor Screen 129
Figure 6-7 FX Series Remote Current Status Menu 131
Figure 7-1 FX Series ADC Features Menu 133
Figure 7-2 FX Series Application Policy Applicability Edit Screen 140
Figure 7-3 FX Series Authorization Realm Edit Screen 141
Figure 7-4 FX Series Specific Optimization Edit Screen 142
Figure 7-5 FX Series Layer 5 Policy Configuration Edit Screen 143

Figure 7-6 FX Series TCP/UDP Ports Table 146
Figure 8-1 FX Series Operations Menu 147
Figure 8-2 FX Series Backup and Restore Screen 148
Figure 8-3 FX Series Change Passwords Screen 149
Figure 8-4 FX Series Upgrade Fast Codes Screen 150
Figure 8-5 FX Series Shutdown/Restart Screen 151
Figure 8-6 FX Series Packet Capture Screen 152
Figure 8-7 FX Series Software Version Display Screen 154

Using This Document

This guide was prepared to assist you in the installation, configuration and management of the FX Series Appliances. This document contains the same information that is available thru the on-line help contained with the FX Series web based administrative screens. This document supports Release 6.2.0 of the FX Series Appliances.

Document Organization

Release Notes
This section delineates the major changes from the prior release.

Theory of Optimization
This section discusses the characteristic of data transmission that will cause slow response and higher bandwidth requirements. It also delineates techniques that can reduce the slowness and help reduce bandwidth requirements.

FX Series Technology
This section provides a brief description of the hardware and optimization techniques available through the FX Series of appliances.

FX Series Installation Patterns
This section provides instruction on how to install the available configurations for all FX Series appliances. It covers In-Line Mode, Routed Mode, WCCP Mode and Mesh Configurations

FX Series Network Settings

This section discusses how to set the basic networking parameters, such as Management Settings, Traffic Interfaces, Quality of Service, Multicator Settings, and Redundancy.

FX Series ADC Specific Settings

This section discusses ADC specific network settings and current performance status, including General Settings.

FX Series Remote Specific Settings
This section discusses Remote specific network settings and current performance status, including General Settings and In-Path Settings.

FX Series Status
This section provides real time monitors and current statistics for the FX Series ADC, the FX Series Remote and for Quality of Service.

FX Series Optimization
This section discusses optimization issues and definitions of optimization techniques available on the FX Series appliances for web based or enterprise applications.

FX Series Operations
This section discusses tools to perform operational tasks, including Backups, License Management, Shutdown/Restart, and Updating Software for all FX Series Appliances. This section also describes how to obtain FX Series support, documentation, and downloads for the FX Series Appliances.

Symbols used in this manual:  Important Note  Informational Note

Contacting Product Support

Go To:

http://www.comtechefdata.com/support

Comtech EF Data Product Support representatives for FX Series Products are available.

Comtech EF Data offers an annual subscription plan providing unlimited telephone support for the coverage period, software upgrades and other important support provisions. Contact Product Support for more information.

Key FX Series Appliance Information

This information should be recorded and saved for future reference for each FX Series Appliance. It should be updated for any upgrades or changes. Providing this information to Support will assist the support team in resolving issues and questions more quickly.

Comtech Serial Number

This can be found on the outside of the appliance.

The following information can be found on the Web GUI Web GUI (bottom left on all screens).

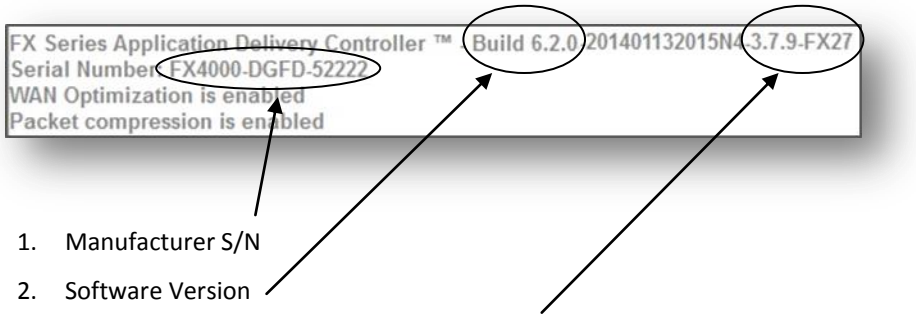

1. Manufacturer S/N
2. Software Version
3. Base Platform Image and Service Pack Version

FX Series End User License Agreement

This is a legal agreement between you (either an individual or an entity) and Comtech EF Data Corporation.

HARDWARE LICENSE and WARRANTY
This product is covered by Comtech EF Data's standard H/W warranty

SOFTWARE LICENSE
This SOFTWARE is protected by the copyright laws of the United States and international copyright treaties as well as other intellectual property laws and treaties. This SOFTWARE product is licensed not sold.

The FX Series Appliance SOFTWARE you have licensed is defined as the SOFTWARE which operates on an appliance. The FX Series Client SOFTWARE you have licensed is defined as the SOFTWARE which operates on an intelligent, single computer, for use in accessing and accelerating Web, Browser or TCP-based applications.

GRANT OF LICENSE: You have the right to install the FX Series Appliance SOFTWARE on all appliances for which you have licensed copies. For each copy of the FX Series Client SOFTWARE this license confers you have the right to install the SOFTWARE on a designated computer for use in accessing and accelerating Web, Browser or TCP-based applications. The SOFTWARE is in "use" on a computer when it is loaded into temporary memory (i.e. RAM) or installed into permanent memory (e.g., hard disk, CD-ROM, or other storage device) of that computer. You may not install the SOFTWARE on more appliances or on more computers than you have licensed copies.

Additionally, you have the right to make one (1) archival copy of the SOFTWARE for each appliance and for each computer which has the SOFTWARE installed in accordance with the terms of this Agreement and subject to the Use Restrictions as set forth below. The copyright notice, as contained in the original CD-ROM, must be affixed to any archival copy.

COPYRIGHT: The SOFTWARE is owned by Comtech EF Data Corporation or its suppliers and is protected by United States copyright laws and international treaty provisions. Therefore, you must treat the SOFTWARE like any other copyrighted material (e.g., a book or musical recording). You may not copy any of the written materials accompanying the SOFTWARE.

OTHER RESTRICTIONS: You may not rent, lease or sublicense the SOFTWARE, but you may transfer the SOFTWARE and accompanying written materials on a permanent basis provided you retain no copies and the recipient agrees to the terms of this Agreement. You may not modify, create a derivative work, reverse engineer, decompile, or disassemble the SOFTWARE. If the SOFTWARE is an update or has been updated, any transfer must include the most recent update and all prior versions. This license and your right to use the SOFTWARE automatically terminate if you fail to comply with any provision of this license agreement.

SUPPORT AND UPGRADES**:** This Agreement does not entitle Licensee to any support, upgrades, patches, enhancements or fixes for the Product (collectively, "Support"). Licensee must make separate arrangements for Support and pay any fees associated with such Support. Any software upgrades, patches, enhancements or fixes provided as part of Support for the Software that may be made available by Comtech EF Data's Maintenance agreement shall become part of the Software and subject to this Agreement.

LIMITED WARRANTY

LIMITED WARRANTY: Comtech EF Data warrants that (a) the SOFTWARE will perform substantially in accordance with the accompanying written materials for a period of ninety (90) days from the date of receipt provided that it is used on the computer hardware and with the operating system for which it was designed. Any implied warranties on the SOFTWARE are limited to ninety (90) days. These warranties commence on the date you first obtain the product and extends only to you, the original customer. Some states/countries do not allow limitations on duration of implied warranty, so the above limitations may not apply to you.

CUSTOMER REMEDIES: Comtech EF Data's entire liability and your exclusive remedy shall be, at Comtech EF Data's option, either (a) return of the price paid, or (b) repair or replacement of the SOFTWARE that does not meet Comtech EF Data's Limited Warranty and which is returned to Comtech EF Data with a copy of your receipt. IN NO CASE WILL COMTECH EF DATA'S LIABILITY EXCEED THE AMOUNT OF THE LICENSE FEE. This Limited Warranty is void if failure to the SOFTWARE has resulted from accident, abuse, or misapplication. Any replacement SOFTWARE will be warranted for the remainder of the original warranty period or thirty (90) days, whichever is longer. Outside the United States, these remedies are not available without proof of purchase from an authorized non-U.S. source.

NO OTHER WARRANTIES: The warranty and remedies set forth above are exclusive and in lieu of all other, oral or written, expressed or implied. Comtech EF Data disclaims all other warranties, expressed or implied, including, but not limited to, implied warranties or merchantability and fitness for a particular purpose, with regard to the SOFTWARE, and the accompanying written materials. Comtech EF Data does not warrant that the SOFTWARE's functions will meet your requirements or that its operation will be uninterrupted or error free. This limited warranty gives you specific legal rights. You may have others which vary from state/country.

NO LIABILITY FOR CONSEQUENTIAL DAMAGES: In no event shall Comtech EF Data be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or any other pecuniary loss) arising out of the use of or inability to use this Comtech EF Data product, even if Comtech EF Data Inc. has been advised of the possibility of such damages. Because some states/countries do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

EXPORT: You acknowledge that the laws and regulations of the United States restrict the export and re-export of the SOFTWARE. You agree that you will not export or re-export the SOFTWARE in any form without the appropriate United States and foreign government approval.

U.S. GOVERNMENT RESTRICTED RIGHTS

The SOFTWARE and documentation are provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1) (ii) of the Rights in Technical Data and Computer SOFTWARE clause at DFARS 252.227-7013 or subparagraphs (c) (1) and (2) of the Commercial Computer SOFTWARE-Restricted Rights at 48 CFR 52.227-19, as applicable. Manufacturer is Comtech EF Data (Stampede), 80A Rhoads Center Drive, Dayton, Ohio 45458. This Agreement is the entire agreement between you and Comtech EF Data relative to the SOFTWARE and supersedes all prior written statements, proposals or agreements relative to its subject matter. If you acquired this product in the United States, this Agreement is governed by the laws of the State of Ohio. Should you have any questions concerning this Agreement, or if you desire to contact Comtech EF Data, address your questions to: Attention: Contracts Division.

Patents and Trademarks

See all of Comtech EF Data's Patents and Patents Pending at http://patents.comtechefdata.com.
Comtech EF Data acknowledges that all trademarks are the property of the trademark owners.

Webmin is a web-based system administration tool created by Jamie Cameron. All recent versions of Webmin may be freely distributed and modified for commercial and non-commercial use.

Comtech EF Data
2114 West 7th Street
Tempe AZ 85281

WORLD WIDE WEB: http://www.comtechefdata.com

Conventions and References

Metric Conversion

Metric conversion information is located on the inside back cover of this manual. This information is provided to assist the operator in cross-referencing non-Metric to Metric conversions.

Recommended Standard Designations

Recommended Standard (RS) Designations have been superseded by the new designation of the Electronic Industries Association (EIA). References to the old designations may be shown when depicting actual text displayed on the Web Server (HTTP) or Command Line Interface pages for the FX Series appliance).

Trademarks

Product names mentioned in this manual may be trademarks or registered trademarks of their respective companies and are hereby acknowledged.

Environmental

The FX Series Appliance must not be operated in an environment where the unit is exposed to extremes of temperature outside the ambient range 0° to 50°C (32° to 122°F); precipitation, condensation, or humid atmospheres above 95% relative humidity; altitudes (unpressurized) greater than 2000 meters; excessive dust or vibration; flammable gases; or corrosive or explosive atmospheres. Operation in vehicles or other transportable installations which are equipped to provide a stable environment is permitted. If such vehicles do not provide a stable environment, safety of the FX Series appliance may not be guaranteed.

Comtech EF Data Warranty Policy

Comtech EF Data products are warranted against defects in material and workmanship for a specific period from the date of shipment, and this period varies by product. During the warranty period, Comtech EF Data will, at its option, repair or replace products that prove to be defective. Repairs are warranted for the remainder of the original warranty or a 90 day extended warranty, whichever is longer. Contact Comtech EF Data for the warranty period specific to the product purchased.

For equipment under warranty, the owner is responsible for freight to Comtech EF Data and all related customs, taxes, tariffs, insurance, etc. Comtech EF Data is responsible for the freight charges only for return of the equipment from the factory to the owner. Comtech EF Data will return the equipment by the same method (i.e., Air, Express, Surface) as the equipment was sent to Comtech EF Data.

All equipment returned for warranty repair must have a valid RMA number issued prior to return and be marked clearly on the return packaging. Comtech EF Data strongly recommends all equipment be returned in its original packaging.

Comtech EF Data Corporation's obligations under this warranty are limited to repair or replacement of failed parts, and the return shipment to the buyer of the repaired or replaced parts.

Limitations of Warranty

The warranty does not apply to any part of a product that has been installed, altered, repaired, or misused in any way that, in the opinion of Comtech EF Data Corporation, would affect the reliability or detracts from the performance of any part of the product, or is damaged as the result of use in a way or with equipment that had not been previously approved by Comtech EF Data Corporation.

The warranty does not apply to any product or parts thereof where the serial number or the serial number of any of its parts has been altered, defaced, or removed.

The warranty does not cover damage or loss incurred in transportation of the product.

The warranty does not cover replacement or repair necessitated by loss or damage from any cause beyond the control of Comtech EF Data Corporation, such as lightning or other natural and weather related events or wartime environments.

The warranty does not cover any labor involved in the removal and or reinstallation of warranted equipment or parts on site, or any labor required to diagnose the necessity for repair or replacement.

The warranty excludes any responsibility by Comtech EF Data Corporation for incidental or consequential damages arising from the use of the equipment or products, or for any inability to use them either separate from or in combination with any other equipment or products. A fixed charge established for each product will be imposed for all equipment returned for warranty repair where Comtech EF Data Corporation cannot identify the cause of the reported failure.

Exclusive Remedies

Comtech EF Data Corporation's warranty, as stated is in lieu of all other warranties, expressed, implied, or statutory, including those of merchantability and fitness for a particular purpose. The buyer shall pass on to any purchaser, lessee, or other user of Comtech EF Data Corporation's products, the aforementioned warranty, and shall indemnify and hold harmless Comtech EF Data Corporation from any claims or liability of such purchaser, lessee, or user based upon allegations that the buyer, its agents, or employees have made additional warranties or representations as to product preference or use.

The remedies provided herein are the buyer's sole and exclusive remedies. Comtech EF Data shall not be liable for any direct, indirect, special, incidental, or consequential damages, whether based on contract, tort, or any other legal theory.

RMA Policy

To return a Comtech EF Data product (in-warranty and out-of-warranty) for repair or replacement, please follow these guidelines.

Contact the Comtech EF Data Customer Support Department during normal business hours. Be prepared to supply the Customer Support representative with the model number, serial number, and a description of the problem. Request a Return Material Authorization (RMA) number from the Comtech EF Data Customer Support representative.

Pack the product in its original shipping carton/packaging to ensure that the product is not damaged during shipping.

Ship the product back to Comtech EF Data. (Shipping charges should be prepaid.)

Online RMA Support

An RMA number can be requested electronically by accessing Comtech EF Data's online Support page (www.comtechefdata.com/support.asp). From this page:

Click the Service hyperlink, and then read the Return Material Authorization section for detailed instructions on Comtech EF Data's return procedures.

Click [Send RMA Request] on the Support page or the RMA Request hyperlink provided in the Service | Return Material Authorization section; fill out the *Billing Information*, *Return Information*, and *Unit to be Returned* sections completely, then click [Send email]
Or –
Send an e-mail providing this same detailed information to the Customer Support Department at service@comtechefdata.com.

Some Stampede products, programs, or services referred to in this publication may not be available in all countries in which Stampede does business. Additionally, some Stampede products, programs, or services may not be available for all operating systems or all product releases. Contact your Comtech EF Data/Stampede representative to be certain the items are available to you.

Release Notes

Version 6.2.2 Functionality Enhancements

These new features have been developed in response to customer feedback and market analysis for the purpose of increasing the interoperability of the FX Series with other CEFD products, to ease moving between screens and to broaden the scope of environments where FX Series can be deployed. This release adds the following new features:

- The FX Series Administrative User Interface has new pull-down navigation structure instead of menu icons.

  This allows enhanced and reorganized tabs for quick and easy direct access to all sub menus. This Header with pull downs allows direct access and is positioned at the top of all Web GUI screens.

- The dynamic ACM screen is gone and is now part of the 'Link' screen
- Multi-Level QoS supports multiple modems with ACM. There are now 3 levels of QoS instead of 1
- ACM can now be configured to support non-Comtech modems
- Point-to-multipoint packet compression is supported
- SNMP can now be configured from the Administrative User Interface
- Ability to monitor and configure active-flows has been added
- Network status screen now clearly illustrates connectivity problems
- QoS Queue and Link Monitor screens now show up to 30 days of history
- Operations->Backup-Restore now allows you to restore just optimization setting without affecting management settings
- Improvements to source IP preservation in NAT environments
- Ability to configure access control on HTTP URLs in L7 application policies
- The default for the Administrative Web GUI is set to run over HTTP. (SSL default is now disabled). **Note:** If version 6.1.1 was installed, then SSL will still be set to “Enabled” as the default to use when logging in to the Administrative Web GUI.

Version 6.1.1 Functionality Enhancements

Package Release 6.1 new features were developed in response to customer feedback and market analysis for the purpose of increasing the interoperability of the FX Series with other CEFD products and to broaden the scope of environments where FX Series can be deployed.

This release adds the following new features:

- Quality of Service
  - QoS only license is now rate limited at 700 Mbps instead of 500 Mbps
  - New protocol filter options for SCTP, PTPv1, PTPv3
  - Added support for VLAN priority in the QoS filters

- FAST Codes

  This release introduces new “trial license” Fast Codes for 30/60/90 day for Packet Compression and WANOP.

- Reporting

  The Status->View Current Status->ACM QOS->By VSAT Modem has two changes.
  New column header for *'Queue Name'* indicates which queue a filter is directed
  'Filter Hits' column header has been changed to 'Filter Matches'

Version 6.1 Functionality Enhancements

These new features have been developed in response to customer feedback and market analysis for the purpose of increasing the interoperability of the FX Series with other CEFD products and to broaden the scope of environments where FX Series can be deployed.

This release adds the following new features:

- Enhanced QOS Monitor Functionality

  CurrentStatus-> ACM QOS-By VSAT modem.
  This now shows ingress packets and bytes which will be non-zero if packet compression is happening.
  CurrentStatus->ACM QOS-Throughput by QOS Queue.
  This now shows packet compression savings percentage.
  Configuration->QOS-Queues pick list.
  This function has been dramatically reworked. You can now change CIR, MIR, Priority, and enable/disable packet compression directly from the view.

- Enhanced SNMP Functionality

  MIB
  Now has MIB which will allow full management of the FX.
  Same MIB is used for both FX-Remote and ADC.
  New wramp SNMP configuration wizard
  Now prompts for destination 'trap' community and 'read/write' community. (Previous MIB was not read/write and did not emit traps (traps are an SNMP term for alerts)

- Enhanced Operations Functionality

  Operations->Shutdown Restart.
  Now has new 'Restart acceleration service and reset cache. This is now the only way to completely reset the cache files

- Status Monitor Enhancements

  Real-Time Monitor
  Real-time monitor now does a 'quick' reset of cache that does not require reboot.
  CurrentStatus->NetworkStatus-Of WAN Interface.
  This is a new status feature that is the only way to ascertain the MAC address of the WAN interface.

- Header Compression/Packet Aggregation

  FX aggregates packets into an Ethernet frame and sends it to a peer, where the packets are restored.
- The default for DDS has been changed to 'Enabled' on the ADC. Previously it was 'Disabled' by default.

Version 6.0.3 Functionality Enhancements

These features have been developed in response to customer feedback and market analysis for the purpose of increasing the interoperability of the FX Series with other CEFD products and to broaden the scope of environments where FX Series can be deployed.

This release adds the following new features:

- Multiprotocol Label Switching (MPLS) is now supported in ACM Filter Definitions.

  If MPLS is selected, then the "MPLS Label" and "MPLS experimental bits" fields will be enabled as filter criteria.

- The default for Dynamic ACM Polling Method Parameters is changed.

  The default setting is now the Modem type, with the pull-down choices including:
  CDM-750, CDM-625, CDM-760, CDM-800, CDM-840, and CTOG-250.
  The default is the CDM-750

- L5 functionality has been enhanced with the following improvements

  Pre-connect option has been removed from the L5 form
  Enable acceleration has been added to the L5 form

  It now includes the ability to define a "*" policy for L5. A "*" is a port range of 1-65535.

- Other changes include:

  VLAN Mode has been added to the general screen for (Trunk or Access)
  Fail-to-Wire option has been added to the general screen (on or off)
  ACM QOS Section of the Current Status screen has an added report "Throughput by QoS with an updated description of "By Modem"

Version 6.0.2 Functionality Enhancements

These features have been developed in response to customer feedback and market analysis for the purpose of increasing the interoperability of the FX Series with other CEFD products and to broaden the scope of environments where FX Series can be deployed.

This release adds the following new features:

- Mesh Network Configuration

  Mesh network optimization is now supported with two appliances at each site.

- Multicator modifications

  Configuration settings have been simplified and located on one main screen on the Web Admin Guide. The Multicator icon will appear on a single screen if an in-path interface is enabled.

- QOS modification for FTP

  FTP is now an option on the QOS filter screen. If FTP is selected, the FX automatically tracks the data ports associated with FTP transfers by monitoring the activity on the FTP control port, which is defaulted to port 21 upon initial selection. The FTP control port may be changed.

- WCCP is now enabled on FX Series Remote

  WCCP functionality for the FX Series Remote is now available and follows the configuration/installation patterns which have been available on the FX Series ADC.

- Redundancy modifications

  The process for setting up these options has been simplified.

Version 6.0.1 Functionality Enhancements

This release added the following features:

Management Port

- This release supports a dedicated management port. The administrative WEB GUI has been enhanced to configure management port settings. Management traffic flows over a separate routing table from the accelerated data traffic. The Administrative Web GUI can now optionally run over HTTP/S.

Trunked VLAN Support

- FX Series supports a trunked network, where multiple 802.1Q tagged VLANs flow thru the same physical connection. To accomplish this, many aspects of the FX Series Remote FX Series ADC data interception and acceleration was modified to retain the VLAN properties. Any accelerated data is transmitted over the network on the same VLAN as the original, non-accelerated data.

. FX Series Release 6.0.1 provides:

- Support for 1024 active VLANs for IDs 2-4095.
- Support Virtual Routing and Forwarding (VRF) environments.
- Support display of tallies on a per-VLAN basis at ADC only (not Remote).
- Accelerated VLAN traffic will maintain original VLAN affinity.
- Private HTTP caches on a per VLAN basis
- Cached HTTP data will be segregated between VLANs.
- VLAN addition and deletion configuration changes can be made without service loss or downtime. A restart is not required for the changes to take effect.

Transparency

FX Series Release 6.0.1 provided:

- Ability to communicate between appliances using the same port as the original client connection.
- Ability to communicate between appliances using the original client source addresses.
- Ability to support active-active ADC configurations.
- Ability to optionally disable multiplexing of client connections.

Dynamic ACM QoS

- Special support was added to FX Series Release 6.0.1 ADC to continuously acquire the data rate of a modem via SNMP connection. When the data rate changes the QoS rules are dynamically adjusted. New fields were added to the Dynamic ACM page to configure the IP address of the modem, and user name and password.

FaST Code Support

- In prior releases, a “license” file was uploaded to the FX to enable functionality.
- In FX Series Release 6.0.1 and above, this methodology is now superseded by FAST Codes.
- The FX Series CLI and Administrative Web GUI have been enhanced to allow Fast code upgrades.

Routed Mode Deployment Option

- The main configuration screen now allows you to put the FX in either “bridged” or “routed” mode. In routed mode, policy based routing (PBR) must be set up on the Cisco router to specifically direct traffic to the FX Series Appliances.

Reliable Multicast Fan-Out

- “Multicator” feature is a powerful new content distribution system. This feature allows a user to upload a file to an FX device via ftp, the file is then reliably multicast to a group of receivers. The receivers then upload the content to a local ftp server. The Multicator employs the “Content Distribution Control Protocol” (CDCP) to ensure that only one multicast transmission is in progress.

Base Platform Image ‘3’ Upgrade Kit

- The new features of Release 6 require new software packages and a new kernel from previous FX releases. This upgrade kit will define procedures for updating existing FX appliances from a USB flash drive.

Management via SSH

- The Base Platform Image “3” provides support over SSH and will also allow the Administrative WEB GUI to function over SSL.

1 Overview - FX Series

1.1 Stampede FX Series Product Line Update

Value Proposition

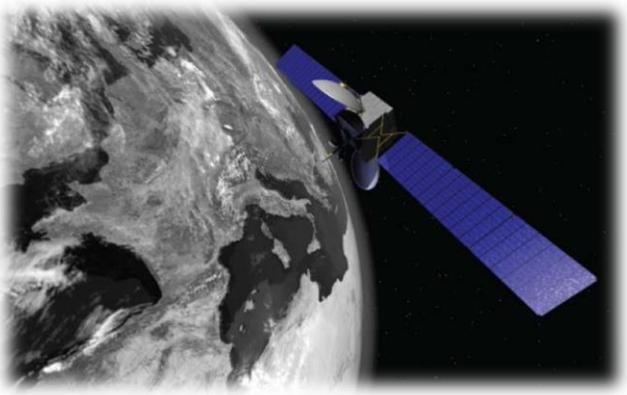

"Reduce OPEX, Improve User Experience"

Reduce OPEX by:

- Shrinking the Data
- Keeping the Pipe Full

Improve User Experience by:

- Getting the Data there faster
- Getting the Right Data there

The Challenges for ISPs with Satellite Links

Data consumed by individual users and enterprises is increasing exponentially. ISPs must cost-effectively keep up with the enormous demand for limited bandwidth - while conserving it.

Assuring Delivery of Web Applications for Bottom Line Results

Data center simplification and the growing migration to web-enabled applications are driving the need for a new class of multi-function optimization devices. The Stampede FX Series combines both one-sided application delivery and two-sided WAN optimization into a single platform. The FX Series delivers unprecedented application performance, optimization, transparency, availability and management for existing networks.

Typical Users

- Internet Service Providers (ISPs)
- Enterprise
- Offshore/Maritime
- Telecommunications Operators
- Satellite Operators
- Managed Service Providers

Common Applications

- High-speed content delivery
- HTTP and TCP optimization & acceleration
- Corporate networks
- Mobile Backhaul

Header Compression/Packet Aggregation

As real-time traffic moves to IP, there is a proliferation of traffic with small payloads. In this case, the header bytes can be 2 to 4 times the number of payload bytes. For small voice packets, compression can result in reducing the required data rate to 30 – 50% of the original. The FX will compress headers, and optionally compress payloads. The FX aggregates compressed packets into an Ethernet frame and sends it to a peer, where the packets are restored. Header compression is integrated into the traffic shaping, and maximum latency per queue can be enforced when aggregating packets

Key Benefits

- Provides up to 80% bandwidth savings in both directions
- Provides up to N times efficiency when using the Multicator
- Enables measurable reduction in response time for users
- Delivers CAPEX for OPEX payback typically in 3-4 months
- Scales easily for small, medium and high volume networks
- Ensures the best traffic flow with Advanced Traffic Shaping
- Matches the modem link rates with ACM tracking
- Real time voice sessions with the use of Header Compresses/Packet Aggregation.

1.2 Technologies that Optimize Satellite Bandwidth Acceleration

Traffic Shaping with ACM Tracking

Traffic is classified and prioritized at layers 2-5. With three levels of filters and the ability to shape with CIR's and MIR's at each level, traffic can be managed across multiple geographic locations using point to point and point to multipoint links. The traffic shaper supports links with ACM by reading the data rate from the modem and adjusting to that rate. There are presets for EF Data modems that support ACM. This feature is available as either stand-alone, or as part of the full WAN optimization product

Transparent Assured Delivery

With flexible options for in-line or Cisco's Web Cache Communication Protocol (WCCP), the FX Series devices deliver unprecedented transparent optimization. End-to-end assurance is maintained for all applications providing complete transparency and the ability for existing Quality of Service (QoS) and network visibility management programs to continue monitoring the health of your network.

Optimize VLAN Trunked Data

All appropriate Layer 5 and Layer 7 optimizations are available for tagged VLAN data, preserving or recreating the VLAN tags for optimized traffic. This includes HTTP caching as well as de-duplication. Caches are maintained by appliance and by VLAN. Appropriate traffic can be shared between VLANs on the same appliance. In addition, the FX-1010 will support up to 8 LAN ports, each of which is tagged and passed to the WAN trunk.

Multicator

The FX Series supports a reliable multicast. This is designed to work in a mesh network, but will also work in a hub/spoke network. In the mesh, any device can be a transmitter with the remaining devices being receivers. Multiple devices can be transmitters. The transmitter function is time shared, with a second device being given permission to transmit after the first is complete. This can work in a hub-spoke network where typically the ADC would be the transmitter, although this is not required. The process is to FTP a file from the client into the transmitter's inbox, that file is transmitted reliably in a multicast to all of the receivers. Once transmitted, the receivers FTP the file to a specified server.

Redundancy and Fail Over

Redundancy is critical to 24/7 availability, and the FX appliance is designed to handle redundancy and fail over in two different ways; inline and routed. The inline configuration is used when operating in conjunction with a CEFD modem operating with 1:1 redundancy. WCCP (Web Caching Communication Protocol) is used in routed mode to allow N devices to serve the function of any M devices, resulting in M: N redundancy. The inline configuration has a primary and a redundant device in series, the redundant takes over whenever the primary fails.

Management

The FX platforms provide total insight through real-time information including over 100 real-time statistics providing extensive details on all inbound and outbound traffic. Historical data for days or months are easily viewed via online graphs, simplifying capacity planning, trending, network issues, and application troubleshooting. Management information can be obtained via an intuitive Web GUI or SNMP. The updating for the FX Series Remotes is automatic. The FX Series remotes poll the FX Series ADC for updates. When the ADC is updated; each remote will download the update and automatically update itself.

Flexibility

The FX Series platforms provide a comprehensive range of flexible options for total transparent 24/7 operation within your existing or growing network infrastructure. No matter what your application acceleration or WAN optimization requirements are today or in the future, the FX Series platform solutions will handle all your business critical applications with ease. Whether your installation requires small, medium or large branches or the consolidation of multiple remote or enterprise data centers, we have the solution for your organization's needs.

Compatible with Advanced VSAT Solutions

The Stampede FX Series products can be added to an Advanced VSAT Solutions network for WAN optimization and application acceleration. The results can be significant improvements in user experience and a reduction by 20-80% in required bandwidth for TCP traffic.

Solutions

Deploy the Stampede FX Series (ADC) as a single-sided solution to optimize traffic from your outbound channel. For a two-sided solution, add the FX Series Remote (REM) appliance and achieve the ultimate in application acceleration and WAN optimization.

Productivity and Performance

The Stampede FX Series WAN optimization improves access to your applications by reducing the amount of data transferred on the link through use of various compression and caching schemes as well as accelerating reliable protocols.

1.3 Single-Sided Solution

1.3.1 Load Balancing via WCCP

The Web Cache Communications Protocol (WCCP) allows satellite network service providers to transparently inject acceleration into their satellite network infrastructure by redirecting traffic flows in real-time to network devices such as the FX Series. WCCP has built-in load balancing, scaling, fault tolerance, and service-assurance (failsafe) mechanisms to ensure network devices can scale and have high-availability. For fault tolerance, if one of the FX Series appliances incurs a hardware failure, the WCCP-enabled router will stop sending traffic to that device and redirect traffic to the other FX Series appliances with zero down-time.

Load balancing via WCCP intelligently distributes the TCP and HTTP workload across multiple FX Series appliances. For flexible scalability, service providers can simply add an FX Series appliance to the cluster, and WCCP will split the traffic load among all the FX Series appliances. Up to thirty-two FX Series appliances can be set up within a cluster and dynamically load balanced.

WCCP enables network service providers to implement the FX Series into their network with greater deployment flexibility, without requiring the FX Series to be physically in-line. The FX Series can be deployed "virtually" in-line, hence, not all traffic is required to pass through the FX Series appliance. The network administrator programs the router to redirect traffic to the FX Service appliance in-bound and out-bound based on the router policies. This allows the administrators to make changes to their network environment by simply changing the router policies.

Stampede's FX Series (running WCCP) localizes content, and responds to content requests in order to reduce the amount of data going over the WAN. This improves application delivery response times, and allows the WAN link to support more traffic. Using WCCP, traffic is transparently redirected to the FX Series appliance for TCP and HTTP acceleration, compression, caching and other optimization services.

With WCCP configured, the router redirects traffic to the FX Series to perform the application acceleration and WAN optimization functions. When an end-user makes a request, the router intercepts the request, and redirects the request to the FX Series inside a generic routing encapsulation (GRE) frame to prevent any modifications to the original packet. The FX Series with WCCP can be used to transparently route traffic, so that you don't have to make changes to Web browsers, and configure the FX Series as a proxy server to offload servers, accelerate application delivery and optimize the network.

1.3.2 Source IP Preservation

Source IP Preservation is a technology that is used to support security policies that require a specific source IP address, or range of IP addresses. It is also used to prevent the FX Series appliance from being blacklisted.

For example, in the event where a situation is deemed inappropriate, such as a SPAM event, the sending device Source IP address will be blacklisted. To avoid this problem, the FX Series uses the end-user's Source IP address when making a request to a Web or application server. The FX Series configuration method makes implementing Source IP Preservation easy within a WCCP or inline environment. The FX Series is usually configured to use the IP address of the client when making requests to content servers, whereas, other FXs make requests to Web servers using their own IP address. IP addressing problems can occur when, for example, an end-user is involved with illegal online activity and the IP address of the FX is recorded in the Web server's logs. If the IP address of the FX is used to make the client request to the server, it will likely be placed on a blacklist, and therefore cause considerable network problems. By spoofing the IP address of the client, the FX Series is able to avoid this problem.

1.3.3 Connection Management

(a) Connection management removes the burden of establishing and terminating TCP connections from the web servers, allowing the server to handle more traffic. Stampede manages network connections in several ways to optimize the flow of data and reduce the impact on the network, application servers and end-user devices. The FX Series appliance maintains a consistent pool of connections between itself and the servers. The servers are then offloaded from managing the connections, and are isolated from inadvertent session disconnects.

(b) The FX Series appliance limits active flows. Active flows are the number of UDP and TCP connections that can be established concurrently between remote clients and content servers which flow through the FX appliances. This should normally never be an issue, but can be a problem in a denial of service attack. The FX series reports the number of active connections in the status monitor.

(c) With Stampede's FX Series Remote appliances working with the FX Series head-end appliance, a persistent connection between the client and server is always maintained, even when the browser may close and reopen a session. These sessions are also multiplexed across multiple connections, improving throughput and response time. This persistent connection is extremely important for AJAX and Web 2.0 applications which constantly open and close sessions as they poll and access various Web services. Stampede eliminates this potentially network intrusive overhead.

1.3.4 ACM QoS

The Quality of Service Function with ACM option is intended to work with modems that support ACM. The FX Series ADC and Remote have the ability to read the current data rate from the modem, and will adjust the output data rate to match the modem data rate. The FX Series data rate is calculated based a per Ethernet frame basis

The FX is also designed to work with the modem in a 1:1 Redundant with fail over mode and work with the modems when they are in a 1:1 redundant configuration.

Output Data Rate

All data rates are Ethernet frame rates. The total data rate is a parameter that can be set, or under the optional ACM mode, can be updated dynamically and continuously from the modem in the link.

Traffic Classification

Traffic can be classified on combinations of Protocol, VLAN, Source/Destination IP Port number, Source/Destination subnet, MPLS labels/EXP and DSCP bits.

Classified traffic is directed into specified Queues. Queues are assigned priority. There are two levels of Groups and a third level of Queues that can be configured. Traffic coming into the appliance is separated by Filters into the level 1 Groups. This traffic can subsequently be separated by Filters into a second level of Groups, and then filtered into Queues where traffic will be released to the WAN based on the QoS and shaping rules defined.

Traffic shaping

Traffic is shaped using drain algorithms on the specified queues. Queues of equal priority are treated in a fair-weighted manner. Connections within a specified Queue are also treated in a fair-weighted manner.

The drain algorithms are strict priority or Min-Max. In Strict Priority, available bandwidth is allocated on the basis of priority.

Min-Max gives more control. Bandwidth is allocated up to a committed information rate based upon priority. Once the committed information rate is reached for all classes, excess bandwidth is allocated based on the same priority, up to a defined maximum for each Queue.

1.3.5 GZIP Compression

The most common use of compression in Web environments is accomplished by enabling GZIP functionality at the Web server. GZIP compression is handled on-the-fly from the servers to the clients. This reduces bandwidth consumption and improves application delivery and client response time. The FX Series uses GZIP compression to reduce the payload size to deliver more data across the satellite link, enabling more applications to be delivered and the ability to support more users. GZIP compression removes non-essential information from data being moved from one location to another, and then reassembles the data to its original form after the transfer is complete.

Squeezing the data reduces network traffic and accelerates the delivery of time-sensitive information. GZIP compression uses standard techniques to compress data sent to browsers. While compression exists in many forms throughout Web deployments, the FX Series is able to more effectively apply compression resulting in better compression ratios. GZIP is not normally used for attachment compression or for inbound compression from the browser. In addition, GZIP cannot be used to compress HTTP headers or image data. In a single-sided mode, the FX Series appliance utilizes GZIP to compress information that can be processed by standard browsers.

Stampede utilizes various compression techniques to reduce the amount of data that must be sent across the network. In two-sided deployment, the FX Series bi-directional compression provides compression for:

- All HTTP Headers
- Application Cookies
- All Text and Data Objects
- JPEG files with Image Reduction, yielding very acceptable quality
- All attachments and file uploads and downloads

1.3.6 Image Reduction and Smoothing

Image Reduction and Smoothing reduces the amount of data required to represent an image without significantly altering the visual perception of the image. This is accomplished in two ways. Smoothing reduces the high frequency components or the sharpness of an image. A moderate amount of smoothing can significantly reduce the amount of data. The quality factor of a JPEG image relates to the precision of the samples. Sample precision can be reduced without visible detection.

The goal of the JPEG quality and smoothing values is to reduce the amount of data while maintaining a usable image. Depending on the JPEG, the compression is often in the range 9:1. A number between 1 and 100 specifies the tradeoff between size of the jpeg data and quality of the original image. A higher number will retain a higher quality but will not conserve as much bandwidth. If no value is specified then the FX Series value is inherited from a higher level policy; a default value of 50 is used if no higher level policy is defined. Images that have been transformed are typically not significantly changed by running through the algorithm again. What this means is that if an image has been compressed with particular smoothing and quality factor, if the same factors are used again, the image is not significantly changed.

1.3.7 Static Caching

Caching brings information closer to the end-user by storing recently accessed data in local memory or on hard disk, reducing the time it takes to bring back needed information, Improving the users' experience by speeding the page load times. While today's browsers maintain their own cache, they tend to be overly conservative. This means they will error on the side of requesting a new piece of data or object, usually when it really hasn't been changed. This not only impacts response time to the end-user, but also saturates bandwidth with unnecessary data transmissions.

The FX Series uses caching to maintain copies of routinely accessed data to eliminate unnecessary requests to Web and application servers, and from going over limited satellite links. By keeping local copies of frequently requested content, the FX Series allows organizations to significantly reduce their upstream bandwidth usage and cost, while improving performance. The FX Series acts as an intermediary from end-users requesting content (such as a file, web page, or other resource) from servers.

Some of the key benefits include:

- Reducing bandwidth consumption
- Keeping servers behind the FX Series anonymous for security purposes
- Delivering fast access to content

1.3.8 TCP Optimization

Advanced protocol optimizations drive significant improvements in bandwidth efficiencies and time savings (reducing payload and latency). WAN optimization and application acceleration technologies are deployed to improve satellite network performance and increase the amount of applications and users that can be delivered over the satellite link. The FX Series manages all TCP sessions, and handles the establishing and tearing down of TCP connections locally (at LAN speeds) to avoid satellite network congestion problems. This helps to increase link utilization and improve the user experience. TCP termination offloads the responsibility from servers having to handle the overhead imposed by the volume of TCP connections from web applications.
Additionally, application level multiplexed TCP streams take advantage of all other TCP or protocol optimization done at the link level, and application-level handshakes are eliminated by consolidating transaction requests.

Benefits include:

- Increases server capacity
- Reduces the amount of traffic sent over satellite links
- Keeps the satellite links maximized for optimum utilization
- Dramatically reduces transaction TCP turns (requests and responses) that bottleneck satellite links

1.4 Two-Sided Solution

1.4.1 Cache Differencing

Cache Differencing takes the concept of caching one step further and maintains identical copies of the browser's cache at the local device and on the FX Series appliance. The FX Series then uses intelligent differencing technology to understand what data has actually changed, and then transfers only the changed data. The local device functions normally, but with less data being transferred, you realize improved utilization of the satellite network, and increased end-user productivity.

Traditionally, pages can be marked as cacheable and will have expiration dates. When they expire they must be retrieved from the original server, resulting in additional traffic and data being transmitted across the satellite network. Within a two-sided environment, the FX Series Remote appliance caches all pages returned to the browser (even pages that are marked as non-cacheable) and performs validation when needed to ensure that no stale data is returned to the browser. When the browser asks for a page or an item that has expired or been marked as non-cacheable, the FX Series remote appliance sends a validation request to the FX Series appliance at the head-end. If the FX Series appliance is aware of the last page the client cache contains and can compute differences in the page, it sends just the differences to an expired page or non-cached page. If the differences are too big, or if the FX Series appliance no longer has retained the last version that the client has, then the entire page is returned and subsequently cached for future possible differencing. The client in turn reconstructs the requested page, caches it, and returns it to the browser. Checksums are calculated by the FX Series appliance at the head-end and verified at the FX Series remote appliance so that pages will never be delivered incorrectly. While this technique adds value on expired pages, it is extremely effective for dynamic page generation.

An important aspect of Stampede's Cache Differencing is the ability to perform differencing not only on HTML GET requests but also on POST requests. This is significant because a) responses to posts are always marked non-cacheable, and b) most applications that are based on SOAP and XML (including most AJAX applications) issue SOAP requests via the HTML POST command.

1.4.2 Multiplexing of Large Data Objects

The FX Series multiplexes large data objects using Comtech EF Data's patented TurboStreaming™ (multiplexed TCP sessions, patent # 7,543,072) that enables HTTP browser traffic to be intermixed across multiple "pipelines". All browser activity is optimized, including the network-intensive polling associated with Web 2.0 and AJAX applications. A key advantage of TurboStreaming is that communication resources can be shared across multiple applications, and all HTTP requests and responses from any application (including multiple browsers) are intermixed simultaneously across multiple concurrent sessions.

TurboStreaming serves as a platform for the consolidation and aggregation of all Web-based traffic from a given user. Multiple HTTP protocol streams are logically aggregated across a few TCP sessions. Individual objects or pieces of objects can be split into any size and then multiplexed with other object data and reconstructed as needed SNSPs that deliver mixed payloads consisting of business-critical applications and data, streaming media, and other network-intensive traffic. The end result is improved throughput and faster response time for the end-user.

TurboStreaming enables the browser to open multiple pipelines (10s or even 100s) that communicate with the FX Series remote appliances. All of this data, from all browsers and all browser windows, is intelligently multiplexed over multiple TCP sessions back to the head-end FX Series appliance. This fully utilizes all available bandwidth, and enables the browser to function at its full potential. This is only possible because of advanced, industry leading two-sided acceleration technology.

1.4.3 Partial Content Update Caching

Intelligently caches Microsoft® updates and other prevalent software updates on the client side saving significant bandwidth attributed to "Patch Tuesday". The FX Series caching methodology handles the rather complicated procedures employed by Microsoft and other AV vendors to request updates by requesting "partial objects". This reduces the amount of data sent over satellite links to reduce bandwidth consumption and provide faster response times for end-users.

The FX Series Remote can dramatically curb bandwidth consumption by caching software updates published frequently by Microsoft, Symantec, Adobe, Apple and many other leading software vendors. The delivery of these updates is performed when software that resides on client devices downloads the new content in the background by requesting "partial content" over HTTP. The complex nature of "partial-content" HTTP requests thwarts the capabilities of most caching devices, however the FX Series Remote appliance caching engine can handle these requests. Once the content is cached by the FX Series Remote, subsequent retrievals by the updating agents that request "partial-content" will be satisfied by the FX Series Remote appliance, eliminating the need to repetitively transfer the same updates over satellite links.

1.4.4 Network Protocol Optimization

The FX Series provides application-aware modules for HTTP, CIFS, MAPI, POP3, SMTP, and FTP that dramatically reduce costly handshakes and intelligently apply compression to lower bandwidth consumption and reduce latency.

Stampede specializes in optimizing protocols by consolidating multiple transactions into a single transaction, which eliminates round-trips, performing cache differencing on dynamically generated content, and bi-directional data compression. In addition, our patented technology (TurboStreaming) enables the transfer of previously compressed objects up to 5 times faster through intelligent multiplexing across multiple TCP sessions.

- TCP and HTTP applications have chatty protocols that put added delay in satellite networks, as do delay-sensitive such as Microsoft Exchange and CIFS.
- IT managers are placing thousands of applications on their satellite links. Many of these applications are mission-critical, and compete over a limited amount of bandwidth.

1.4.5 Dynamic Data Deduplication

Dynamic Data Deduplication segments the incoming data stream, uniquely identifies the data segments, and then compares the segments replacing repetitive streams of payload data with signatures prior to transmission over the satellite links. This feature is not application protocol specific and can be applied to most TCP application traffic. The FX Series intelligently monitors the data stream and is able to distinguish protocol headers which change frequently from payload data which is often static. The FX Series extracts this payload data and segments it into blocks, storing each block into persistent memory known as a "byte cache". Blocks of data are replaced with a signature for that data. This generates significant data reduction.

1.4.6 Header Compression/Packet Aggregation

As real time traffic moves to IP, there is a proliferation of traffic with small payloads. In this case, the header bytes can be 2 to 4 times the number of payload bytes. For small voice packets, compression can result in reducing the required data rate to 30 – 50% of the original. The FX aggregates packets into an Ethernet frame and sends it to a peer, where the packets are restored. Header compression is integrated into the traffic shaping, and maximum latency per queue can be set.

1.4.7 Multicator

Theory of Operation

A powerful new content distribution system can now be set up with the “Multicator” feature. This feature allows a user to upload a file to an FX Series device via ftp, the file is then reliably multicast to a group of receivers. The receivers then upload the content to a local ftp server. The Multicator employs the “Content Distribution Control Protocol” (CDCP) to ensure that only one multicast transmission is in progress.

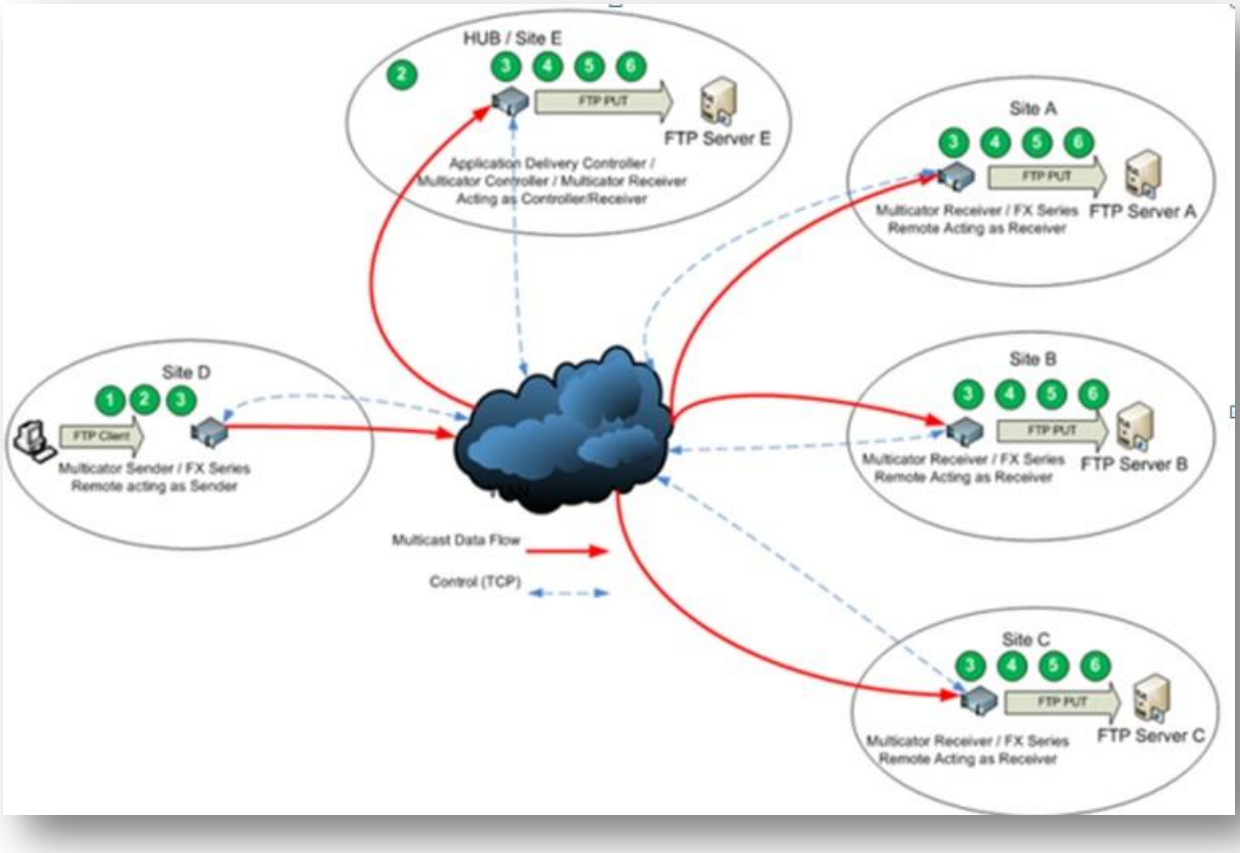

Figure 1-1 FX Series Multicator Theory of Operation

Sequence of Events

1 Files are deposited on the Remote Sender (Site D) using a standard FTP client

2 The Sender then notifies the Controller that it has data to send and is granted permission to reliably multicast the data across the WAN

3 Under control of the Multicator Controller, the Sender establishes a reliable multicast connection to the Receivers.

4 The Sender sends the file to each of the Receivers (Sites A, B, C, and E)

5 Each Receiver verifies receipt to the Controller

6 Each receiver FTPs the file to the respective server.

1.5 FX Series Appliances

1.5.1 Theory of Operation

The FX Series ADC software can run on the FX-4010, the FX-4000, the FX-1005 or the FX-1000. The FX Series ADC applies deflate compression, image transformation, static and dynamic content caching. To the client, the FX Series ADC appears to be the back-end server.

The FX Series Application Delivery Controller (ADC) devices accelerate application delivery and reduce the amount of traffic over satellite links. ADCs are single-sided (asymmetric), requiring an appliance only in the head-end. The FX Series ADC serves as a proxy for TCP management, acceleration and offloading server and network resources for out-bound traffic. TCP acceleration removes the time, quantity and complexity associated with multiple short-lived connections that slow network performance and add overhead to Web server CPU resources. An ADC terminates the client-side TCP session requests, and multiplexes many short-lived sessions into a single longer-lived session between the FX Series ADC and the Web servers.

WANOP Optimization and Data Compression

In addition to a one-sided configuration, the FX Series ADC can reside at the service provider head-end, and work together with FX Series Remote appliances located at each remote site. These products provide two-sided WAN optimization and application acceleration to alleviate the adverse effects that latency and performance errors have upon satellite network performance. They are referred to as WAN Optimization Controllers (WOCs).

In two-sided optimization, if a connection to the FX Series ADC is not able to be achieved by a remote appliance, then the remote appliance will go into a “pass-through” mode where the requests will be directed to the target content server.

Header Compression

As real time traffic moves to IP, there is a proliferation of traffic with small payloads. In this case, the header bytes can be 2 to 4 times the number of payload bytes. For small voice packets, compression can result in reducing the required data rate to 30 – 50% of the original. The FX aggregates packets into an Ethernet frame and sends it to a peer, where the packets are restored.

1.5.2 Reporting

Reports

Important FX Series ADC appliance events are recorded so that the following reports can be viewed:

- **Acceleration Statistics**
  Aggregate Statistics
  By L7 HTTP Policy
  By L5 Application Policy
  Current Connections
- **Throughput Statistics**
  ADC Aggregate Throughput
  Remote Aggregate Throughput
- **Port Statistics**
  By Port Definition
- **WCCP Status**
  By WCCP Definition
- **ACM QoS**
  By Modem
- **Routes**
  By Table
- **Network Status**
  By Interface
  Of WAN Interface
- **HTTP Log Analysis**
  By Month
  Report Download
- **Multicator**
  Multicator Statistics

Monitors

- **WANOP Monitor**
  This Monitor provides a real time view of vital WANOP statistics for both the FX Series ADC and the FX Series Remote.
- **QOS Monitors**
  The QoS Monitors provide a real time view of vital QoS statistics, including current, average and elapsed stats. Monitors include one for Links and a second for Queues

1.5.3 Deployment Options

The FX Series ADC can run in single sided mode ADC only, in-line mode, routed mode or in WCCP mode. The installation instructions for these are in the FX Series Installation Patterns Section.

1.6 FX Series Remote Appliance

1.6.1 Theory of Operation

The FX Series Remote software can run on the FX4010, the FX-4000, the FX-1005, the FX-1010 and the FX-1000. The FX Series Remote accelerates traffic by intercepting user requests and forwarding them to the FX Series ADC. The FX Series ADC applies deflate compression, image transformation, static and dynamic content caching. The FX Series Remote applies static content caching, dynamic content caching, deflate compression, Dynamic Data De-duplication, persistent connections, connection multiplexing, client side connection termination, and TurboStreaming. To the client, the FX Series Remote appears to be the back-end server. When in a two-way configuration the FX Series Remote will communicate with the FX Series ADC via the port that the client is connecting by default. If the FX Series Remote is configured to connect to a specific FX Series ADC then port 4922 will be used. If a connection to the FX Series ADC is not able to be achieved then the remote appliance will go into a “pass-through” mode where the requests will be directed to the target content server.

Most FX Series Remote configuration is accomplished with an easy-to-use browser-based tool to set polices on the FX Series ADC appliance. The configuration policies are designed to provide full inheritance properties, meaning that most configuration settings are shared between all FX Series Remote appliances, but individual over-rides can be set for specific FX Series Remote appliances. Examples of policy-based settings include:

- Bandwidth reservation and prioritization
- HTTP application optimization
- Compression and caching settings for HTTP, CIFS, POP3, SMTP, and FTP

Header Compression

As real time traffic moves to IP, there is a proliferation of traffic with small payloads. In this case, the header bytes can be 2 to 4 times the number of payload bytes. For small voice packets, compression can result in reducing the required data rate to 30 – 50% of the original. The FX aggregates packets into an Ethernet frame and sends it to a peer, where the packets are restored.

Wanop Optimization and Data Compression

All TCP traffic between the FX Series Remote is compressed using intelligent data dictionaries to ensure that repeated patterns are eliminated from subsequent accesses. Several techniques are utilized to guarantee that the TCP communications between the FX Series Remote and the FX Series head-end appliance are fully optimized, including:

RFC3649

"High-speed TCP for Large Congestion Windows"

Streaming

Moves data streams over multiple concurrent TCP connections between FX Series Remote appliances and FX Series head-end appliance. This insulates the FX Series from intermittent packet loss, as data is almost always going at full speed over at least one of the connections.

HTTP Optimization

The optimization techniques of FX Series client acceleration are built into the FX Series Remote appliance, resulting in highly optimized delivery of HTTP applications to remote site users without having to deploy software on individual computers. Some of the optimizations that FX Series Remote appliance can apply to HTTP applications include:

- Caching of static objects, Cache differencing of dynamic content and Cookie Compression

1.6.2 Report*ing*

Reports

Important FX Seri*e*s Remote appliance events are consolidated at the FX Series ADC appliance. These events are recorded so that the following consolidated reports can be viewed on the Remote Appliance:

- **Acceleration Statistics**
  Aggregate Statistics
  Current Connections
- **Throughput Statistics**
  Aggregate Throughput
- **WCCP Status**
  By WCCP Definition
- **ACM QoS**
  By Modem
- **Routes**
  By Table
- **Network Status**
  By Interface
  Of WAN Interface
- **Multicator**
  Multicator Statistics

Monitors

- **WANOP Monitor**
  This Monitor provides a real time view of vital WANOP statistics for both the FX Series ADC and the FX Series Remote.
- **QOS Monitors**
  The QoS Monitors provide a real time view of vital QoS statistics, including current, average and elapsed stats. Monitors include one for Links and second for Queues.

1.6.3 Deployment Options

The FX Series appliances can run in in-path mode, in routed mode or in WCCP mode.

The installation instructions for these are in the FX Series Installation Patterns Section.

1.7 Mesh Networking with the FX Series

1.7.1 Theory of Operation

In addition to the single sided and the two sided client/server or Hub/Remote star network, we've now introduced a full mesh network. We accelerate traffic from the FX Series Remote to the FX Series ADC, with both appliances at each site.

NOTE: The FX Series Mesh can utilize the FX1005 appliances in a dual rack installation.

The FX Series Remote accelerates traffic by intercepting user requests and forwarding them to the FX Series ADC. The FX Series ADC applies deflate compression, image transformation, static and dynamic content caching.

The FX Series Remote applies static content caching, dynamic content caching, deflate compression, Dynamic Data De-duplication, persistent connections, connection multiplexing, client side connection termination, and TurboStreaming. To the client, the FX Series Remote appears to be the back-end server.

1.7.2 Mesh Capability with two FX Series appliances at each node

- All optimizations are handled – Remote to ADC
- Traffic shaping, is done with the FX Remote, not the FX ADC
- The first ADC picks up the traffic and will accelerate/optimize it.

  The configurations for each appliance are done separately and have a cable connected between the Remote LAN port and the ADC WAN port as shown below.

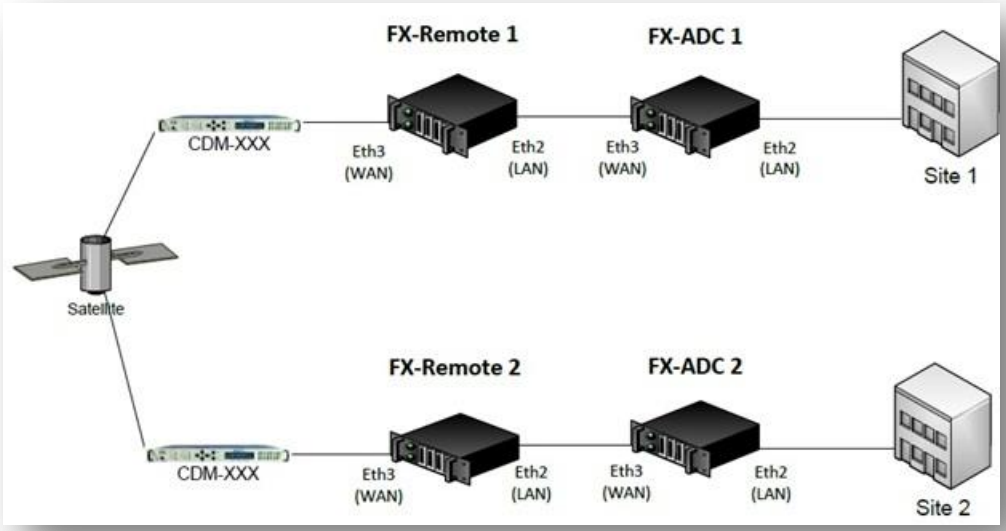

Figure 1-2 FX Series Basic Mesh Connectivity Diagram

- Hub/Spoke with meshing between FX Remotes, with the FX ADC hub available for web browsing and other applications.

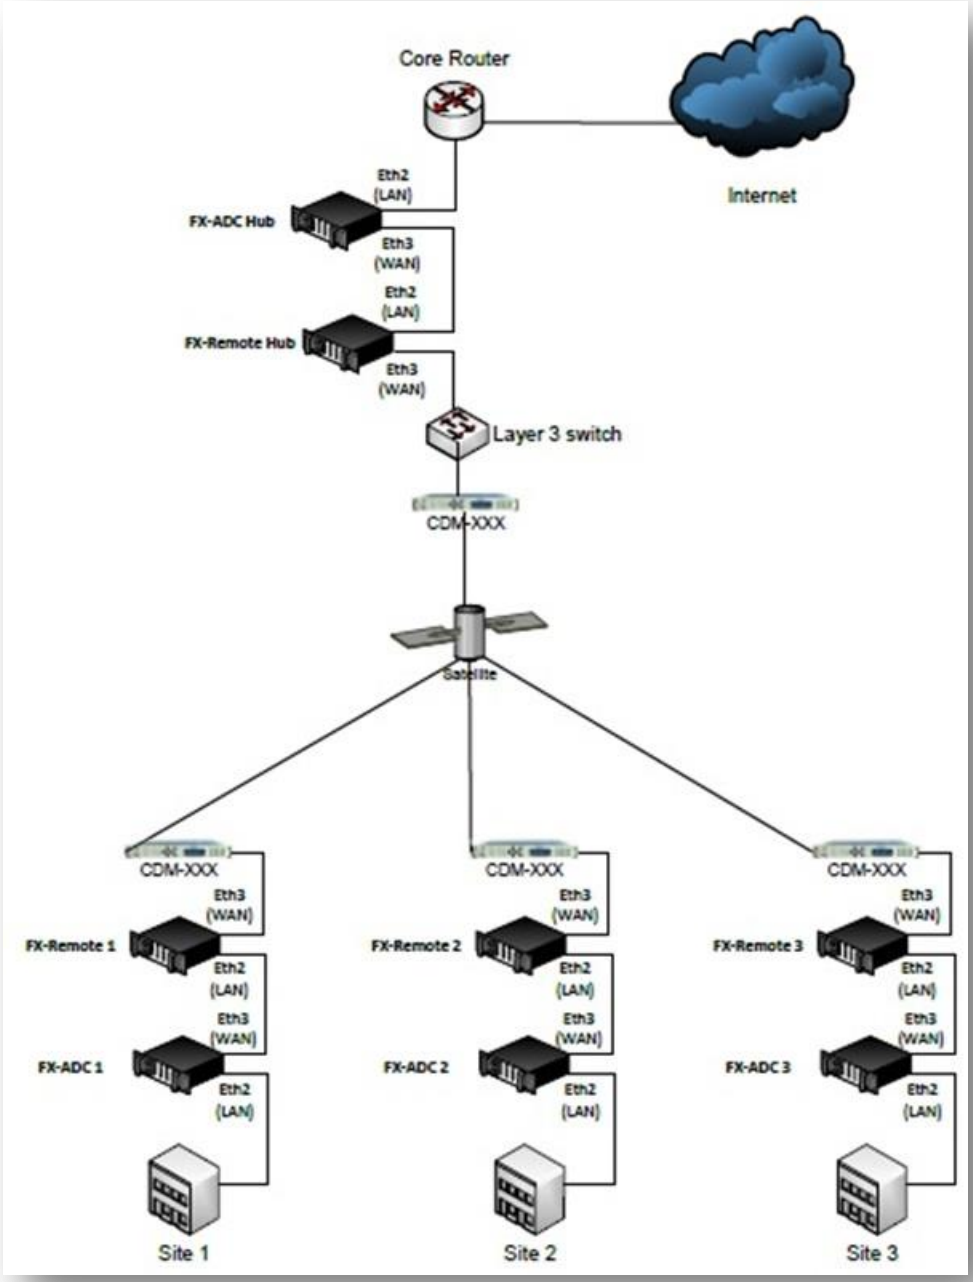

Figure 1-3 FX Series Hub Spoke Mesh Connectivity Diagram

1.7.3 Mesh configuration with Redundancy

The Redundancy configuration could be set up at each site to provide total redundancy

The fail to wire capability is structured between the two like devices and between the Remotes and the ADC as shown below.

The appliances are connected in series as shown below.

See the mesh deployment installation pattern for details.

.

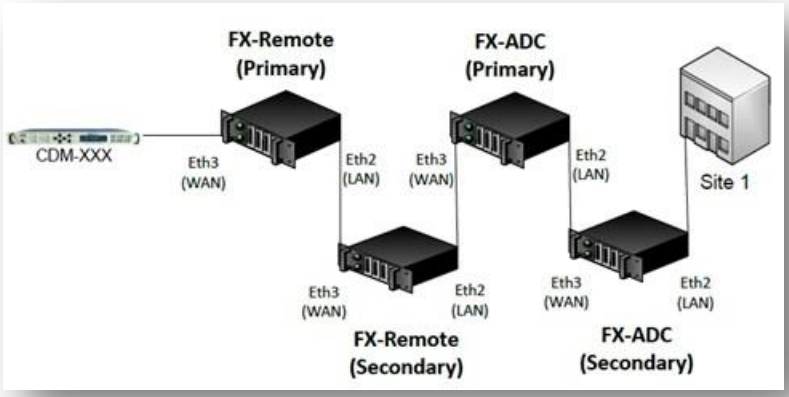

Figure 1-4 FX Series Mesh with Redundancy Connectivity Diagram

1.8 FX Series Appliances Data Sheet

Deploy the Stampede FX Series (ADC) as a single-sided solution to optimize traffic from your outbound channel. For a two-sided solution, add the FX Series Remote (REM) appliance and achieve the ultimate in application acceleration and WAN optimization.

1.8.1 Single Sided with the Application Delivery Controller (ADC)

	FX-1005 ADC	FX-4010-ADC
Max Accelerated Sessions	3,000	30,000
Data Rate Options Mbps	1, 2, 4, 6, 10, 15	10, 15, 25, 45, 70, 155, 310
Load Balancing via WCCP	✓	✓
Connection Management	✓	✓
Advanced Traffic Shaping with ACM (d)	✓	✓
Source IP Preservation	✓	✓
Optimize VLAN Tagged Data	✓	✓
GZIP Compression (b)	✓	✓
Image Reduction (c)	✓	✓
Content Caching		
Static Caching	✓	✓
Redundancy - In-Path and Routed Modes	✓	✓

1.8.2 Two Sided with the ADC and the Remote

	FX-1005 REM/ADC	FX-1010 REM	FX-4010 REM/ADC
Max Accelerated Sessions (a)	6,000 (a)	6,000	30,000 (a)
Data Rate Options Mbps	1, 2, 4, 6, 10, 15	2, 4, 6, 10, 15, 25	10, 15, 25, 45, 70, 155, 310, 700 (f)
Header Compression Rate (PPS) (e)	35,000		700,000
Load Balancing via WCCP	✓		✓
Connection Management	✓	✓	✓
Traffic Shaping with ACM (d)	✓	✓	✓
IP Source Preservation	✓	✓	✓
Optimize VLAN Tagged Data	✓	✓	✓
Multicator	✓	✓	✓
Content Reduction			
Bi-directional Compression	✓	✓	✓
Image Reduction (c)	✓	✓	✓
Dynamic Data De-duplication	✓	✓	✓
Content Caching			
Static Caching	✓	✓	✓
Cache Differencing	✓	✓	✓
TCP Optimization	✓	✓	✓
Multiplexing Data Streams	✓	✓	✓
Auto Updates to the Remotes	✓	✓	✓
See Data Sheet Notes: (a) (b) (c) (d) (e)on next page			

Figure 1-5 FX Series Appliances Data Sheet

Data Sheet Notes:

(a) When used as an ADC, the FX-1005 will handle 3000 concurrent sessions.

(b) Maximum accelerated WAN rates are a function of compressibility. If all content is being GZIP compressed with a ratio of greater than 4:1, the maximum WAN rate may not be accelerated.

(c) The number of images handled per second is a function of image size.

	FX-4010-ADC	FX-1005-ADC
Image Size	Images Per Second	Images Per Second
10 KB	1800	80
50KB	1000	35
500KB	100	35

(d) Available as either a stand-alone feature or part of the WAN optimization product. As a stand-alone feature, the maximum data rate is 700 Mbps, when purchased with the WAN optimization; the data rate is limited to the WAN optimization rate.

(e) Packets per second (PPS) is 50% outbound and 50% inbound. Header compression is currently only available in point-to-point configurations and is not currently supported in the FX-1010. Header Compression is currently available as either a standalone feature added to the base configuration or part of the WAN Optimization product. When purchased without the WAN Optimization feature the maximum rate is 700KBps. When included with WAN Optimization, the data rate is limited to the WAN Optimization rate.

1.8.3 Configuration Models

The Base Configuration with QOS only is an option with no WAN Optimization.

Option 1 Add Header Compression (rates up to700KBps) with no WAN Optimization.

Option 2 Wan Optimization including Header Compression with rates as shown in the tables.

1.8.4 FX Series Hardware Specification

FX Series 1005 **FX Series 1010** **FX Series 4010**

Model	FX-1005	FX-1010	FX-4010
Form Factor	1RU	1RU	1RU
Weight	2.6 lbs. (1.2kg)	13.3 lbs. (6.0 kg)	15 lbs. (6.8 kg)
Dimensions (h x w x d)	1.7" x 8.5" x 7.4" (43 x 215 x 188 mm)	1.7" x 17.0" x 15.6" (44 x 431 x 395 mm)	1.7" x 16.8" x 14.0" (43 x 427 x 356 mm)
Memory	4 GB	4 GB	16 GB
Storage	(1) 160 GB SATA	(1) 160 GB SATA	(1) 1 TB SATA III
Network Interface (GE) Ports/Fail-to-Wire Pairs	4/1	11/0	4/1
Serial Ports	1	1	1
USB Interface Ports	2	2	2
Rack Mount Kits	1 or 2 units in 1RU		
Power Supply – UL Approved, FCC Compliant	Requires a 60 W/12V power adapter with lock	200 W ATX power supply unit with input range of 90~264V@ 47-63 Hz	Single Power (200 W) Auto (100V-240V)
Power Supply Safety/EMC Certifications	EN 61000/IEC 6100-Compliant Australian AS/NZS Class A FCC Part 15 Subpart B Canada ICES-003 Class A Europe/CE Mark ROHS	EN 61000/IEC 61000-Compliant Australian AS/NZS Class A FCC Part 15 Subpart B Canada ICES-003 Class A Europe/CE Mark ROHS	En 60950/IEC 60950-Compliant Canada – CUL Listed Germany –TUV Listed Europe/CE Mark CCC Certified ROHS
Environment	Operating temp 0 - 40°C, Storage temp -20 - 60°C, Humidity 5 - 90%	Operating temp 0 - 40°C, Storage temp -20 - 60°C, Humidity 5 - 90%	Operating temp 10 - 35°C, Storage temp -40 - 70°C, Humidity 8 - 90%

Figure 1-6 FX Series Hardware Specifications

1.8.5 FX-4010 Physical Description

Front Panel

The front panel has the power button, the reset button and 5 LEDs to visually indicate certain vital states of the appliance.

1. Power Button
2. Reset Button

LED indicators from left to right:

1. Power 0n/Off
2. HDD (on - activity/off - no activity)
3. Management Port (on - Linking / Off - not linking)
4. Auxiliary Port (on - Linking / Off - not linking)
5. Temperature Warning

Back Panel

Using suitable RJ-45 cable, you can connect FX Series FX-4010 System to a computer, or to any other piece of equipment that has an Ethernet connection; for example, a hub or a switch. Moreover, LAN (eth 2) / WAN (eth3) are configured as LAN Bypass when failure events occur.

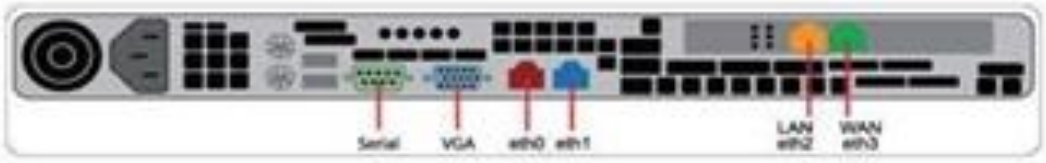

Figure 1-7 FX Series FX-4010 Back Panel

From left to right

1. Power-In Socket
2. Inputs for mouse and keyboard
3. (2) USB 2.0 Ports
4. Serial Port
5. VGA Port
6. MGT/Control (Eth0)
7. AUX Port (Eth1)
8. LAN (Eth2)
9. WAN (Eth3)

LED indicators for MGT (Eth0) and AUX (Eth1) Ports
On/Flashing indicates that the port is linking.
Off indicates that the port is not linking.

LED indicators for LAN and WAN Ports (3) LEDs per port

Link Activity: Turns on any link speed, blinks on activity (green)
100: Turns on Mbit/s link (green).
1000: Turns on Mbit/s link (green).
Bypass: LED 1000 and LED 100 of LAN port 0 are turned on
Disconnect: LED 1000 and LED 100 of WAN port 1 are turned on

1.8.6 FX4010 DC Physical Description

Front Panel

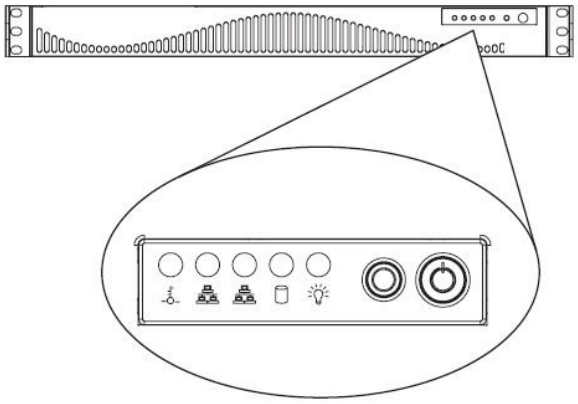

The front panel has the power button, the reset button and 5 LEDs to visually indicate certain vital states of the appliance.
(NOTE: Actual Picture not available at time of printing)

Front panel from right to left.

1. Power Button
2. Reset Button

LED indicators:

3. Power 0n/Off
4. HDD (on - activity/off - no activity)
5. Management Port (on - Linking / Off - not linking)
6. Auxiliary Port (on - Linking / Off - not linking)
7. Temperature Warning

Back Panel

Using suitable RJ-45 cable, you can connect FX Series FX-4010 System to a computer, or to any other piece of equipment that has an Ethernet connection; for example, a hub or a switch. Moreover, LAN (eth 2) / WAN (eth3) are configured as LAN Bypass when failure events occur. (NOTE: Actual Picture not available at time of printing)

From left to right

1. Inputs for mouse and keyboard
2. (2) USB 2.0 Ports
3. Serial Port
4. VGA Port
5. MGT/Control (Eth0)
6. AUX Port (Eth1)
7. LAN (Eth2) – NOT Shown in Picture
8. WAN (Eth3) – Not Shown in Picture
9. Power-In Socket

LED indicators for MGT (Eth0) and AUX (Eth1) Ports

On/Flashing indicates that the port is linking.

Off indicates that the port is not linking.

LED indicators for LAN and WAN Ports (3) LEDs per port

Link Activity: Turns on any link speed, blinks on activity (green)
100: Turns on Mbit/s link (green).
1000: Turns on Mbit/s link (green).
Bypass: LED 1000 and LED 100 of LAN port 0 are turned on
Disconnect: LED 1000 and LED 100 of WAN port 1 are turned on

Operating Environment (System)

Operating Temperature Range	• 5°C ~ 35°C (41°F ~ 95°F)
Non-Operating Temperature Range	• -40°C ~ 60°C (-40°F ~ 140°F)
Operating Relative Humidity Range	• 8% ~ 90% (non-condensing)
Non-Operating Relative Humidity Range	• 5% - 95% (non-condensing)

Power Supply

410W DC-DC power supply (24-pin) with cable harness

DC Voltage	Voltage Range = -36V to -72V Nominal Voltage = -48V Max Input Current = 18A @ -48V
DC Output	5V + 3.3V ≤ 180W
+5V	35.0 Amp
+5V standby	3.0 Amp
+12V	32.0 Amp
-12V	0.5 Amp
+3.3V	20.0 Amp

Regulatory (Power Supply)

Power Supply	USA - UL listed, FCC Canada - CUL listed Germany - TUV Certified
Safety / EMC	Europe/CE Mark EN 60950/IEC 60950-Compliant CCC

1.8.7 FX-1005 Physical Description

Front Panel

Figure 1-8 FX Series FX-1005 Front Panel

Power/Status/HDD LED (left vertical icons)

Power (Green): If the LED is on it indicates the system is powered on. If it is off, it indicates the system is powered off.

Status (Green/Amber): If the LED is Green, it indicates that the system's operational state is normal. If it is Amber, it indicates that the system is malfunctioning.

HDD (Yellow): If the LED blinks, it indicates data access activities; otherwise, it remains off.

LED indicators for Network Ports:

MGT Port (Eth 0) **AUX Port (Eth 1)** **LAN Port (Eth 2)** **WAN Port (Eth 2)**

LED Indicator	Interpretation
SPEED	
Amber	The connection speed is 1000Mbps
Green	The connection speed is 100Mbps
Off	The connection speed is 10Mbps.
LINK/ACT	
On/Flashing (Yellow)	The port is linking.
Off	The port is not linking.

Back Panel

Figure 1-9 FX Series FX-1005 Rear Panel

Reset Switch

Use a pointed object to press the reset button to reboot the system without turning off the power.

Console Port

By using suitable rollover cable (also known as Cisco console cable), you can connect to a computer terminal for diagnostic or configuration purpose

Two USB 2.0 Ports

It connects to any USB devices, for example, a flash drive

4 Gigabit LAN ports

Using suitable RJ-45 cable, you can connect FX Series 1005 System to a computer, or to any other piece of equipment that has an Ethernet connection; for example, a hub or a switch.
Moreover, LAN (Eth2) and LAN (Eth3) are configured as LAN Bypass when failure events occur.

1) MGT - Management (Eth 0)
2) AUX - Auxiliary (Eth 1)
3) LAN - (Eth 2)
4) WAN - (Eth 3)

DC-in 12V Jack

The system requires a 60W/12V power adapter with lock.

Power-on Switch

It is a switch to turn on or off the power.

Summary of Specifications

Network Interface /Fail to Wire

Power Supply – UL Approved

(4) GbE ports, (1) pair bypass

200 W (Auto 100V – 200V)

1.8.8 FX Series FX-1005 Hardware Mounting Options

Tabletop Mounting (Standard)

(a) To mount the FX-1005 on the table, use the rubber feet in the tabletop mounting pack.
(b) Follow the following procedures as a guideline: (may be pre-attached)
(c) Place the rubber feet on the mounting spots at the bottom of the FX-1005 .
(d) Place the FX-1005 on the table using the rubber feet.

Double Unit Rack Mount (Optional Accessory)

To mount two FX-1005 systems onto the rack, use the mounting kit with the screw pack.

(a) Follow the following procedures as a guideline:
(b) Attaching two screws having a washer under the head to the inner side of the system's chassis.
(c) Align the screws of one system with the mounting slots of the other system and mount the two systems side by side by clipping them together
(d) Make sure that the attachment between the two systems is secure and the mounting screws are locked in place.
(e) Use the screws provided to fix the short ear-bracket to the left and right sides of the system as shown in the picture.
(f) Use the mounting hardware included to attach and secure the bracket to the rack.

Installing the ear-bracket to the rear side is an alternative rack mounting

NOTE: The short-ear bracket could also be mounted at the rear side of the system. Thus, the rear panel of the system could be mounted in the front of the rack mounting equipment.

Single Unit Rack Mount (Optional Accessory)

NOTE: Place the power adaptor in the bracket first before installing the adaptor holder.

(a) To mount the FX-1005 onto the rack, use the mounting kit with the screw pack.
(b) Follow these procedures as a guideline:
(c) Attach the adaptor mounting bracket to the system by fastening 5 screws
(d) Place the adaptor in the adaptor mounting bracket.
(e) Make sure that the power adaptor's AC socket is not blocked. Align the AC socket with the holes on the mounting bracket.
(f) You could use the adaptor holder to hold your adaptor to prevent it from sliding back and forth.

(g) Use 3 screws provided to fix the bracket to the left and right side of the system.
(h) Use the mounting hardware included to attach and secure the bracket to the rack.

1.8.9 FX-1010 Physical Description

Front Panel

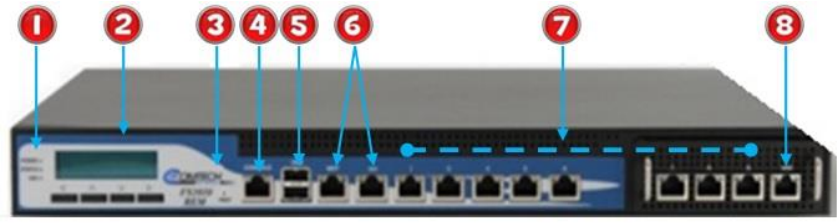

Figure 1-10 FX Series FX-1010 Front Panel

1 Power/Status/HDD LED

Power:
If the LED is on it indicates that the system is powered on. If it is off, it indicates that the system is powered off.

Status:
If the LED is green, it indicates that the system's operational state is normal. If it is red, it indicates that the system is malfunctioning.

HDD:
If the LED is on, it indicates that the system's storage is functional. If the LED blinks, it indicates data access activities. If it is off, it indicates that there is no hard disk present or functional.

2 System Panel: LCD System Panel
The LCD System Panel is programmed to display WOC on the first line and "Active" on the second.

3 Reset Switch:
The reset switch can be used to reboot the system without turning off the power.

4 Console Port:
By using suitable rollover cable or RJ-45 to DB-9 Female (Cisco console cable), you can connect to a computer terminal for diagnostic or configuration purpose. Default terminal Configuration Parameters: 115200 baud, 8 data bits, no parity, 1 stop bit, no flow control.

5 Two USB 2.0 Ports:
It connects to any USB devices, for example, a flash drive.

6 Management Port and Auxiliary Ports:
The Management Port is a Fast Ethernet port that can be connected for configuration or troubleshooting purpose. It conforms to the IPMI (Intelligent Platform Management Interface) and can be implemented on this port through the Open Platform Management Architecture (OPMA) interface.

7 Eight Gigabit LAN ports (Ports 1-8)

8 WAN Port - LAN/WAN Port LEDs

Right LED:
If the LED is orange, it indicates that the connection speed is 1000Mbps. If the LED is green, it indicates that the connection speed is 100Mbps. And if it is off, it indicates that the speed is 10Mbps.

Left LED:
If the LED is on, it indicates that the port is linked. If it blinks, it indicates there is traffic.
Using suitable RJ-45 cable, you can connect FX-1010 system to a computer, or to any other piece of equipment that has an Ethernet connection; for example, a hub or a switch.

Back Panel

Figure 1-11 FX Series FX-1010 Rear Panel

4 System CPU Fans

Power-on Switch

AC Power-in socket -

200W ATX power supply unit with input range of 90~264V@47-63Hz.

Power Supply Fan

Summary of Specifications

Network Interface (11) 10/100/1000,

Power Supply – UL Approved 200 W (Auto 100V – 200V)

Rack Mounting

Rack mounting hardware is included with FX-1010 appliance

2 Initial Installation Information

2.1 Pre-Installation Information

2.1.1 Unpacking

Inspect shipping containers for damage. If shipping containers are damaged, keep them until the contents of the shipment have been carefully inspected and checked for normal operation.
The FX Series appliance is packaged in pre-formed, reusable, cardboard cartons containing foam spacing for maximum shipping protection.

Unpack the appliance as follows:

Step Procedure

1. Remove the appliance, and the power cord and cables from the carton.
2. Save the packing material for storage or reshipment purposes.
3. Inspect the appliance for any possible damage incurred during shipment.
4. Check the equipment and accessories against the packing list to ensure the shipment is correct.

Parts List

Acceleration Appliance
Quick Start Guide
1 - Power Cord
2 - Cat5e 7ft UTP Snagless Cable
1 - Cat5e Crossover Orange/Red 7ft UTP Snagless Cable
1 – Null Modem 6ft Cable

2.1.2 User Interfaces

The FX Series supports a basic menu-driven interface, which is accessible using the console port (eth0) or a web-based graphical user interface (GUI). Initial network configurations are managed thru the console connection, and the optimization and general operations functions are managed via the GUI. There are three alternate methods to connect to the FX Series Appliance

1. Attach a Monitor, keyboard and mouse to device.
2. Connect the supplied serial cable with a setting of (19200,N,8,1)
3. Attach a cross-over cable to the eth1 interface which has a static IP address of 169.254.55.55

(*See FX Series Console Management Functions in Appendix)*

To connect to the Web GUI using a PC with a Browser access: http://yourFxHostname:10000 or http://IP:10000 . The default USERID and PASSWORD are “comtech” and “comtech” for both the console and the GUI.

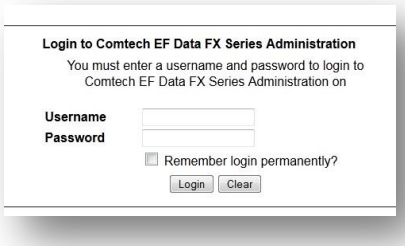

2.1.3 Documentation

The latest FX series Documentation can be found on the Comtech EF Data Web site at: **http://www.comtechefdata.com/support**

2.2 How to Configure Appliance Management Address

2.2.1 All Installation Patterns

Log into the appliance via console or SSH session using the username: comtech and password: comtech

1 Select option 1 “Configure appliance”
2 Select option 1 “Configure network settings”
3 Select option 4 “Configure TCP/IP for eth0 Ethernet Port”
4 Select option 1 “Configure DHCP”
    a. Enter no and press Enter key (disable DHCP for this interface)
5 Select option 2 “Configure IP Address”
    a. Enter the IP address of the appliance and press enter
6 Select option 3 “Configure Netmask”
    a. Enter the subnet mask and press the Enter key
7 Select option 0 “Return to previous menu”
8 Select option 3 “Configure Default Gateway”
    a. Enter the default gateway IP address and press the Enter key

Verify network connectivity by doing a ping of the appliance address from an external device.
Depending on the environment, there are several network installation patterns that can be used. These are each documented in following installation patterns sections.

- Section 2.3 How to configure FX Series Installation Pattern (In-Line Mode)
- Section 2.4 How to configure FX Series Installation Pattern (Routed Mode)
- Section 2.5 How to configure FX Series Installation Pattern (WCCP Mode)
- Section 2.6 How to Configure Two FX Series Appliances in a Mesh Configuration

2.3 How to configure FX Series Installation Pattern (In-Line Mode)

2.3.1 Cable the Appliance

The eth2 (LAN) and eth3 (WAN) ports both need to be connected to a switch or router in which:
The eth2 (LAN) port is on the link closest to the back end servers if the appliance is an FX-ADC or Clients if the appliance is a FX-Remote.
The eth3 (WAN) port is on the link closest to the satellite modem.

FX Series ADC

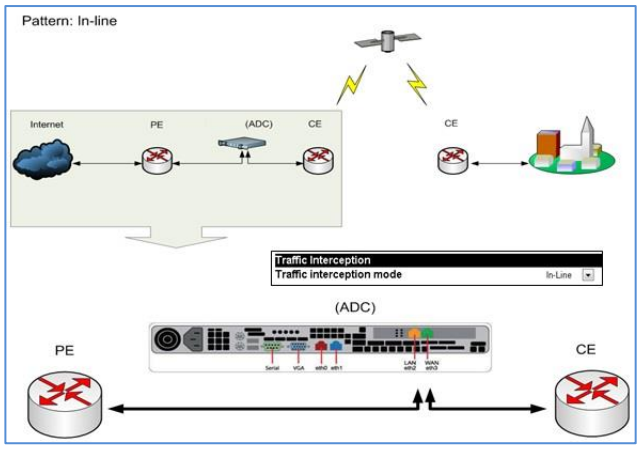

FX Series Remote

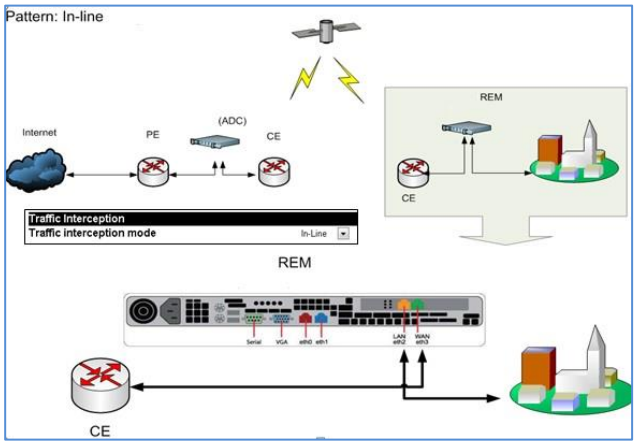

2.3.2 Configure the Appliance

Login to the appliance through the browser interface at:
http://{IP_address_of_the_appliance}:10000

1. Enter the default user name “comtech” and the default password “comtech”.
    a. Click Login.
2. Go to **Configuration -> General Settings**
    a. Change “Traffic interception mode” to “In-Line”.
    b. Click Save.
3. Go to **Traffic Interfaces -> In-Path Interfaces**
    a. Click the Add button to add a new in-path interface.
        - Enter the VLAN ID of 0. If no VLAN tagging is to be used. Use the VLAN ID of the VLAN if traffic is to be VLAN tagged
        - Enter the IP Address.
        - Enter the Netmask.
        - Enter the Gateway.
        - Enter any static routes needed in the “Routes” field.
    b. Click “Add In-Path Interface”.

4. Go to **Traffic Interfaces-> LAN Interfaces.**
    a. Click on the “Add” button to add a new LAN interface.
        - Select “eth2” in the “Physical interface” selection box.
        - In the “In-Path Interface(s)” selection box, click the IP address of the in-path interface *from above*.
        - In the “Untagged in-path interface” selection box, select the IP address of the in-path interface *from above*. Select “None” if the VLAN tag is to be propagated across the WAN.
    b. Click “Add LAN Interface”.

    Go to Operations -> Shutdown and Restart
    Click on the “Restart Service Button”

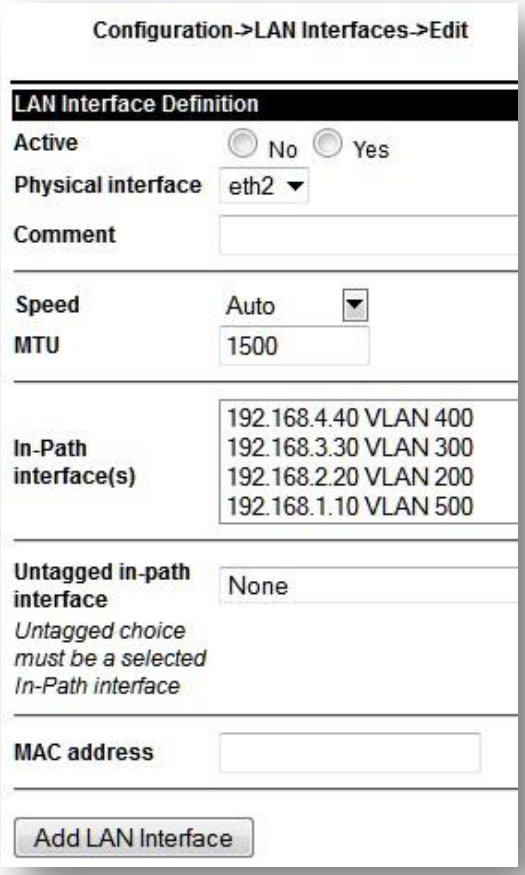

2.4 How to configure FX Series Installation Pattern (Routed Mode)

2.4.1 Cable the Appliance

Connect the eth3 (WAN) port to a switch or a router.

2.4.2 Configure the Appliance

Login to the appliance through the browser interface at: *http://{IP_address_of_the_appliance}:10000*

1. Enter the default user name “comtech” and the default password “comtech”.
    a. Click Login.

FX Series ADC

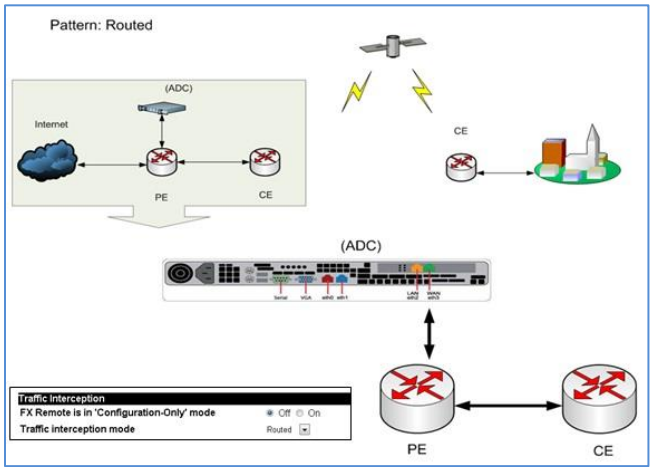

FX Series Remote

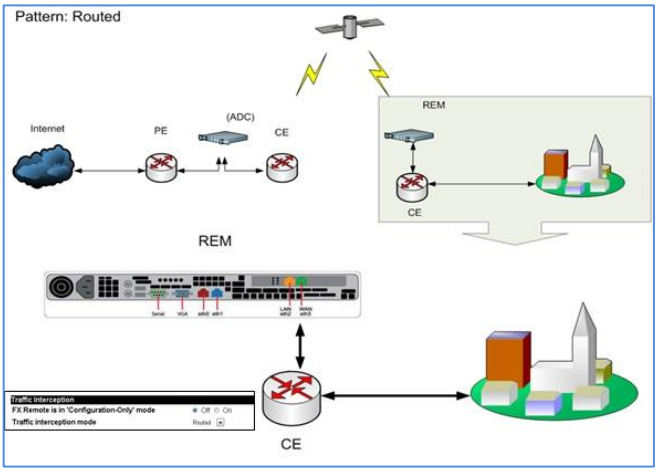

Go to **Configuration -> General Settings**

a. Change “Traffic interception mode” to “Routed”
b. Click Save

2. Go to **Traffic Interfaces -> In-Path Interfaces**
    a. Click the Add button to add a new in-path interface.
        - Enter the VLAN ID of 0, if no VLAN is to be tagged. Use the VLAN ID of the VLAN if traffic is to be VLAN tagged.
        - Enter the IP Address.
        - Enter the Netmask
        - Enter the Gateway
        - Enter any static routes needed in the “Routes” field.
    b. Click “Add In-Path Interface”

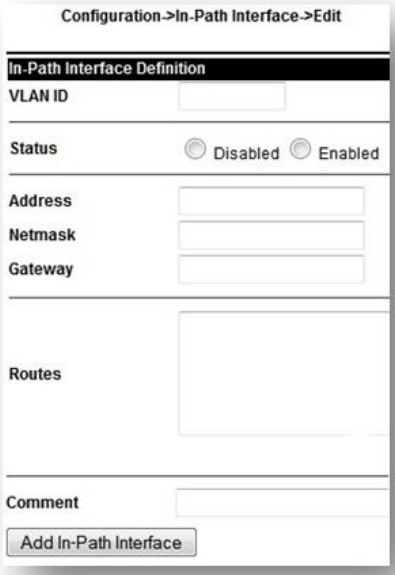

3. Go to **Traffic Interfaces -> LAN Interfaces.**
    a. Click on the “Add” button to add a new LAN interface.
        - Select “eth3” in the “Physical interface” selection box.
        - In the “In-Path Interface(s)” selection box, click the IP address of the in-path interface *see above*.
        - In the “Untagged in-path interface” selection box, select the IP address of the in-path interface *see above*. Select “none” if the VLAN tag is to be propagated across the WAN.
    b. Click “Add LAN Interface”.

    Go to Operations -> Shutdown and Restart Click on the “Restart Service Button”

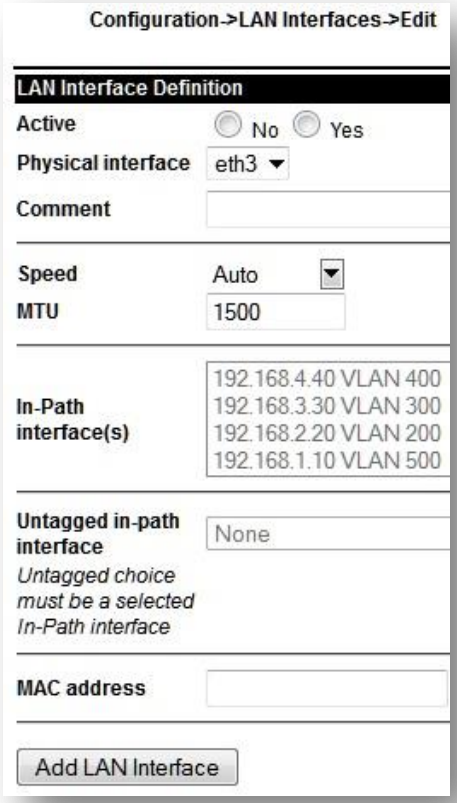

2.5 How to configure FX Series Installation Pattern (WCCP Mode)

2.5.1 Cable the Appliance

Connect the eth2 (LAN) port to a switch or a router.

2.5.2 Configure the Appliance

Login to the appliance through the browser interface at: *http://{IP_address_of_the_appliance}:10000*

1. Enter the default user name “comtech” and the default password “comtech”.
   a. Click Login.

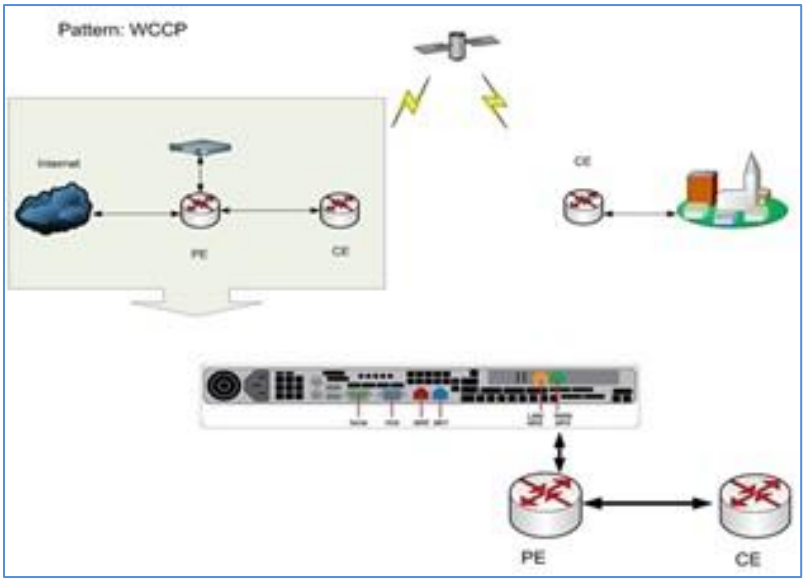

2. Go **to Configuration -> General Settings**
   c. Change “Traffic interception mode” to “WCCP”
   d. Click Save

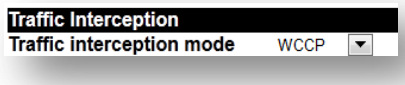

3. Go to **Traffic Interfaces -> In-Path Interfaces**
   a. Click the Add button to add a new in-path interface.
      - Enter the VLAN ID of 0.
      - Enter the IP Address.
      - Enter the Netmask
      - Enter the Gateway
      - Enter any static routes needed in the “Routes” field.
   b. Click “Add In-Path Interface”

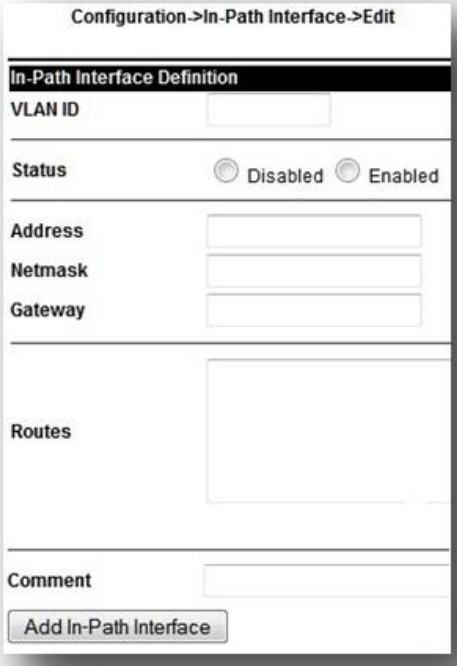

4. Go to **Traffic Interfaces -> LAN Interfaces.**
   a. Click on the "Add" button to add a new LAN interface.
      - Select "eth3" in the "Physical interface" selection box.
      - In the "In-Path Interface(s)" selection box, click the IP address of the in-path interface *see above*.
      - In the "Untagged in-path interface" selection box, select the IP address of the in-path interface *see above*.
   b. Click "Add LAN Interface".

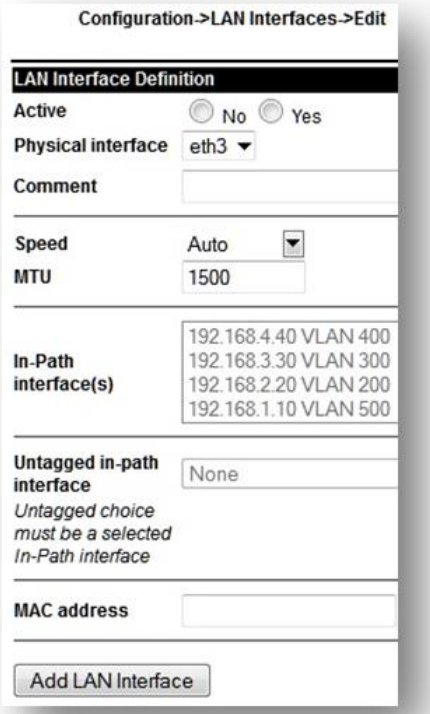

Go to **Status -> Real-Time Monitor.** Click ***"Restart Service".***
Go to Operations -> Shutdown and Restart Click on the "Restart Service Button"

2.5.3 Configure WCCP Settings

Go to **Traffic Interfaces -> WCCP**

1. Click "Add WCCP Definition"
   a. Select the IP address of the interface that will send the WCCP messages to the router in the "Source IP address" selection box.
      Enter the routers IP address in the "Router address" field.
   b. Select the local interface that will receive the GRE traffic if using GRE redirection. If using L2 redirection leave blank.
   c. If using GRE redirection enter the router identifier of the router in the "Remote GRE tunnel address" field. If using L2 redirection leave blank.
   d. Use default value for "Enablement".
   e. Use default value for "Critical".
   f. Select Redirection method.
   g. Select the same value used for Redirection method.
   h. For L2 use Mask Assignment scheme. For GRE use Hash Assignment scheme. * This may differ on highest end Cisco equipment.
   i. Leave the password field blank, unless one was configured for WCCP on the Cisco device.

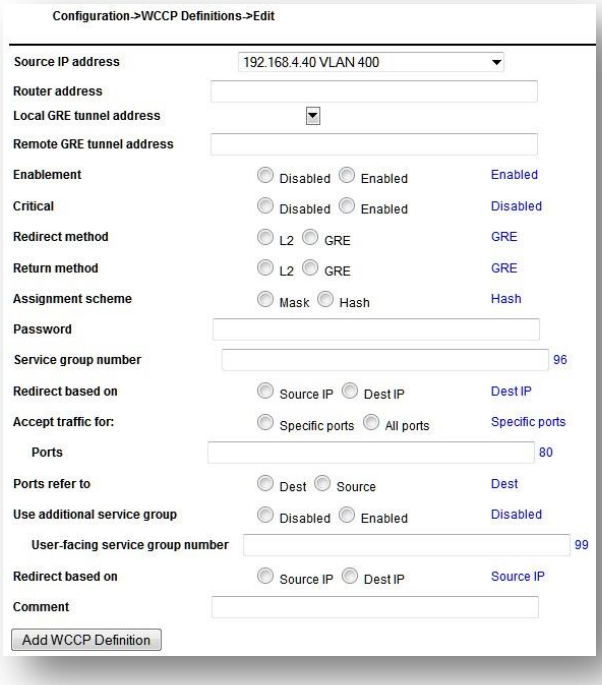

j. If using “web-cache” WCCP redirection, (no source IP address preservation) enter 0 in the “Service group number” field. If using source IP address preservation use the default value.
k. Use the default values for “Redirect based on”, Accept traffic for”, and “Ports”.
l. If using source IP address preservation select “Enabled” for “Use additional service group”. If using “web-cache” default this value.
m. Click “Add WCCP Definition”

*Go to **Operations -> Shutdown and Restart***

*Click on the **“Restart Service Button”***

NOTE: For more information on “redirection” or configuring on FX1000 appliances or older FX Series 4000 appliances, see the WCCP section: Other WCCP Configurations

2.6 How to Configure Two FX Series Appliances in a Mesh Configuration

This configuration consists of two FX Series appliances, one configured as a FX Series ADC and the second appliance configured as a FX Remote.

NOTE: Two FX1005 appliances can be installed in a rack using the Double Unit Rack Mount (See Section 1.8.8 above - FX Series FX-1005 Hardware Mounting Options)

2.6.1 Cable the Appliances

A short cable is provided for the connection from the Remote's LAN port to the ADC's WAN port (the yellow connection in the figure below:

(See picture of rear panel in section 1.8.6 above)

1. Connect the Eth3 (WAN) port of FX Remote to the satellite connection. [Blue]
2. Connect the Eth2 (LAN) port of the FX-Remote to the Eth3 (WAN) port of the FX-ADC. [Yellow]
3. Connect the Eth2 (LAN) port of the FX-ADC to the user network. [Green]

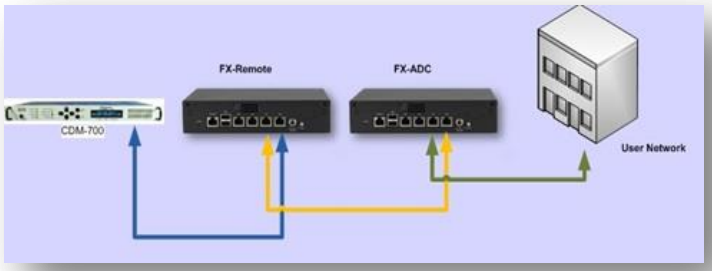

Figure 2-1 FX Series Mesh Connection Diagram

2.6.2 Configure the appliances

To configure the appliances, see Sections 2.2 How to Configure Appliance Management Address and Section 2.6 How to Configure Two FX Series Appliances in a Mesh Configuration.

2.6.3 Mesh installation with Redundancy capability

This scenario consists of two sets of appliances at the site. The configuration for each appliance is done separately to backup FX Series Appliances with fail to wire configuration. The configuration setups are similar as described in Section 3.6 below and is shown diagramed here.

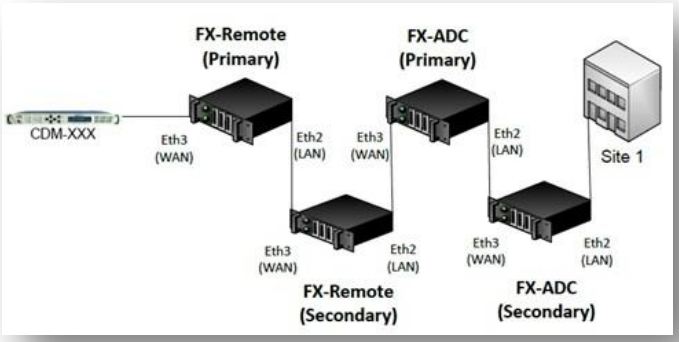

3 FX Series Configuration

Figure 3-1 FX Series Main Configuration Screen

3.1 Standard Configuration Overview

The following screens provide common interfaces for the FX Series appliances.

Application Policies

Customize the optimization techniques that will apply to your enterprise applications.

(See the Optimization Acceleration Settings in Section 7 below).

Authorization Realms

Web Application Policies

Layer 5 Application Policies

Management Settings

Basic Network Interfaces

The FX series reserves two ports, management and auxiliary, for management traffic.

Host Settings

Configure the host name and DNS settings to facilitate management and time synchronization.

General Settings

The General Settings control the method of traffic interception and WCCP. In addition, this section includes settings to configure basic HTTP settings, system time and software updates of FX-Remotes.

See Specific Sections below.

FX Series ADC Specific Settings Section

FX Series Remote Specific Settings Section

Multicator Settings

Configure reliable multicast fan-out settings for a Controller, Transmitters and Receivers.

Redundancy

Configure a redundancy cluster that can share a common configuration.

Traffic Interfaces

In-Path Interfaces
Configure In-Path Interfaces for user data.

LAN Interfaces
Set speed, MTU and VLAN options of the physical LAN interfaces. .

Port Definitions
Configure a list of port definitions.

WCCP Settings
The Web Cache Communication Protocol (WCCP) is a Cisco-developed content-routing technology which allows you to integrate cache engines into your network infrastructure.

Quality of Service Settings

Hierarchy
This screen delineates the configured relationship of Queues, Groups and Links.

Links
Maintain QoS links. Links correspond to a satellite modems.

Groups
Maintain QoS groups which allow you to group multiple QoS queues.

Group Filters
Maintain rules which classify traffic and assign it to one of the QoS groups.

Queues
Maintain QoS Queue definitions.

Queue Filters
Maintain rules which classify traffic and assign it to one of the QoS queues.

NOTE: The screens that have specific functionality for the FX Series ADCs or the FX Series Remotes can be found in these sections in this manual:

FX Series ADC Specific Settings Section
General Settings

FX Series Remote Specific Settings Section.
General Settings

3.2 Management Settings

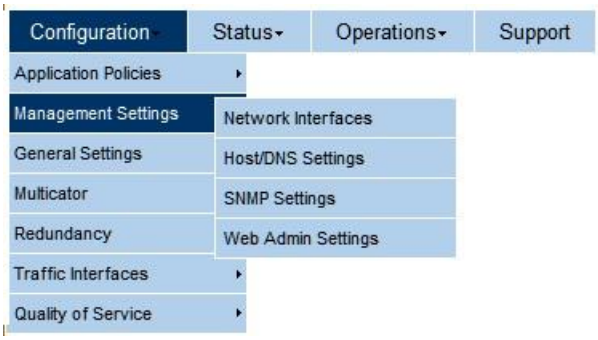

3.2.1 How to Configure Network Interfaces

The FX Series reserves two ports, management and auxiliary, for management traffic. This traffic is isolated from the ports in which accelerated traffic flows. These interfaces are tied to a management routing table which is not used for accelerated traffic.

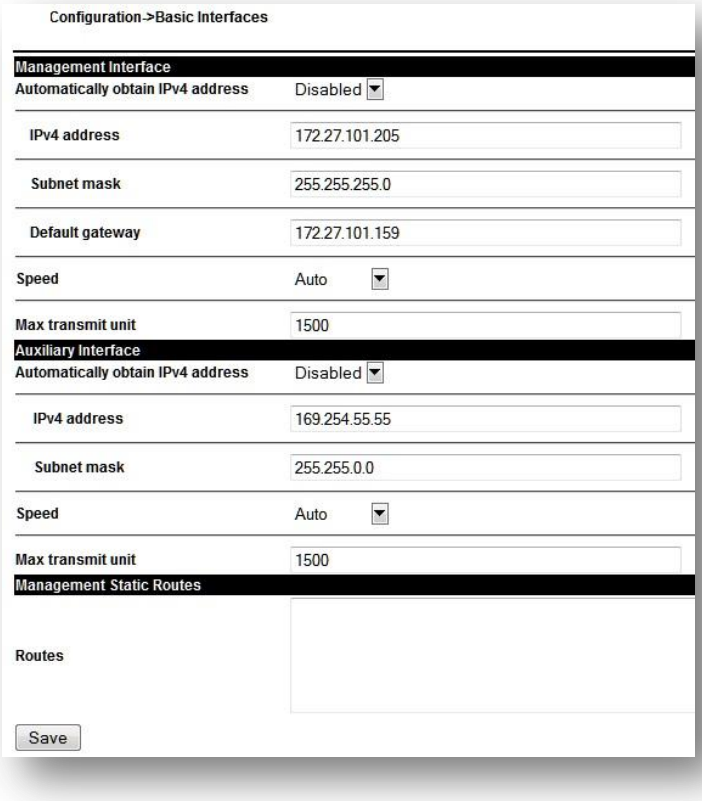

Figure 3-2 FX Series Basic Network Interfaces Screen

Management Interface

The management interface corresponds to the “eth0” Ethernet port. Typically the management interface is connected to a private network where system management tools such as ssh, the management web GUI, and SNMP are utilized.

Automatically obtain IPv4 address:

If set then the FX appliance will use DHCP to obtain an IP address, subnet mask, and default gateway. The factory default for the management interface is to use DHCP.

IPv4 Address:
This is the IP address of the management interface.

Subnet Mask:
This specifies the network that the management interface is on. The default value is 255.255.255.0.

Default Gateway:
This is the IP address of the gateway for which packets that are outside the bounds of the management subnet will be directed. A default gateway address which is on the same subnet as depicted by the "IPv4 address" and "Subnet mask" must be specified even if the gateway does not exist.

Speed:
This presents a pull-down selector of speed/duplex combinations that will be set for this interface. The default value is to automatically negotiate the speed and duplex.

Max Transmit Unit (MTU):
Specifies the max transmit unit. The default value is 1500. The range is 576 to 9000.

Auxiliary Interface

Automatically obtain IPv4 address:
If set, then the FX appliance will use DHCP to get an IP address, subnet mask, default gateway.

IPv4 address:
This is the IP address of the auxiliary interface. The factory default is 169.254.55.55.

Subnet mask:
This specifies the network that the auxiliary interface is on. The default value is 255.255.0.0.

Speed:
This presents a pull-down selector of speed/duplex combinations that will be set for this interface. The default value is to automatically negotiate the speed and duplex.

Max Transmit Unit (MTU):
Specifies the max transmit unit. The default value is 1500. The range is 576 to 9000.

Management Static Routes

Routes:
Enter into the text area static routes which are used by the management interface. Each static route must be entered on a separate line and must have exactly the following format:

Subnet "SubnetMask" "Gateway"

For example, to define a static route such that subnet 172.88.0.0/16 should be routed by gateway 172.27.101.99 you would enter the following:

172.88.0.0 255.255.0.0 172.27.101.99

3.2.2 How to Configure Host/DNS Settings

In most environments, configuring host names and DNS is not required for the FX to operate because for most accelerated traffic, the IP address of the content server is resolved by the originating client before it is processed by the FX. The host settings should be set to facilitate management and time synchronization.

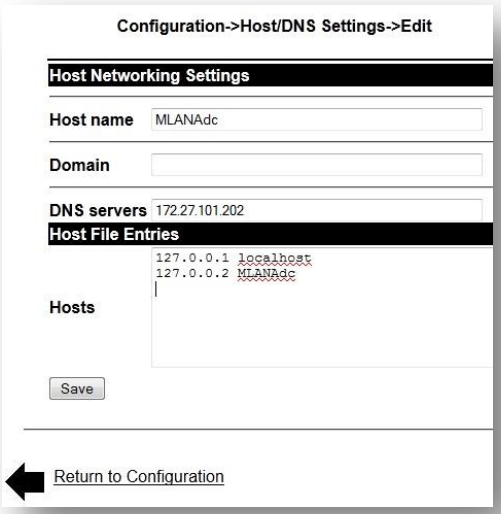

Figure 3-3 FX Series Host/DNS Settings Screen

Host/DNS Settings

Host name:
This is the host name of the appliance. This must be a "short" name and must not contain any periods. A fully qualified name is formulated by appending a '.' followed by whatever is entered into the "Domain" field.

Domain:
This is the DNS domain of the appliance.

DNS Servers:
Enter one or more IP addresses separated by commas.

Host/DNS File Entries

Hosts:
If no DNS is available, this field allows you to map specific host names to an IP address. This may be needed for active-passive redundant configurations. Each entry should be on a separate line. The format of each line is: "nn.nn.nn.nn FullyQualifiedHostName *OptionalShortHostName"*
Where nn.nn.nn.nn is the IpV4 address that you want to assign to FullyQualifiedHostName".

DNS Server Configuration Guidelines:
In order for the fully qualified host name to be accurate within a domain, it is a best practice to set your local DNS server to match the IP address of the FX appliance with the name "Host name" and "Domain" fields on this page. Or you can set a "Host File Entry" with the IP address of the FX appliance with the fully qualified name and the short name.

3.2.3 How to Configure SNMP Settings

The Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing devices on IP networks.

Settings

Basic SNMP Settings

Enable SNMP:
If 'Yes' then the FX will respond to SNMP requests otherwise the SNMP services are not loaded. The default is 'No'.

Read-only community string:
This specifies the group of SNMP monitors that have read-only access to the MIB. The default value is 'public'.

Read-write community string:
Specifies the group of SNMP monitors that have read-write access to the MIB, a typical value is 'private'.

Traps

Enable Traps:
If 'Yes', then the FX will send SNMP trap messages to the address specified in 'Trap destination' when certain events occur.

Trap community:
This field specifies the community that will be included in the trap messages that the FX sends. The default value is 'comtech'.

Figure 3-4 FX Series SNMP Edit Screen

Trap destination:
This field specifies the host name or IP address of the management station that will receive SNMP traps sent by the FX.

System Information

Name:
This is the management name assigned to this FX.
The default value is the serial number of the FX.

Location:
This optional string describes the physical location of the FX.

Contact:
This optional string specifies the contact information, typically an email address, for the FX.

3.2.4 How to Configure Web Admin Settings

Figure 3-5 FX Series Web Management Interface Screen

Administration

Use SSL for the administrative Web GUI:
If “enabled” then HTTP/S must be used when managing the FX from the Web graphical user interface. The default value is “Disabled”.

3.3 Traffic Interface Settings

3.3.1 How to Configure In-Path Interfaces

The In-Path Interfaces settings allow you to maintain in-path interfaces. A list of previously defined in-path interfaces are displayed in the order in which they were defined. An existing entry may be chosen by clicking on the IP Address. The user adds interfaces by clicking the “Add” button. Interfaces can be “enabled” or “disabled” by checking box to the left of the IP address and clicking on the enable or disable button. Clicking on the Interface “Name” will allow you to modify that Interface. Each of those actions will then lead to an In-Path Interface screen.

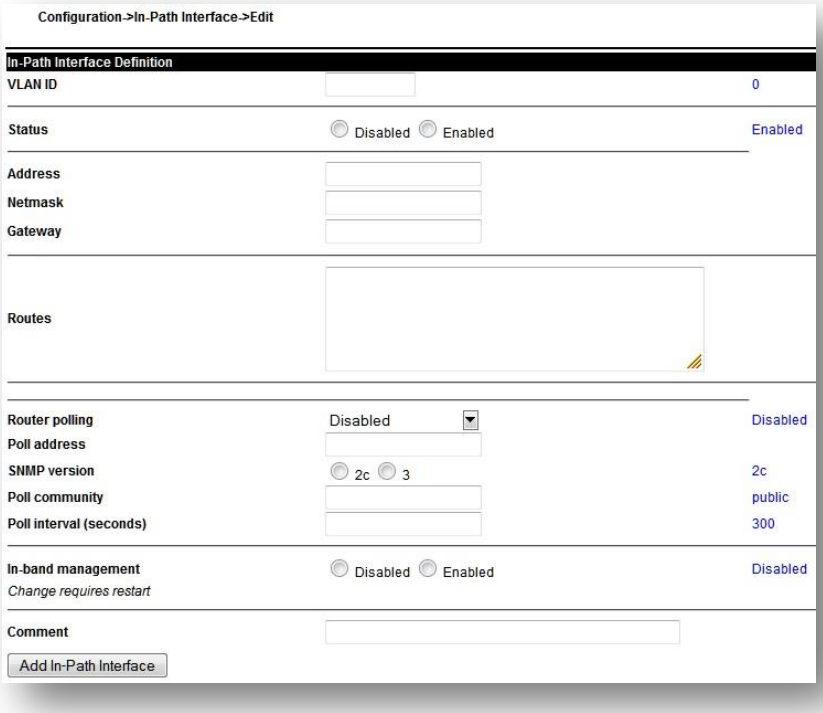

Figure 3-6 FX Series In-Path Interfaces Screen

Best Practices for Routed Mode Configurations

If using two in-path interfaces in routed mode and both interfaces are on the same VLAN, which is normally the case, then only one in-path interface can have a default gateway. The in-path interface with no default gateway defined must have a routing configuration that is comprised entirely of static routes.

In-Path Interface Definition

VLAN ID:

If this interface carries tagged VLAN traffic, then enter the VLAN ID number which is a value between 2 and 4094. Untagged traffic should have a value of 0. The default is 0.

Status:
This is the status of the In-Path interface. It must be set to "Enabled" for the In-Path interface to receive and process data. The default is Enabled.

Address:
This is the IPv4 address of this interface.

Netmask:
This defines the subnet boundaries of this interface.

Gateway:
This is the default gateway for this interface.

Routes:
Enter into the text area static routes which are used by this in-path interface. Each static route must be entered on a separate line and must have exactly the following format:

Subnet "SubnetMask" "Gateway"

For example, to define a static route such that subnet 172.88.0.0/16 should be routed by gateway 172.27.101.99 you would enter either of the following supported formats:

172.88.0.0 255.255.0.0 172.27.101.99

Or 172.88.0.0 /16 172.27.101.99

Router Polling:

If enabled, then SNMP router polling will be used (over the management interface) to ascertain the routes that should be added to the route table that is associated with this interface. The default value is disabled. Router polling is only supported for VLAN 0.

Poll address:
This is the IPv4 address of the router which will respond to the SNMP router poll requests.

SNMP Version:
This is the version of SNMP that will be employed when making the router poll requests. The default value is 2c.

Poll Community:
This is the SNMP community that is associated with the router poll request. The default value is "public".

Poll Interval (seconds):
This is the frequency in seconds that the router tables will be updated based upon the SNMP router polling response.

In-Band Management:

Normally out-of-band management of the FX is accomplished through the management interface, however in some scenarios, out-of-band management is not feasible and management of the device must be performed over an in-path interface. If this is the case, only one in-path interface may be used for in-band management. The default value is disabled.

Comment: This field provides a means to store useful information about the configuration

NOTE: Changing this setting requires a restart of the acceleration service on the "Operations->Shutdown/Restart" page..

3.3.2 How to Configure LAN Interfaces

Overview

The LAN Interface settings allow you to maintain LAN interface definitions. A list of previously defined in-path interfaces are displayed in the order in which they were defined. The user adds interfaces by clicking the "Add" button. Interfaces can be "enabled", "disabled" or "deleted" by checking box to the left of the IP address and clicking on the enable, disable or delete button. Clicking on the Interface "Name" will allow you to modify that Interface.

Working with LAN Interfaces without WAN Optimization

If you are utilizing the FX strictly for ACM QoS or packet compression, in order to configure a non-default MTU of the network interfaces, you must define a LAN interface for all network interfaces for which the traffic to be processed by the FX will flow through, typically these will correspond to "eth2" (LAN) and "eth3" (WAN) physical interfaces. In this case, it is recommended to have the MTU on both LAN interfaces to be the same.

NOTE: You must only do it if you want to have a non-default MTU

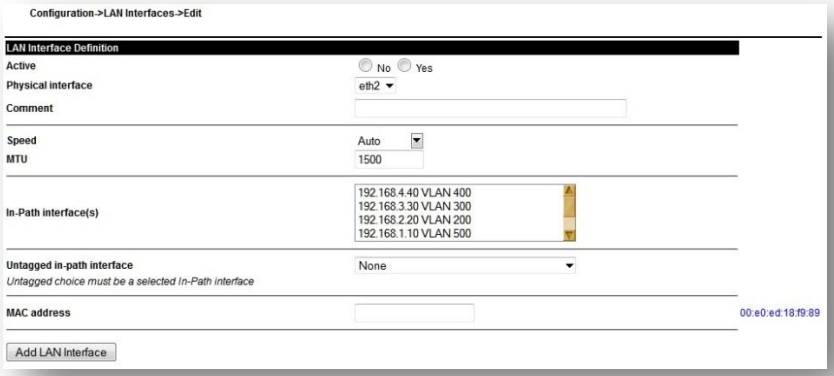

Figure 3-7 FX Series LAN Interfaces Screen

Configure LAN Interfaces

Active:
If adding a new LAN interface, this field allows you to set the initial status.

Physical interface:
Select the physical interface from the pull-down.

Comment:
This field provides a place to store any user defined comment to describe the rationale for this LAN Interface definition.

Speed:
Select speed and duplex from the pull-down.

MTU:
Specify a value between 68 and 9000. Note that when operating in "In-Line" mode, the MTU of the WAN interface will automatically be set to match the setting of the LAN interface.

In-Path interfaces:
Select the in-path interfaces that can be connected to the physical interface. In a trunked environment, there may be multiple in-path interfaces connected to the physical interface. VLAN tags will be preserved.

Untagged in-path interface:
You can specify one and only one of the selections from the "In-Path interfaces(s)" field or "None" If an in-path interface is selected then the VLAN tag associated with that in-path interface will be applied to the traffic received before forwarding it to the WAN interface and removed when forwarding traffic from the WAN interface. If "None" is selected then no tags are added or removed. The "None" value will only be used when connecting to a "Trunked" interface.

MAC address:
This field specifies the Ethernet address of the interface.

This should only be set in redundant configurations where the traffic interception mode is 'in-line'. In this case, this field should be set to the permanent MAC address of the primary. The permanent MAC address of this FX is shown in blue. If the field is left blank, then the permanent address of the FX is used.

3.3.3 How to Configure Port Definitions

The Port definitions screen allows you to define which IP address and port combinations that the FX Series ADC will listen on and what protocols should be accepted over these combinations. A port definition is required for every IP address and port combination on which the FX Series ADC will accept connections in a proxy mode. For each port definition you must specify the protocol that will be used. Port definitions are only needed if you will be directing traffic to the FX Series ADC as a proxy, or from a remote software client that is running acceleration plug-in.

A list of previously defined Port definitions is displayed in the order in which they were defined. An existing entry may be chosen by clicking on the port. You can "Enable", "Disable", or "Delete" one or more Port definitions by selecting the checkbox to the left of the port column and clicking on the desired button. By clicking on "Add" you can add a new port definition which will bring up this port definition screen.

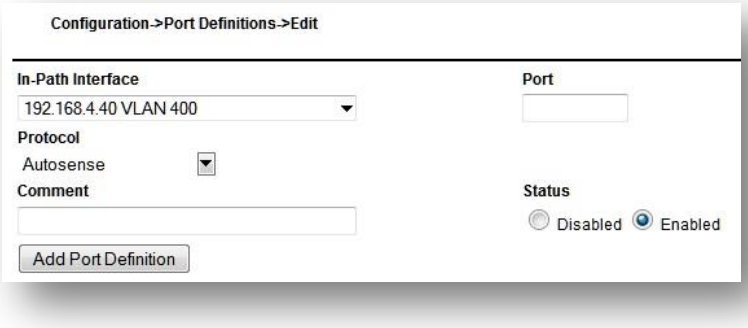

Figure 3-8 FX Series Port Definitions Screen

Configure Port Definitions

In-Path Interface:
Specify the IP address that is associated with this port definition.

Port:
This field will be filled in automatically as you set the "Protocol" field. After setting these fields, you can then override the port field to create a unique IP Address / Port combination.

Protocol:
This specifies the protocol that will run over this port. There are the following choices:

HTTP:
This choice specifies that you want the ADC to function as either a forward or reverse proxy on this port.

Accelerated HTTP/L5:
This choice specifies that you want the ADC to use this port to service the HTTP Acceleration Protocol (HAP) that has been extended to also accelerate non-HTTP TCP/IP based protocols at layer 5. Accelerated HTTP/L5 is only available if you have deployed the acceleration plug-in to your remote users.

Autosense:
This setting supports the AOD injection where both HTTP traffic and accelerated HTTP can flow over the same port.

Comment:
This provides a place to store any user defined comment to describe the rationale for this port definition.

Status:
This allows you to control whether this port definition is enabled or disabled.
Example Port Definitions
By default, port definitions are not required to function as a one-sided FX Series ADC or as a head-end serving FX Series Remote appliances.

The table below shows port definition setting examples:

IP	Port	Protocol - SLL	Typical Use
any	80	HTTP / Acc. HTTP - Autosense	Transparent redirection with capability to inject AOD.
any	8080	HTTP	Forward proxy of HTTP traffic from standard browsers
any	4917	Acc. HTTP/L5	Accelerated traffic between standard and advanced clients and FX Series ADC

Setting up an HTTP Forward Proxy
A forward proxy requires that an end user specifically set their browser proxy settings such that port 80 traffic is specifically directed to the IP address of an in-path interface of the ADC on a specific port (usually 8080). In order to get this to work some additional steps are required:

1. An L5 Policy must be defined for port 8080 traffic.
2. This L5 policy must have the "certified application" set as "HTTP Traffic".
3. This L5 policy must have "Protocol" defined as "Generic TCP".

The "Protocol" on the "Port Definition" must be defined as "HTTP".

3.3.4 How to Configure WCCP

Overview

The Web Cache Communication Protocol (WCCP) is a Cisco-developed content-routing technology which allows you to integrate cache engines into your network infrastructure.

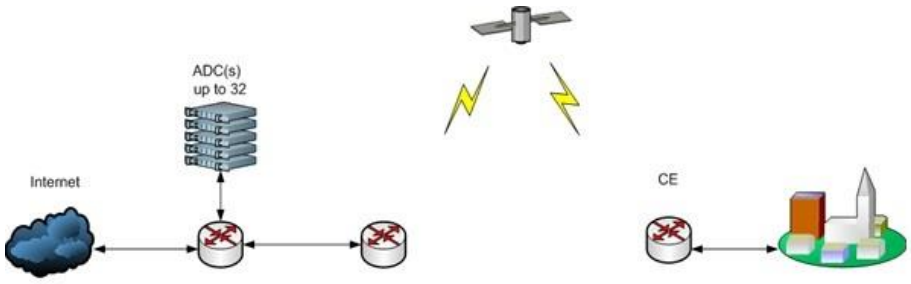

WCCP Configuration Considerations

There are two basic configurations that will be used when installing an appliance.

1) **Web-cache or transparent proxy interception.**
   In this mode of interception the FX will have the same characteristics as a standard proxy. The Cisco device will redirect traffic to the appliance, which will then make request on behalf of the user using the appliance's IP address as the source.
   NOTE: Only one service group is required, service group zero. On the Cisco device this will be configured as "web-cache".
2) **Dynamic service groups or source IP address preservation**
   In this mode of interception the FX will the same characteristics as an in-line device. The Cisco device will redirect traffic to the appliance, which will then make request on behalf of the user using the user's address as the source (spoofing).
   NOTE: This configuration requires two service groups, inbound and outbound.

WCCP Cisco Device configuration

1) We will use either eth0 or eth1 when installing in a WCCP pattern.
   Log into the Cisco device and identify the inbound and outbound interfaces.
   These must correspond to in-path interfaces.
   In the global configuration enable WCCP with the appropriate commands.
   If configuring as a web-cache we will enter the following: "ip wccp web-cache"
2) For a WCCP with source IP address preservation setup.
   We will enter the following commands: "ip wccp 99", "ip wccp 96"
   By default our appliances use service groups 99 for outbound traffic and 96 for inbound traffic.
3) At the interface level if we are configuring a web-cache setup.
   We will enter the following command:
   For all inbound interfaces: "ip wccp web-cache redirect in"
4) At the interface level if we are configuring a source IP address preservation setup.
   We will enter the following commands:
   For all inbound interfaces: "ip wccp 99 redirect in"
   For all outbound interfaces: "ip wccp 96 redirect in"
   It is possible to control which traffic is redirected by subnet using the redirect-list option

WCCP Definitions

This screen allows you to maintain WCCP definitions. A list of previously defined WCCP definitions is displayed in the order in which they were defined. An existing entry may be chosen by clicking on the router address. By clicking on "Add" you can add a new WCCP definition which will bring up the WCCP definition screen.

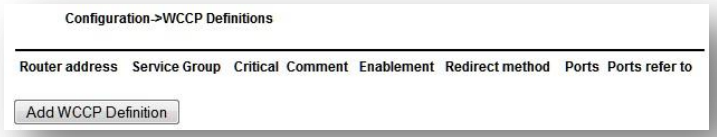

Configure WCCP Definitions

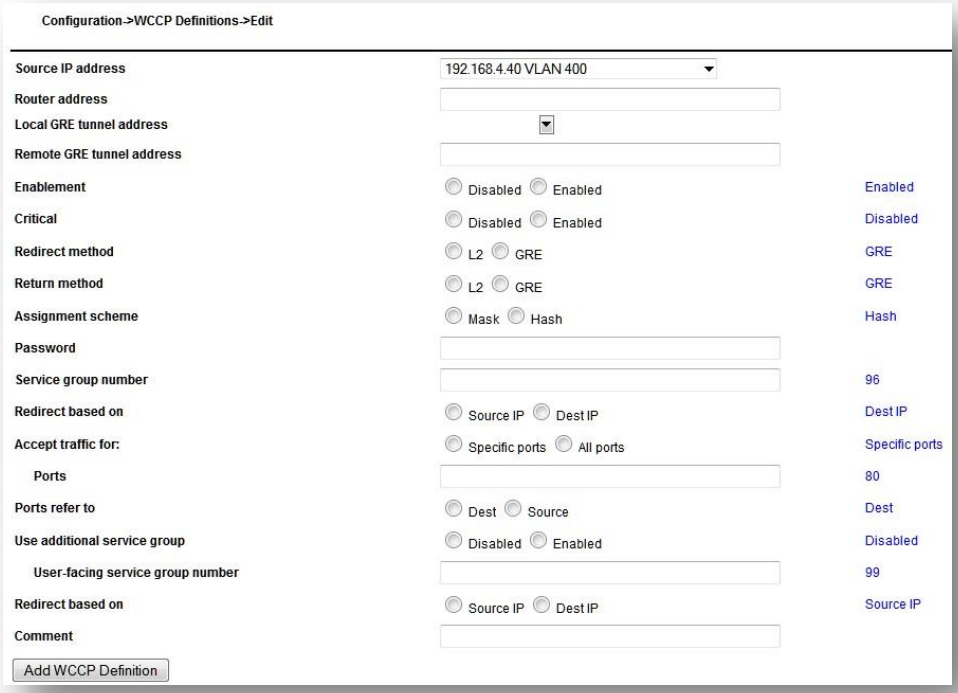

Figure 3-9 FX Series WCCP Definitions Screen

Source IP address:

This is the IP address (must have been already defined as an In-Path interface on the FX) that will be used when sending WCCP messages to the router. The IP address must be selected from the list of in-path interfaces.

Router address:

This is the address of the primary router to which WCCP messages will be directed. This setting must be specified.

Local GRE tunnel address:

This is the IP address of the local end of the GRE tunnel**.** If this field is not specified then the "Source IP address" will be used for the local endpoint. This field is not needed if L2 redirection is specified.

Remote GRE tunnel address:
This is the address of the router which will send the redirected traffic to the FX Series ADC in a GRE tunnel. If this field is not set then the FX will attempt to dynamically learn the address by examining the WCCP packets from the router. This field is not needed if L2 redirection is specified.

Enablement:
This specifies if this WCCP definition should be processed. The default value is enabled

Critical:
If set, and "Use additional service group" is disabled, then this service group is considered critical. Non-critical service groups will not attempt to negotiate WCCP with the router unless all critical members have seen their IP address in the assignment map or hash allotment and are in a usable state.

Redirect method:
This specifies the method in which the router or switch will direct packets to the FX. The choices are "GRE" (Generic Routing Encapsulation) or "L2" which means that the router will simply modify the MAC destination address to point to the FX. The default is "GRE".

Return method:
Although the FX never returns redirected packets to the router, it may be necessary to set this to "GRE" even though "L2" was specified as the redirect method in order to successfully negotiate WCCP.

Assignment scheme:
This specifies how the router or switch will decide which FX to direct the packets. In general, this should be set to "Mask" for switches and "Hash" for router. The default setting is "Hash"

Password:
If WCCP packet signing is required then this password must match the setting of the WCCP router. The default is no password.

Service group number:
This is the WCCP service group that the FX should join. The default value is 96.

Redirect based on:
If "source" then the router will redirect responses from the content server to this member, otherwise the router will redirect client requests that otherwise would have been directed to the content server. Service groups are defined at the router. The default is the "Destination".

Accept traffic for:
This radio button allows you to control if only specific ports or all ports should be redirected to the FX Series FX. If "Specific Ports" selected then these are specified in the "Ports" field. If "All ports" is selected then the WCCP router will direct all TCP and UDP traffic to the FX. The default value is "Specific Ports".

Ports:
This defines the TCP and UDP ports that the router should transparently redirect to the FX. Up to 8 ports may be specified separated by a comma. The default value is 80.

Ports refer to:
This indicates if the ports field pertains to the source port (for responses from content server) or destination port (for requests from clients). The default value is "Source".

Use additional service group:
You can define two service groups within the same WCCP definition. This is normally used if you want the FX to preserve the source IP address of the remote clients when making requests to content servers on behalf of those clients. However, if using the "extra" group then the definition is not deemed as non-critical and will not verify that "critical" service groups are in a usable state. If this is set you must also enable "Preserve client IP addresses" in the "Other" section on the "Configure->General" page. See more detailed description titled "WCCP IP Spoofing Configuration" below. The default value is "Disabled".

User-facing service group:
This is the WCCP service group that the FX should join to receive redirected client traffic. The FX will not attempt to join this group unless it successfully enrolls in the main service group. This prevents the situation where client requests are redirected to the FX when it is not able to receive server responses. The default value is 99.

Redirect based on:
If "Source IP" then the router will redirect responses from the content server to this member, otherwise the router will redirect client requests that otherwise would have been directed to the content server. Service groups are defined at the router. The default value is "Source IP".

Comment:
A comment of up 80 characters can be entered into this field.

WCCP Router Configuration and Status Monitoring

Configuration:

The following is an example of some common WCCP Router "cli" commands.

conf t
ip wccp enable
ip wccp version 2
interface (specify interface carrying traffic)
ip web-cache redirect
CTRL-Z

Status Monitoring:

The following WCCP Router "cli" commands can show status:

show ip wccp
show ip wccp 99 view
show ip wccp 96 detail
term mon
debug ip wccp packets
debug ip wccp events
clear ip wccp

WCCP IP Spoofing Configuration for Routers

The FX can preserve the source IP address of the remote client when making requests on their behalf by joining two service groups. The first service group receives the redirected client requests and is also known as the "User-facing" service group. The second is referred to as the "Server-facing" service group and it receives the redirected server responses. If two or more FXs have joined these service groups, then the router will be instructed to split the load of the user-facing service group based on source IP address, and the responses of the server-facing service group will be split based on destination IP address. This technique ensures that the response will be directed to the same FX that originated the request on behalf of the remote user.

The recommended router configuration is to use three interfaces, each corresponding to a different subnet. To illustrate the setup, we provide an example configuration along with a "show running-config" that is compatible with the default WCCP settings of the FX.

Example:

Interface A: (Ethernet0/0)
This is the user-facing subnet that receives redirected requests from clients.

Interface B: (Ethernet0/1)
This is the server-facing subnet that receives redirected responses from the content server.

Interface C: (Ethernet1/0)
FX subnet

Service group 99
This should be defined to handle redirected outbound requests from the users destined for the subnets on Interface B. "Interface C" must be excluded from this to avoid loop-backs that would otherwise occur when FXs spoof the user IP addresses.

Service group 96
Should be defined to handle redirected responses from content servers that would have otherwise been sent out on "Interface-A**".**

The subnets:
A: User (192.168.103.xxx subnet)
B: Content servers - all other subnets via gateway at 192.168.101.158
C: FXs (192.168.106.xxx)

#show running-config

```
Building configuration...

Current configuration: 948 bytes
!
version 12.2
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname 2600-lab
!
enable password xxxx
!
memory-size iomem 10
ip subnet-zero
ip wccp 96
ip wccp 99
!
!
no ip domain-lookup
ip domain-name example.enterprise.com
ip name-server 192.168.101.202
!
!
interface Ethernet0/0
ip address 192.168.103.224 255.255.255.0
ip wccp 96 redirect out
half-duplex
!
interface Ethernet0/1
ip address 192.168.101.224 255.255.255.0
ip wccp 99 redirect out
half-duplex
!
interface Ethernet1/0
ip address 192.168.106.224 255.255.255.0
```

```
no ip route-cache
ip wccp redirect exclude in
half-duplex
!
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.101.158
ip http server
ip pim bidir-enable
```

WCCP IP Spoofing Configuration for Switches

Switches tend to have less CPU power than a router but on the other hand they have the ability to handle traffic flow decisions in hardware. In order to leverage the hardware switching capabilities the following configuration settings are recommended:

- On the FX, use “L2” Redirection method
- On the FX, use “Mask” assignment scheme
- On the FX, do not define separate service group definition records, instead set the “Use additional service group field”, this is because the Cisco L2 expects the same WCCP source port to be used to conduct WCCP negotiations.
- On the switch, use “redirect in” to direct packet flow to the appliance.

On the switch, never use “redirect-out”
On the switch, do not use “redirect exclude in”

In the same subnet scenario described above, the following is an example of a configuration for an intelligent switch:

#show running-config

```
.
.
.
!
ip routing
ip wccp 96
ip wccp 99
!
interface Vlan1
ip address 192.168.101.225 255.255.255.0
ip wccp 96 redirect in
!
interface Vlan3
ip address 192.168.103.225 255.255.255.0
ip wccp 99 redirect in
!
interface Vlan5
ip address 192.168.105.225 255.255.255.0
!
interface Vlan6
description for 106 subnet
ip address 192.168.106.225 255.255.255.0
!
```

Other WCCP Configurations

Using “redirect-list” to select specific redirection
For testing purposes, or to gradually stage traffic redirection to the FX Series ADC, a Cisco router will support redirection by either access control lists or group lists. For example:
ip wccp 99 redirect-list *access-list*

LAN and In-Path Interface Requirements for WCCP

The FX Series uses the eth2 physical interface to conduct the WCCP protocol with the router or switch and also to receive redirected requests and responses. Therefore in order to configure WCCP, you must define an in-path interface with an IP address that is on the same subnet as the Cisco router or switch. The gateway for this in-path interface must be that of the Cisco router or switch interface the appliance is connected to. The VLAN ID must be 0. Following this, a LAN interface must be defined for eth3 and assigned to the aforementioned In-Path interface.

Configuring WCCP on earlier models

On FX-1000 and some earlier models of FX-4000, the specialized fail-to-wire network interface card requires that an eth3 LAN interface be defined in order to run WCCP over eth2, even though it’s not actually used. Therefore, on an FX-1000, an in-path interface with a non-existent VLAN must be defined. This in-path interface must subsequently be assigned to a LAN interface for eth3. It is not required that the eth3 physical interface be cabled to anything.

3.4 Quality of Service

3.4.1 Theory of Operations for QoS and Traffic Shaping

Structure

There are two levels of groups and a third level of queues that can be configured. Traffic coming into the appliance is separated by filters into the level 1 groups. This traffic can subsequently be separated by filters into a second level of groups, and then filtered into queues where traffic will be released to the WAN based on the QoS and shaping rules defined. The groups as well as the queues have a MIR and a CIR defined that is used to control that amount of traffic that is allowed to pass through the groups and queues to the WAN.

Links

Links represent physical connections within the network. Links are used to bring any data rate limitations into the traffic shaper that is imposed by the physical connection. Each link can be a point-to-point connection or it can be a point-to-multipoint connection. The network can have a mix of links. Each of the far side points will have a peer FX. In the case of the point-to-multipoint link, there is one peer for each remote "multipoint". The link rate is the data rate of the outbound modem. ACM as well as CCM is supported in the link, if ACM is enabled, then the FX will read the data rate from the modem.

Traffic Shaping

The traffic shaper consists of two levels of groups and an additional level of queues. Traffic from the queues is enabled onto the WAN interface. The groups are used to separate the traffic into the appropriate queues. Each group can have a MIR and CIR data rate associated with it, which gives further refinement on the traffic shaping.

A point-to-multipoint link can have Level 1 groups that span multiple remote points, all within the same link. When doing WAN/OP functions, the peer FX will be automatically discovered, and once discovered, WAN optimization will occur. However, when doing header compression, the remote peer must be configured. This is done on the Level 2 Group configuration screen.

Each group has a CIR (Committed Information Rate) and MIR (Maximum Information Rate) associated with it. Data moving through a group will have the same priority. Each peer group's CIR will be met if possible, if not, then each will get their share using Stochastic Fair Weighting. If the CIR's can be met, then each peer group will be allowed up to its MIR amount of traffic, again, sharing traffic using Stochastic Fair Weighting.

Some care must be taken in setting the MIR's and CIR's. The sum of the CIR's of the children need to be less than or equal to the CIR of the parent. This means that the sum of the CIR's of the queues that children of a Level 2 Group must be less than that groups CIR, and the sum of the CIR's of level 2 Groups that are children of a Llevel1 Group must be less than that groups CIR. A child's MIR must be less than or equal to its parents MIR.

Once data makes it to a queue, it can be released to the WAN based on group priority, CIR and MIR.

Traffic Control Properties

Queue drain algorithm:

This specifies the drain method for scheduling outbound packets for all Queue definitions. The drain algorithm is set by the link and applies to all queues associated with that link.

Strict Priority:

In the Strict Priority Drain algorithm, higher priority queues are depleted before lower priority queues pass traffic. Traffic is capped at the link rate.

Min-Max:

The Min-Max drain algorithm is priority based. If there is enough data rate available, each queue will receive their respective CIR. If there is not sufficient data rate to satisfy all requested CIRs, then traffic will be dropped starting with the lowest priority queue and progressing through the queues in ascending priority until the requested CIR is met.

When traffic is dropped from queues with the same priority, then each of the equal queues will have traffic dropped proportionally.

Once all requested CIRs are met, if there is additional data rate that can be filled, it will be allocated to the queues in order of priority starting with the highest. Each queue is given additional data rate up to the requested rate, or MIR, whichever is lowest. If there are queues at the same priority, they are granted additional data rate proportionally

3.4.2 QoS Configuration Hierarchy Screen

This screen represents all of the hierarchy within the QoS system in the appliance. Configuration is not done in this screen, but it does provide links to the configuration screens for the parameters displayed here.

Important Considerations

Internal Signaling Traffic:

Internal signaling traffic represents control messages sent between FXs when header compression is in use. This traffic is sent at the highest priority and is not considered in the CIR. The amount of internal signaling is a function of the types of FX units on each side of the link. If there is are FX4010s on both sides of the link, then the internal signaling in both directions will be 30 kbps duplex and 15 kbps simplex. If all of the FXs are FX1005s then the amount of internal signaling will be 6 kbps duplex and 3 kbps simplex. Finally if there is an FX4010 on one side and all FX1005's on the other, the amount of signaling will be 18 kbps, 15 kbps + 3 kbps. In addition, if the compression tunnels are not completely configured, then the amount of traffic can be significantly higher.

ARP Traffic:

Generally it is important to set up queues and filters such that ARP traffic is sent at the highest priority.

In the table, there is one row for each queue. The table represents the hierarchy of the QoS tree, with parents to the left and children to the right. Because there is one row per queue, the groups and the links will likely show up multiple times in the table. If a level 2 group has 3 children queues, then it will show up in 3 lines. If a level 1 group has 8 descendent queues, it will show up on 8 lines.

Link		Group level 1				Group level 2				Queue					
Name	Data Rate	Filter	Name	CIR	MIR	Filter	Name	CIR	MIR	Filter	Name	CIR	MIR	Prio	HC
Link_1	25,000	log-A	LogicA	10,000	20,000					arps	Signaling Q	0	20	1	No
Link_1	25,000	log-A	LogicA	10,000	20,000					voip -Default-	VoIP	1,000	2,000	2	Yes
Link_1	25,000	log-A	LogicA	10,000	20,000					HTTP	Web Browsing	0	20,000	3	No
Link_1	25,000	log-A	LogicA	10,000	20,000						-Default- *	0	20,000	8	No
Link_2	100,000	acme	acme	60,000	100,000		Argo Acme Crew	10,000	40,000	AAC_VoIP	AAC_VoIP	4,000	10,000	3	Yes
Link_2	100,000	acme	acme	60,000	100,000		Argo Acme Crew	10,000	40,000	AAC_HTML	AAC_Web Browsing	5,000	20,000	4	No
Link_2	100,000	acme	acme	60,000	100,000		Argo Acme Management	20,000	40,000	AAM_VoIP	AAM_VoIP	5,000	10,000	2	Yes
Link_2	100,000	acme	acme	60,000	100,000		Argo Acme Management	20,000	40,000	AAM_HTML	AAM_Web Browsing	10,000	20,000	3	No
Link_2	100,000	acme	acme	60,000	100,000		Cyclops Acme Crew	10,000	40,000	CAC_VoIP	CAC_VoIP	3,000	10,000	3	Yes
Link_2	100,000	acme	acme	60,000	100,000		Cyclops Acme Crew	10,000	40,000	CAC_HTML	CAC_Web Browsing	4,000	10,000	4	No
Link_2	100,000	acme	acme	60,000	100,000		Cyclops Acme Management	20,000	40,000	CAM)VoIP	CAM_VoIP	3,000	10,000	2	Yes
Link_2	100,000	acme	acme	60,000	100,000		Cyclops Acme Management	20,000	40,000	CAM_HTML	CAM_Web Browsing	5,000	20,000	3	No
Link_2	100,000	xylex	xylex	40,000	100,000		Argo Xylex Crew	5,000	15,000	axc_VoIP	axc_Voip	1,000	6,000	3	Yes
Link_2	100,000	xylex	xylex	40,000	100,000		Argo Xylex Crew	5,000	15,000	axc_HTML	axc_Web Browsing	4,000	10,000	4	No
Link_2	100,000	xylex	xylex	40,000	100,000		Argo Xylex Management	15,000	30,000	axm_VoIP	axm_VoIP	3,000	6,000	2	No
Link_2	100,000	xylex	xylex	40,000	100,000		Argo Xylex Management	15,000	30,000	axm_HTML	axm_Web Browsing	12,000	20,000	4	No
Link_2	100,000	xylex	xylex	40,000	100,000		Cyclops Xylex Crew	5,000	15,000	cxc_VoIP	cxc_VoIP	1,000	6,000	2	Yes
Link_2	100,000	xylex	xylex	40,000	100,000		Cyclops Xylex Crew	5,000	15,000	cxc_HTML	cxc_Web Browsing	4,000	10,000	4	No
Link_2	100,000	xylex	xylex	40,000	100,000		Cyclops Xylex Management	15,000	30,000	cxm_VoIP	cxm_VoIP	3,000	6,000	2	Yes
Link_2	100,000	xylex	xylex	40,000	100,000		Cyclops Xylex Management	15,000	30,000	cxm_HTML	cxm_Web Browsing	12,000	20,000	3	No

Figure 3-10 FX Series QoS Hierarchy Screen

Configuration Parameters

Link

Name:

This is the name given to the link. There are two basic ways to add a link. One is to navigate to the Configuration\Quality of Service\Link page and select the add button. The other is to navigate to Configuration\Quality of Service\Link\Add... Once a link is added, it will show up here by name and its name will link back to the link configuration page. It is possible and normal for a specific link to show up on multiple lines.

Data Rate:
This gives the configured data rate for the link. If the modem is configured for ACM, this will indicate the rate that is read from the modem.

Group Level 1

Filter:
This is a list of all filters defined to get to the group immediately to the right in the table. It is also a hyper link that will take you to the group filter page. There are two basic ways to add a filter. One is to navigate to the Configuration\Quality of Service\Group Filters page and select the add button. The other is to navigate to Configuration\Quality of Service\Group Filters\Add.

Name:
This is the name given to the group. There are two basic ways to add a group. One is to navigate to the Configuration\Quality of Service\Group page and select the add button. The other is to navigate to Configuration\Quality of Service\Group\Add. The name is a hyper link that group's configuration page. At each level, there is a default group. This is the group where traffic that is not otherwise selected by a filter will go. This group by default will have the name "Default". This name can changed and another group can be selected as the default. The default queue will be designated by the asterisk "*" appended to the name.

CIR:
This is the Committed Information Rate of the group. Setting this to a number higher than the configured Link rate will result in an error.

MIR:
This is the Maximum Information Rate of the group.

Group Level 2

Filter:
This is a list of all filters defined to get to the group immediately to the right in the table. It is also a hyper link that will take you to the group filter page. There are two basic ways to add a filter. One is to navigate to the Configuration\Quality of Service\Group Filters page and select the add button. The other is to navigate to Configuration\Quality of Service\Group Filters\Add.

Name:
This is the name given to the group. There are two basic ways to add a group. One is to navigate to the Configuration\Quality of Service\Group page and select the add button. The other is to navigate to Configuration\Quality of Service\Group\Add. The name is a hyper link that group's configuration page. At each level, there is a default group. This is the group where traffic that is not otherwise selected by a filter will go. This group by default will have the name "Default". This name can changed and another group can be selected as the default. The default queue will be designated by the asterisk "*" appended to the name.

CIR:
This is the Committed Information Rate of the group. Setting this to a number higher than the configured parent rate will result in an error.

MIR:
This is the Maximum Information Rate of the group.

Queue

Filter:
This is a list of all filters defined to get to the queue immediately to the right in the table. It is also a hyper link that will take you to the queue filter page. There are two basic ways to add a filter. One is to navigate to the Configuration\Quality of Service\Queue Filters page and select the add button. The other is to navigate to Configuration\Quality of Service\Queue Filters\Add.

Name:
This is the name given to the queue. There are two basic ways to add a queue. One is to navigate to the Configuration\Quality of Service\Queue page and select the add button. The other is to navigate to Configuration\Quality of Service\Queue\Add. The name is a hyper link that queue's configuration page. At each level, there is a default queue. This is the queue where traffic that is not otherwise selected by a filter will go. This queue by default will have the name "Default". This name can changed and another queue can be selected as the default. The default queue will be designated by the asterisk "*" appended to the name.

CIR:
This is the Committed Information Rate of the queue. Setting this to a number higher than the configured parent rate will result in an error.

MIR:
This is the Maximum Information Rate of the queue.

Prio:
This is the priority of the queue. Priorities range from 1 as the highest, to 8 as the lowest.

HC:
This indicates if Header Compression is enabled for the queue. If yes, it could be header or header and payload. Only 31 queues can have HC enabled.

Warning Icon:
If a red exclamation icon appears in an object name, this indicates a configuration warning. These warnings that the configuration will still function but can't meet the specified criteria due to conflicting parameters. If you hover over the icon, text will appear that advises of the configuration conflict. The following messages are possible:

- Too much CIR configured. Increase CIR for this object, or decrease CIR for the children
- Too much CIR configured. Decrease CIR for this object, or increase CIR for the parent.
- Too much MIR configured. Decrease MIR for this object, or increase MIR for the parent.

3.4.3 How to Configure QoS Links

This screen allows you to maintain QoS links. Links correspond to a satellite modem (or a set of satellite modems in 1:1 redundancy) that will be polled to ascertain the current transmit-rate capacity. If no satellite modems are to be polled, then a 'clear sky rate' may be specified. Each link sets a cap on the maximum transmit-rate for all of the 'Groups' that are members of this link.

A list of previously defined links will be displayed. An existing entry may be chosen by clicking on the link name. Other buttons at the bottom of the screen are as follows:

Figure 3-11 FX Series QoS Links Screen

- Add to create a new QoS link.
- Enable / Disable/ Delete one or more links that are selected by the checkbox to the left of the link name.

Links

Name:
This field is a logical name that is used as a reference for the 'Member of' field when 'Groups' are defined. This field must be unique and must be entered.

Enabled:
This selects whether the link definition is enabled or not. If disabled then all filters associated with the groups that are members of this link become inoperative.

Comment:
This field provides a place to store any user defined comment to describe the rationale for this link definition.

Clear sky data rate (kbps):
This is the output data rate that will be used if the FX is unable to read a rate from the modem or if modem polling is disabled.

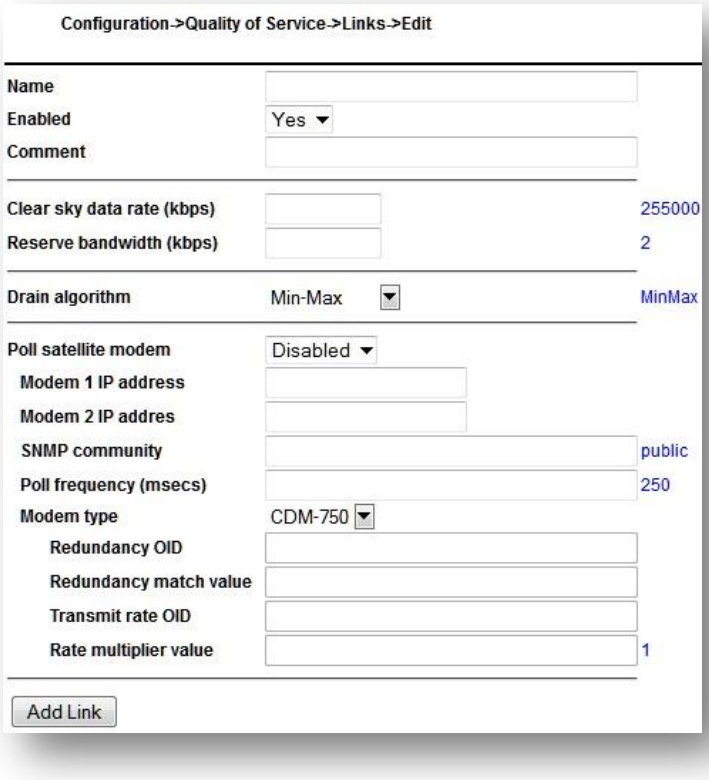

Figure 3-12 F Series QoS Link Edit Screen

Drain algorithm:

This field specifies the drain method for scheduling outbound packets for all Queue definitions.

Strict Priority:
In the Strict Priority Drain algorithm, higher priority queues are depleted before lower priority queues pass traffic. Traffic is capped at the link rate.

Min-Max:
The Min-Max drain algorithm is priority based. If there is enough data rate available, each queue will receive their respective CIR. If there is not sufficient data rate to satisfy all requested CIRs, then traffic will be dropped starting with the lowest priority queue and progressing through the queues in ascending priority until the requested CIR is met.

When traffic is dropped from queues with the same priority, then each of the equal queues will have traffic dropped proportionally. Once all requested CIRs are met, if there is additional data rate that can be filled, it will be allocated to the queues in order of priority starting with the highest. Each queue is given additional data rate up to the requested rate, or MIR, whichever is lowest. If there are queues at the same priority, they are granted additional data rate proportionally.

Poll satellite modem:

Enable or Disable polling of the satellite modem over the FX management interface. The default value is 'Disabled'. For 1:1 redundancy, both a primary and secondary modem may be specified, in which case the transmit data rate capacity of the modem which identifies itself as the "active" will be used for bandwidth allocation calculations. The FX assumes that both the primary and secondary modems are configured with the same SNMP community and are the same modem type.

Primary 1 IP address:
This is the IP address of the primary satellite modem. (This field must be entered if polling is enabled).

Secondary 2 IP address:
For 1:1 redundancy, this is the IP address of the secondary modem. If the secondary modem responds, then its rate will be used. If a second modem is specified, then the FX assumes that both are configured with the same SNMP community.

SNMP community:
This is the read-only community of the satellite modems.

Poll frequency (msecs):
This is the number of milliseconds that the FX waits between polls to ascertain the data rate from the modem. (Default: 250)

Reserve bandwidth (kbps):
The data rate that the FX delivers data is the Ethernet frame rate. This parameter sets the amount of bandwidth that will be held in reserve. The FX will deliver data at the rate read from the modem minus this rate. This will allow for any mismatch between the Ethernet frame rate, and the rate that the modem reports.

Modem type:
This allows you to set the CEFD satellite modem type. (Default: CDM-750). If 'Other' is chosen then the following fields become accessible to ascertain the transmit data rate and redundancy mode:

Redundancy OID:
The OID of the SNMNP query string used to ascertain whether the modem is in 'active' or 'standby' mode.

Redundancy match value:
Specifies a string to compare against to indicate the matching response to the 'Redundancy mode OID' to determine if the satellite is in active mode.

Transmit rate OID:
The OID of the query string used to ascertain the current transmit data rate capacity. The response to this query must be a numeric value.

Rate multiplier value:
The response to the 'Transmit rate OID' query is multiplied by this value to determine the data rate in bits per second. The default multiplier value is 1.

Redundancy

States (*As shown in the QoS Real-Time Monitor)***:**

ONLINE:
Active modem, link speed will be determined by this modem

OFFLINE:
This modem is considered a backup; this state is only entered when the value polled from redundancy OID does not match the "Redundancy match value"

NOT RESPONDING:
Modem does not respond to SNMP polls

The FX will continually poll all modems specified, transmit link speed will be determined by the "ONLINE" modem. The online modem is determined as follows:

CASE 1: Only one modem is specified, this modem reports one of two states "ONLINE" or "NOT RESPONDING"

CASE 2: Two modems are specified, polling the "Redundancy OID" matches the "Redundancy match value".

The FX assumes that a modem is unresponsive (NOT RESPONDING) if either of the following cases is satisfied:

CASE 1: The modem has not yet been polled.

CASE 2: After 20 consecutive poll failures the modem status will change to "NOT RESPONDING" in the QoS Real-Time Monitor and report a data rate of zero Kbps. If the modem fails to respond in one second, it is considered an unsuccessful poll. Unsuccessful polls will have an aggregate poll time of one polling cycle plus two seconds for the ONLINE modem and one second for the OFFLINE modem.

The FX will attempt to determine data rate and redundancy state (if requested) over the management interface. Modems polling will continue regardless of state, state will change when a modem changes state or begins responding to SNMP polls.

NOTE: If both modems report "offline" or "not responding" the output QoS rate will become the clear sky data of the LINK.

How to Verifying ACM QoS Connectivity

To verify that the FX is correctly ascertaining the data rate from the modem, navigate to “Status->View Current Status->ACM QoS->By Modem” This causes the current data rate from all modems which were defined to be displayed, as well as tallies of successful and unsuccessful poll operations.

```
 IP                      State        RateKbps  Spolls  Fpolls Link Guid
 ----------------        -----        -------  ------  ------ --------------------------------
#
# QoS Queues
#
There are 5 qosqs, of which 0 have packet compression enabled
                  Name   Pri    Max Rate (bps)  Cur Rate (bps)           Drops  PktOnQ       BytesSent        PktsSent    IngressBytes     IngressPkts            Group
                  ----   ---    --------------  --------------           -----  ------       ---------        --------    ------------     -----------            -----
-- Internal Signals --     0       255,000,000               0               0       0               0               0               0               0
          Signaling Q     1            20,000             504               0       0         482,892           8,050               0               0       vGjLuXERLkc
       Realtime media     2         2,000,000           6,320               0       0       6,071,173          94,262               0               0       vGjLuXERLkc
           Test Queue     2        20,000,000           2,464               0       0       2,449,896           7,755               0               0       vGjLuXERLkc
            -Default-     8       255,000,000               0               0       0               0               0               0               0       vGjLuXERLkc
#
# Filters
#
There are 4 filters
             Filter Name                AssignedToQueue           Prot.             Matches    Criteria
             -----------                ---------------           -----             -------    --------
                    arps                    Signaling Q             ARP                   0
          RealTime media                 Realtime media             TCP                   0    SrcPorts=554,5505,5000-5010,64064,1935,1111,8134
           Test Filter 2                     Test Queue             UDP                   0
               -Default-                  Realtime media             Any                   0
```

Figure 3-13 FX Series ACM QOS Status by Modem Report

3.4.4 How to Configure QoS Groups

This function allows you to maintain QoS groups. QoS groups allow you to group multiple QoS queues (or sub-groups of QoS queues) into the same link. A common packet compression peer address is also specified on a per QoS group basis. An existing entry may be chosen by clicking on the group name. Clicking on a column header will sort the list based upon the contents of the column. Other buttons at the bottom of the screen are as follows:

- Add to create a new QoS group
- Enable/ Disable/Delete one or more groups that are selected by the checkbox to the left of the group name.

NOTE: If disabled, then this group and all queues or sub-groups that are members of this group are not used.

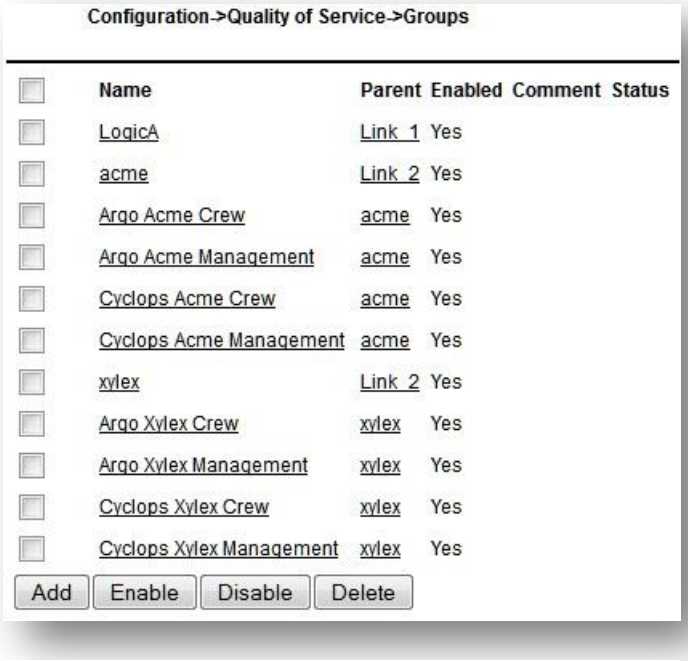

Figure 3-14 FX Series QoS Groups

Name:
This field is a logical name that is used as a reference for the 'Member of' field when 'Groups' are defined. This field must unique and must be entered.

Enabled:
This selects whether the group is enabled or not. If disabled then this group and all queues or sub-groups that are members of this group are not used.

Comment:
This provides a place to store any user defined comment to describe the rationale for this group.

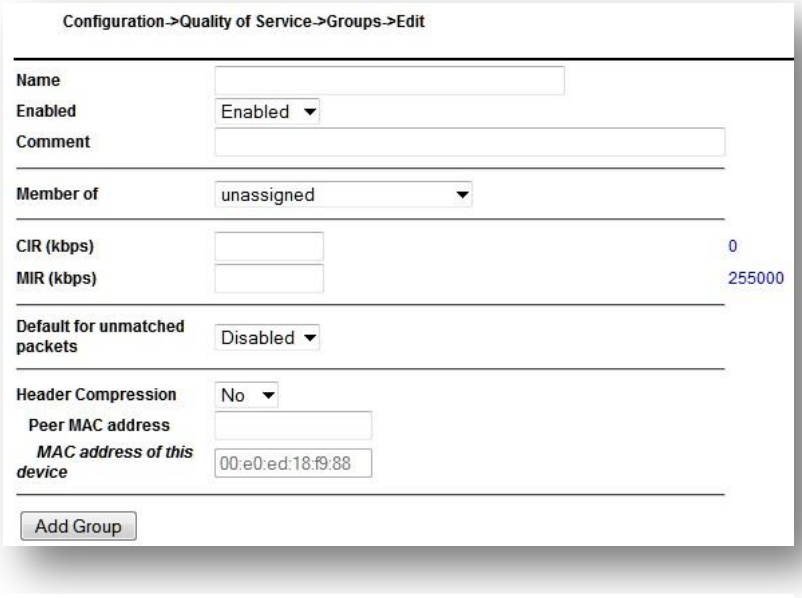

Figure 3-15 FX Series QoS Group Edit Screen

Member of:
This selects the link or higher-level group that this group is a member of.

CIR (kbps):
This specifies the "Committed Information Rate" in kbps (1000 bits per second). The range is 0 up to the licensed rate. If the FX WAN optimization feature is not licensed then up to 700000 can be specified. The default is 0. This field is disabled if "Strict Priority" was configured as the drain algorithm.

MIR (kbps):
This specifies the "Maximum Information Rate" in kbps (1000 bits per second). The range is 0 up to the licensed rate. If the FX WAN optimization feature is not licensed then up to 700000 can be specified. If 0 is specified, some packets may still be sent at a very low rate, to inhibit all traffic then a "DROP" queue should be defined. The default is the max licensed rate. This field is disabled if "Strict Priority" was configured as the drain algorithm.

Default for unmatched packets:
If set, then this is the group that packets which have not matched any of the group filters at this level will be directed to. If no group is designated as 'Default', then unmatched packets will be directed to an arbitrary group.

Header compression:
If set to 'yes' then the traffic associated with queues that are members of this group, which also have packet compression enabled, is aggregated and encapsulated in Ethernet frames and sent directly to the MAC address specified in the 'Peer MAC address' field. You can subsequently enable/disable header compression on a per QoS queue basis. The default setting is 'No'.

Peer MAC address:
This setting specifies the MAC address of the WAN Interface of the FX which will receive the encapsulated compressed and aggregated packets. The MAC address must be specified in format xx:xx:xx:xx:xx:xx where each 'xx' is a hex digit. If no peer MAC address is entered in a second level group, then the inherited peer MAC address from the first level group is used, this address is shown in blue.

MAC address of this device:
This is a 'display-only' field that shows the MAC address of the WAN interface of this device which can be copied and pasted when configuring the peer.

Group Filters:
This is a read-only list of group filters that are currently assigned to this group.

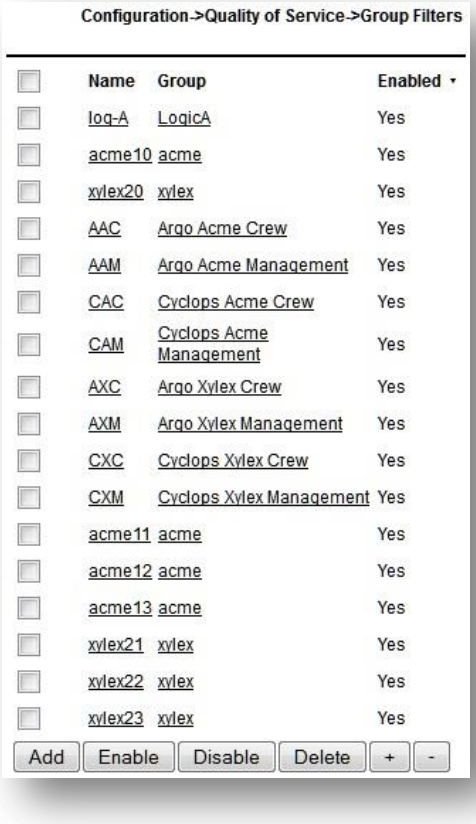

Figure 3-16 FX Series QoS Group Filters

3.4.5 How to Configure QoS Group Filters

Group filters are rules which classify traffic and assign it to one of the QoS groups. A list of previously defined group filters is displayed in order of rank. Clicking on a column header sorts the list based upon the contents of the column. Other buttons at the bottom of the screen are as follows:

- Add:
  Create a new QoS queue
- Enable/Disable/Delete:
  Enable/Disable/Delete one or more QoS queues that are selected by the checkbox to the left of the queue name.
- All filters that reference the deleted queues become unassigned.
- '+' – Increase the rank of a group filter.
- '-'– Decrease the rank of a group filter.

If the user checks a box and clicks “Add” the rank is set based upon the rank of the selected item. Each of those actions will bring up a ‘Group Filter screen. The rules in the pick-list are sorted based upon Rank. “+" and "-" buttons at the bottom of the pick list move selected rules up or down in rank, multiple items can be selected for the rank adjustment. An existing entry may be chosen by clicking on the group filter name.

Clicking on the Queue “Name” will allow you to modify that Queue. Each of those actions will then lead to a Queue definition screen. The Default Queue is pre-defined and may only be deleted if other queues exist. Its initial priority is 8, the lowest priority.

Name:
This identifies the customer/function of the filter. This field must be entered and must be unique.

Enabled:
This selects whether the group filter is enabled or not.

Group:
Selects which group that traffic that matches the criteria specified in this group filter should be directed to.
This field may be left unassigned during definition, but must be eventually be assigned for them to take effect.

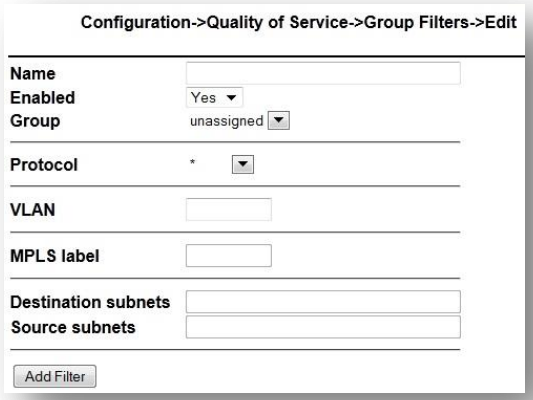

Figure 3-17 FX Series QoS Group Filters Edit Screen

Protocol:
Select between *, IP, MPLS. The default is * (all protocols). If ‘IP’ is selected then the ‘MPLS label’ field is inaccessible. If ‘*’ is selected then all fields except ‘VLAN’ are inaccessible. If ‘MPLS’ is selected, then the ‘Destination and Source subnets’ fields, are inaccessible.

VLAN:
Enter either 0 or a VLAN ID between 2 and 4094. Only one may be selected. 0 indicates untagged traffic as the selection criteria. The default is any VLAN.

MPLS label:
If MPLS was selected as the protocol then a decimal value between 0 and 1048575 may be entered. If no value is entered then all MPLS labels will match the filter criteria. If there are multiple MPLS labels, the filter will only match the first label encountered in the packet.

Destination subnets:
This is specified in CIDR format. Multiple subnets may be separated by a comma. The default is ‘*’ (any subnet). Acceleration tunnels utilized by FX WAN Optimization may not maintain the original application destination address, therefore this field should not be used when classifying FX Wanop traffic unless an application policy is defined to prevent tunnel sharing between different destination subnets.

Source subnets:
This is specified in CIDR format. Multiple subnets may be separated by a comma. The default is ‘*’ (any subnet)

3.4.6 How to Configure QOS Queues

These settings allow you to manage QoS Queues. A list of previously defined queues will be displayed in order of priority. Clicking on a column header will sort the list based upon the contents of the column. An existing entry may be chosen by clicking on the queue name which leads to the Queues Add/Edit Screen. Other buttons at the bottom provide these functions:

- Add – Create a new QoS queue
- Enable/Disable/Delete – Enable/Disable/Delete one or more QoS queues that are selected by the checkbox to the left of the queue name. All filters that reference the deleted queues become unassigned. The Default Queue is always defined. It is initial priority is 8, the lowest priority and is associated with the Default Filter.

Figure 3-18 FX Series QoS Queues

Name:
This field is a logical name that is used as a reference when queue filters are defined. This field must be entered.

Enabled:
This selects whether the queue is enabled or not.

Group:
This specifies the group that this queue is a member of.

Scheduling Discipline:
This specifies the order in which packets in this queue are scheduled for transmission. The choices are:

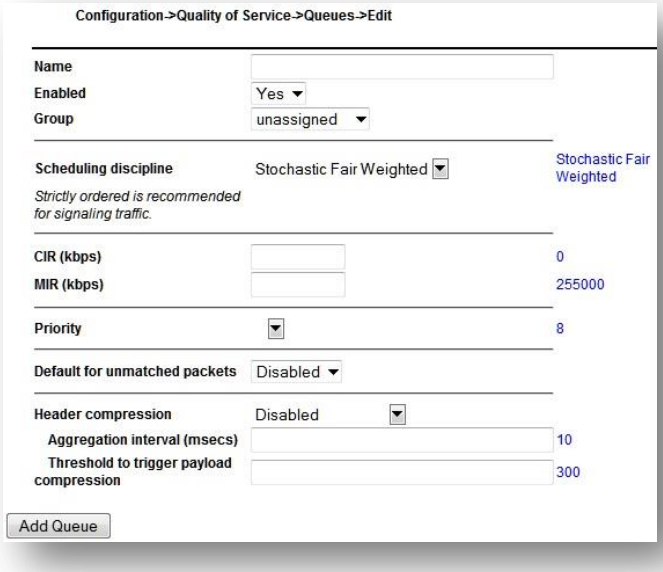

Figure 3-19 FX Series QoS Queue Edit Screen

Stochastic Fair Weighted
The scheduler will attempt to evenly distribute outbound traffic based on hashing the source and destination addresses.
This is the default choice and prevents one traffic flow from consuming all bandwidth assigned to this queue at the expense of other flows assigned to this queue.
When there is high link congestion this method may introduce miniscule delays.

Strictly ordered
Packets are sent in the order that they are received. This may be a good choice for signaling traffic where there can be absolutely no disruption in packet transmission.

DROP
Packets directed to this queue are discarded.

CIR (kbps):
This specifies the "Committed Information Rate" in kbps (1000 bits per second). The range is 0 up to the licensed rate. If the FX WAN optimization feature is not licensed, then up to 700000 kbps can be specified. The default is 0. This field is disabled if "Strict Priority" was configured as the drain algorithm.

Note: To minimize jitter, set the CIR for high priority traffic high enough to accommodate peak usage requirements.

MIR (kbps):
This specifies the "Maximum Information Rate" in kbps (1000 bits per second). The range is 0 up to the licensed rate. If the FX WAN optimization feature is not licensed then up to 500000 can be specified. If 0 is specified, some packets may still be sent at a very low rate, to inhibit all traffic then a "DROP" filter should be defined. The default is the max licensed rate. This field is disabled if "Strict Priority" was configured as the drain algorithm.

Priority:
This is the drain priority for the queues. Classes of equal priorities will be treated the same, with rates split proportionally between them. The minimum value (highest priority) is 1; the maximum value (lowest priority) is 8.

Default for unmatched packets:
If set, then this is the queue that packets which have not matched any of the queue filters in a group will be directed to. If no queue is designated as 'Default' for a group, then unmatched packets will be directed to the lowest priority queue in the group.

Header compression:
This enables/disables header compression and packet aggregation for this queue. For header compression to occur, you must also configure header compression in the group this queue is a member of. If header compression is enabled then WAN optimization features such as caching can't be performed on traffic associated with this queue. When configuring header compression, you can choose to perform header compression, or both header and payload compression.

Aggregation interval (msecs):
This specifies the packet aggregation flush interval in milliseconds if packet compression is enabled. The minimum and default value is 10 msec. The maximum is 1000 msec.

Threshold to trigger payload compression:
If 'header and payload' compression is selected, this value specifies the minimum payload size to trigger payload compression. The default is 300 bytes.

Filters:
This is a read-only list of filters that are currently assigned to this queue.

3.4.7 How to Configure QoS Queue Filters

The fields on this screen dictate how traffic will be directed to a queue

Name:
This summarizes the customer/function of the filter. This field must be entered and must be unique.

Enabled:
This selects whether the queue is enabled or not

Queue:
Selects which queue that traffic that matches the criteria specified in this filter definition should be directed. This field may be left blank during definition, but must be eventually be assigned.

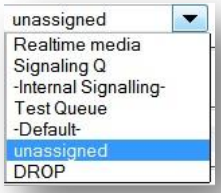

Protocol:
Select between * / IP/ UDP / TCP / FTP / ARP / ICMP / MPLS / SCTP / PTPv1 / PTPv2. Only one choice may be selected. The default is * (all protocols).

If **FTP** is selected, the FX automatically tracks the data ports associated with FTP transfers by monitoring the activity on the FTP control port, which is defaulted to ports 20.21 upon initial selection. To maintain Multicator transmitter functionality when performing QOS on non-standard FTP ports an additional QOS filter must be created for FTP that utilizes port 21. The FTP control port may be changed.

If **MPLS** is selected, only MPLS label, MPLS experimental bits, VLAN and VLAN priority fields may be selected as filter criteria, otherwise these fields are disabled.

If **ARP, PTPv1, PTPv2, or FTP** is selected, then this filter can't be assigned to a queue for which packet compression is enabled.

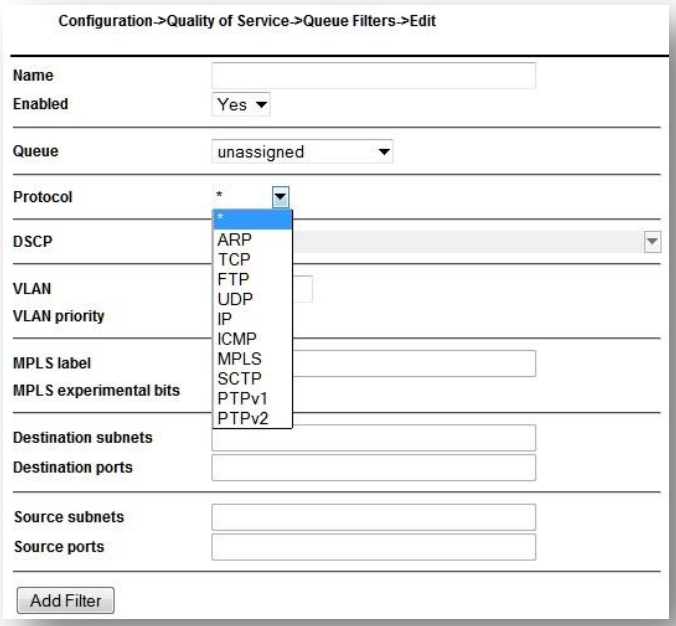

Figure 3-20 FX Series QoS Queue Filter Edit Screen

If **PTPv1 or PTPv2** is selected, only DSCP, VLAN, VLAN priority, Destination subnets, and Source subnets may be selected. If PTPv2 is selected, in addition to PTPv2 running over UDP, the filter will also check for packets on ethertype 88f7, in which case IP specific options do not apply.

If **SCTP** is selected, only DSCP, VLAN, VLAN priority, Destination subnets, and Source subnets may be selected.

DSCP:
Select one of the choices from the pull-down menu of DCSP choices. Only one choice may be selected). The default is * (any)

VLAN:
Enter either 0 or a VLAN ID between 2 and 4094. Only one may be selected. 0 indicates untagged traffic as the selection criteria. The default is * (any VLAN)

VLAN Priority:
Choose between 'any' or a priority value between 0 and 7. Only one choice may be selected. The default is * (any VLAN priority)

MPLS Label:
If MPLS was selected as the protocol then a decimal value between 0 and 1048575 may be entered. If no value is entered then all MPLS labels will match the filter criteria. If there are multiple MPLS labels, the filter will only match the first label encountered in the packet.

MPLS experimental bits:
If MPLS was selected as the protocol then you can choose a value between 0 and 7 as match criteria. If '*' is chosen then the filter does not use the experimental bits in the match criteria, otherwise all bits must match exactly.

Destination Subnets:
This is specified in CIDR format. Multiple subnets may be separated by a comma. The default is '*' (any subnet). Acceleration tunnels utilized by FX WAN Optimization may not maintain the original application destination address, therefore this field should not be used when classifying FX WANOP traffic unless an application policy is defined to prevent tunnel sharing between different destination subnets.

Destination Ports:
Port ranges can be specified by either entering the lowest port followed by '-', followed by highest port; or multiple ports may be entered separated by comma. The destination port is relative to the FX.

Source Subnets:
This is specified in CIDR format. Multiple subnets may be separated by a comma. The default is '*' (any subnet)

Source Ports:
Port ranges can be specified by either entering the lowest port followed by '-', followed by highest port; or multiple ports may be entered separated by comma. The source port is relative to the FX. Acceleration tunnels utilized by FX WAN Optimization do not maintain the original application source port; therefore this field should not be used when classifying FX WANOP traffic.

NOTE: When specifying multiple fields as selection criteria the choices are logically "ANDed" when formulating a match. Multiple selections within a field are logically "ORed" when formulating a match.

3.5 FX Series Multicator

The Multicator is a set of three components of the FX Series which allow controlled reliable content distribution via multicast. These components are as follows:

Multicator Controller (MC)

Within a Multicator deployment, there is one, and only one, FX SERIES ADC appliance which must be designated as an MC.

An MC is the central point where Multicator configuration parameters are stored. The MC ensures that only one multicast transmission is occurring at a time. The MC ensures that if there is a network outage, a multicast which was in progress will resume from the point where the outage occurred. The MC maintains a central log of all Multicator events.

Multicator Transmitter (MT)

This component actually performs the multicast of the content after checking with the MC. Any FX Series ADC or FX Series Remote can function as a Multicator Transmitter (MT) if the license is enabled. The MT employs world renowned open source technology to reliably deliver content via multicast.

NOTE: Any FTP program can be used to upload content to the MT.

Multicator Receiver (MR)

This component receives the content which is transmitted by the MT. Upon completion of a successful reception of new content, the MR uploads this content to a local FTP server.

NOTE: Any combination of MC, MT, and MR may be configured on the same appliance as long as there exists one, and only one MC in the Multicator deployment.

Theory of Operation

A powerful new content distribution system can now be set up with the separately licensed “Multicator” feature. This feature allows a user to upload a file to an FX Series device via FTP. The file is then reliably multicast to a group of receivers. The receivers then upload the content to a local FTP server. The Multicator employs the “Content Distribution Control Protocol” (CDCP) to ensure that only one multicast transmission is in progress.

Sequence of Events

1. Files are deposited on the Transmitter (Sender) using a standard FTP client.
2. The Transmitter then notifies the Controller that it has data to send and is granted permission to reliably multicast the data across the WAN.
3. Under the direction of the Controller, the Transmitter establishes a reliable multicast connection to the Receivers.
4. The Transmitter sends the files to each of the Receivers.
5. Each Receiver sends an acknowledgment of receipt to the Controller.
6. Each receiver uses FTP to send the files to the respective server.

NOTE: In order to configure Multicator, at least one In-Path interface must be defined and operational.

3.5.1 Multicator Settings

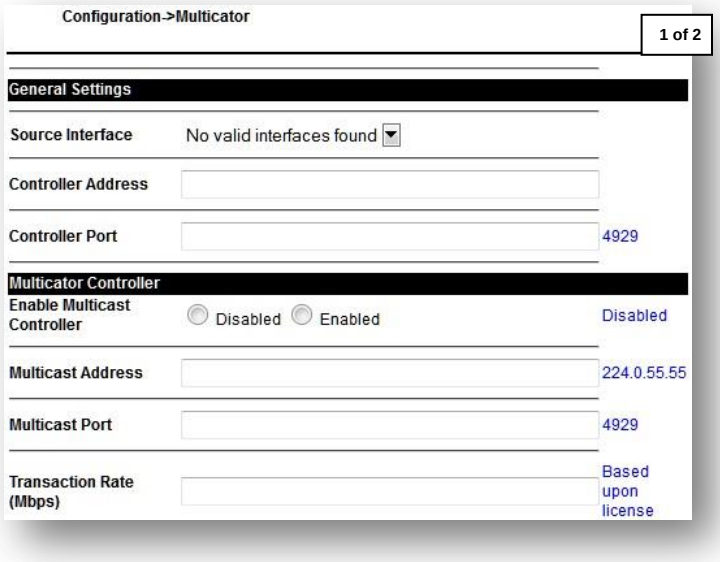

Figure 3-21 FX Series Multicator General/Controller Edit Screen

General Settings

Source Interface:
For transmitters and receivers, this is the interface used when communicating with the controller. For the controller, this is the interface used when transmitters and receivers communicate

Note: In routed mode this should always be the interface designated as the WAN interface

Controller Address:
This is the address of the controller that the transmitter/receiver will communicate with. If this device is a controller and a transmitter or receiver this address should match the "Source Interface" field.

Controller Port:
The port that the controller will use to communicate with transmitters and receivers, if this appliance is a transmitter or receiver it is the port used to communicate with the controller

Multicator Controller Settings

Enable Multicast Controller:
This setting enables the multicast controller on this appliance. Only one controller should be enabled on a network. The default is "Disabled".

Multicast Address:
This is the multicast IPv4 address that will be used to transfer files via reliable multicast. This address is communicated to the transmitters and receivers. The default value is 224.0.55.55.

Multicast Port:
This is the multicast port that will be used to transfer files via reliable multicast. This port is communicated to the transmitters and receivers. The default value is 4929.

Transaction Rate:
This is the max speed that a multicast transmitter will transmit a file. The default value is based off the license.

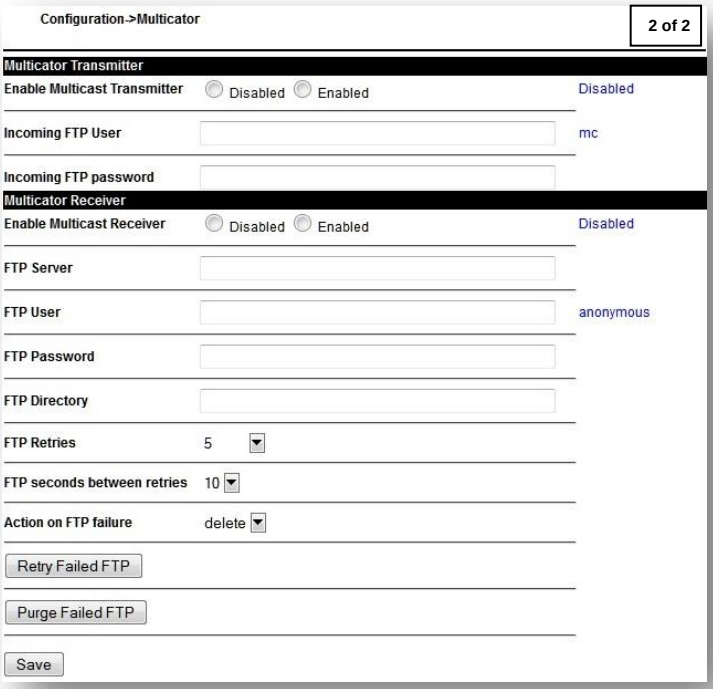

Figure 3-22 FX Series Multicator Transmitter/Receiver Edit Screen

Multicator Transmitter Settings

Enable Multicast Transmitter:
This setting enables the multicast transmitter on this appliance.

Incoming FTP User:
This is the user name that must be used when content is uploaded to the FX appliance. The default is "mc".

Incoming FTP Password:
This is the password that must be used when content is uploaded to the FX appliance. The default password is “comtech”.

Multicator Receiver Settings

Enable Multicast Receiver:
This setting enables the multicast receiver on this appliance.

FTP Server:
This is the IP address of the FTP server into which newly received content will be fanned-out.

FTP User:
This is the user name used when transferring new content to the FTP server.

FTP Password:
This is the password which will be used when transferring new content to the FTP server.

FTP Directory:
This optional parameter is the directory where new content will be transferred. If this directory does not already exist, it will be created before the file is transferred. The default is none.

FTP Retries:
This is the number of times to attempt to send the file to the FTP server before both deleting it and moving on to the next or keeping it and moving on to the next. The receiver will attempt to resend all failed files when the acceleration service is restarted or the "Retry Failed FTP" button is pressed.

FTP seconds between retries:
This pull-down allows you to select number of seconds which will elapse between each attempt to send the file to the FTP server.

Action on FTP failure:
This specifies the action to take if the FTP retry limit is exceeded. If delete is selected the file will be delete, otherwise the file will be stored until the "Retry Failed FTP" button is clicked or the acceleration service is restarted.

Retry Failed FTP:
This button will cause any files that failed FTP transfer to be re-sent to the FTP server. This action is only valid if "Action on FTP Failure" is set to "keep"

Purge Failed FTP:
This button will delete any files which are pending to be re-sent to the FTP server. This action is only valid if "Action on FTP Failure" is set to "keep".

Note: If disk utilization reaches 80%, a purge of all files that failed FTP transfer will automatically occur.

3.5.2 How to set the Multicator General Configuration

1. Log into the browser interface of the appliance.
2. Click the Configuration link; Click the Multicator link.
3. Select the appropriate interface in the "Source Interface" field in the "General Settings" section. In routed mode this should always be the WAN facing interface.
4. Enter the IP address of the controller for the transmitter and receiver in the "Controller Address" field. If this is the controller enter the IP address in the "Source Interface" field.
5. Enter the port the transmitter/receiver will communicate with the controller on. If this appliance is also a controller, this is the port it will listen on.

3.5.3 How to set the Multicator Controller Configuration

1. Log into the browser interface of the appliance.
2. Click the Configuration link; Click the Multicator link.
3. Click the enable radio button in the “Enable Multicast Controller” field in the "Multicator Controller Section"
4. Enter the multicast IP address you wish to use in the “Multicast Address” field. Default: 224.0.55.55
5. Enter the port you wish to use for multicast in the “Multicast Port” field. Default: 4929
6. Enter the rate at which data should be transmitted via multicast in the “Transaction Rate” field. Default: The licensed rate of the appliance.

3.5.4 How to set the Multicator Transmitter Configuration

1. Log into the browser interface of the appliance.
2. Click the Configuration link; Click the Multicator link.
3. Click the enable radio button in the “Enable Multicast Transmitter” field in the "Multicator Transmitter" section.
4. Enter the username used in FTP file submissions to the transmitter in the "Incoming FTP user" field. Default: mc
5. Enter the password used in FTP file submissions to the transmitter in the “Incoming FTP Password” field. Default: comtech

3.5.5 How to set the Multicator Receiver Configuration

1. Log into the browser interface of the appliance.
2. Click the Configuration link; Click the Multicator link.
3. Click the enable radio button in the “Enable Multicast Receiver” field in the "Multicator Receiver" section.
4. Enter the controller port of the Multicator controller. Default: 4929
5. In the “FTP Server” field, enter the IP address of the FTP server that will receive the file delivered to the receiver via multicast.
6. In the “FTP User” field, enter the user name for the FTP server that will receive the file delivered to the receiver via multicast.
7. Default: anonymous
8. In the “FTP Password” field, enter the password for **t**he FTP server where the file received via multicast will be placed.
9. Default: no directory, file deposited in FTP root
10. In the “FTP Directory” field, enter the directory on the FTP server where the file received via multicast will be placed.
11. In the "FTP Retries" field select the number of times the receiver should attempt to deliver a file to the FTP server before abandoning the file transfer. Default: 5
12. In the "FTP seconds between retries" field select the number of second between FTP retry attempts. Default: 10
13. In the “Action on FTP failure” field, select the action to be taken on the file if the file transfer fails and all retry attempts have been exhausted. Default: Keep.

3.6 Redundancy

This section allows you to configure 1:1 redundancy with fail over in which a secondary FX, with the same configuration as the primary FX, polls for the existence of the primary FX, and takes over its non-management IP addresses when the primary does not respond to the poll. When the primary FX comes back up, the secondary FX will relinquish the IP addresses. This section also allows you to configure shared configurations between members of an appliance pool. This is useful to synchronize configurations in a WCCP cluster.

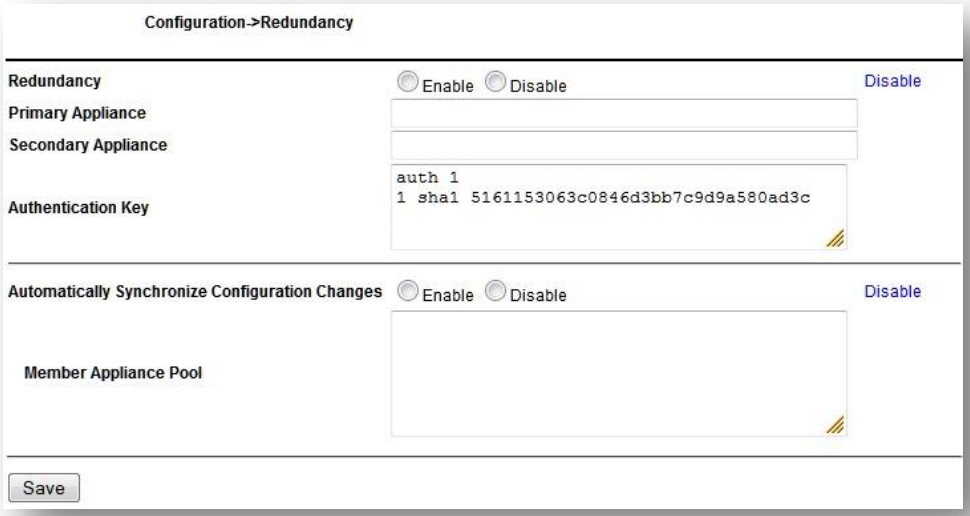

Figure 3-23 FX Series Redundancy Edit Screen

3.6.1 Redundancy Configuration Settings

Redundancy
This must be enabled if either the 1:1 Redundancy with fail over or "Automatically synchronize configurations" feature is required.

Primary Appliance:
This is the host name of the primary (master) FX.

Secondary Appliance:
Enter the 'short' host name or IP address of the secondary FX that will engage if the primary FX becomes inoperable. The secondary appliance continually polls the primary appliance and if the primary appliance does not respond then the secondary appliance asserts control over the realm of IP addresses that external clients connect to. When the primary appliance becomes operational again, the secondary will relinquish control of these IP addresses. This parameter is not required if only synchronizing configuration changes to member pool appliances is being configured.

Authentication Key:
By default, this key is generated automatically. For 1:1 Redundancy with fail over, the secondary appliance's key must match the primary appliance's key. This requires manually copying the primary appliance's key and pasting it into this field on the secondary appliance (after deleting the secondary appliance's generated key).

Automatically Synchronize Configuration Changes:
This field must be enabled for configuration synchronization. Any time a configuration change is applied using the browser administration interface; the change is immediately synchronized with the Secondary Appliance and/or members of the "Member Appliance Pool". The 'Configuring Key-Exchange' procedure below must be performed.

Member Appliance Pool:
If the "Automatically Synchronize Configuration Changes" is enabled, then this field defines the list of host names or IP addresses, separated by commas, of the appliances that will share the same configuration files as the primary appliance. The devices defined in the "Member Appliance Pool" share their configurations and require a valid SSH key to be exchanged with the "primary" appliance.

(See section titled "Configuring Key-Exchange" below).

Save Button:
Clicking on 'Save' will commit the fields on this form to disk. If this is the Initial configuration of high-availability the appliance must be rebooted after the 'Save' completes.

3.6.2 How to Configure Key-Exchange

In order for the FXs to securely communicate with each other in an automated fashion it is necessary to use the FX-Series Appliance Manager" via SSH to configure common cluster authentication keys.

To configure the key exchange between the primary and secondary, log into the FX with "ssh" to access the "FX-Series Appliance Manager" and perform the following sequence on the primary FX:

1. Choose1 "Configure Appliance"
2. Then choose 2 "Configure Passwords"
3. Then choose 2"Configure Redundancy Cluster Key"
4. Enter the IP Address or short host name (as specified on the 'Management->Host Settings' page) of the peer appliance. If prompted with '/root/.ssh/id dsa already exists. Overwrite? (y/n)' Enter 'y'.
5. On prompt: Are you sure you want to continue connecting (yes/no) enter 'yes"
6. On the password prompt enter "comtech"

Repeat this for the secondary and/or each entry in the Member Appliance Pool.

NOTE: For 1:1 Redundancy with failover configurations, the primary appliance and secondary appliance entries must be associated with the auxiliary port. A 'short' host name is required. These may be specified via the DNS server or by configuring the local host table. (See Configuration->Host Settings)

ARP Considerations:
When the FX performs the IP take over it will send out a gratuitous ARP so that other routers are notified of the take-over.

3.6.3 How to Configure 1:1 Redundancy with Fail Over

- Given an FX-4000/FX4010 ADC named '**PrimAdc**' installed and operating in "in-line mode" eth2/eth3) using eth0 as a management port.
- Add a new FX-4000/FX4010 ADC appliance named **'Adc2nd'** for 1:1 redundancy with fail over.

1. Cable the ADC appliances:
   a) Connect **PrimAdc/Adc2nd** auxiliary (eth1) ports with cross-over cable
   b) Remove **PrimAdc** eth3 port connection (Wan) and plug it into **Adc2nd** eth3 (Wan) port
   c) Use cross-over cable and connect **PrimAdc**'s eth3 to **Adc2nd** eth2 (LAN) port.

2. Browse to Adc2nd: **Configuration->Host Settings:**
   a) Set 'Host name' to Adc2nd
   b) In 'Host File Entries' add the following lines:
   10.1.1.10 PrimAdc.com PrimAdc
   10.1.1.11 Adc2nd.com Adc2nd
   Click 'Save'

3. Browse to Adc2nd: **Configuration->Management ->Network interfaces**
   Set Auxiliary Interface
   'Ip v4Address' = 10.1.1.11 'Subnet mask' = 255.255.255.0
   Click 'Save

4. Browse to **PrimAdc: Configuration->Management ->Network interfaces**
   Set Auxiliary Interface
   'Ip v4Address' = 10.1.1.10 'netmask' = 255.255.255.0
   Click 'Save'

5. Browse to **PrimAdc: Configuration->Host Settings**
   In 'Host File entries' add the following line
   10.1.1.11 Adc2nd.com Adc2nd
   10.1.1.10 PrimAdc.com
   Click 'Save'

6. On **PrimAdc** browse to: **FX Series Application Delivery Controller-> Configuration->Redundancy**
   a) Enable 'Redundancy'
   b) Set 'Primary Appliance' to **PrimAdc**
   c) Set 'Secondary Appliance' to **Adc2nd**
   d) Set "Automatically Synchronize Configuration Changes" to "Enabled"
   Click 'Save**'** and refresh the browser screen.
   Should see "Authentication Key" similar to:
   auth 1
   1 sha1 0509160a630240f400ec5e389c942422

   The 'Save' action will synchronize **PrimAdc**'s configuration with **Adc2nd**.

7. On Adc2nd browse to:
   a) FX Series Application Delivery Controller-> Configuration->Redundancy

   Verify that 'Authentication Key' matches that shown on the **PrimAdc.**
   Note: If the Keys do not match, copy and paste the Authentication Key from **PrimAdc** to **Adc2nd** then click 'Save' on **Adc2nd.**

   NOTE: It is necessary to reboot both appliances for Redundancy service to run

3.6.4 How to Synchronize Configurations in a WCCP Cluster

When multiple FX devices are functioning in a WCCP cluster, each device has unique network settings, therefore only selected configuration settings are synchronized between the members of the cluster. The settings which are synchronized include only the following:

- HTTP Application Policies
- L5 Application Policies
- Authorization Realms
- QoS Queue Definitions
- QoS Filters
- QoS Links
- QoS Groups
- QoS Group Filters

Of the above, HTTP application policy synchronizations take effect immediately, throughout the cluster, unless the HTTP policy references a newly defined authorization realm. If an HTTP application policy references a modified realm then a restart of the acceleration service may be required for the intended change to take effect. The other settings require a restart of the acceleration service to take effect.

4 FX Series ADC General Settings

This chapter discusses the General Settings of the FX Series ADC appliances.

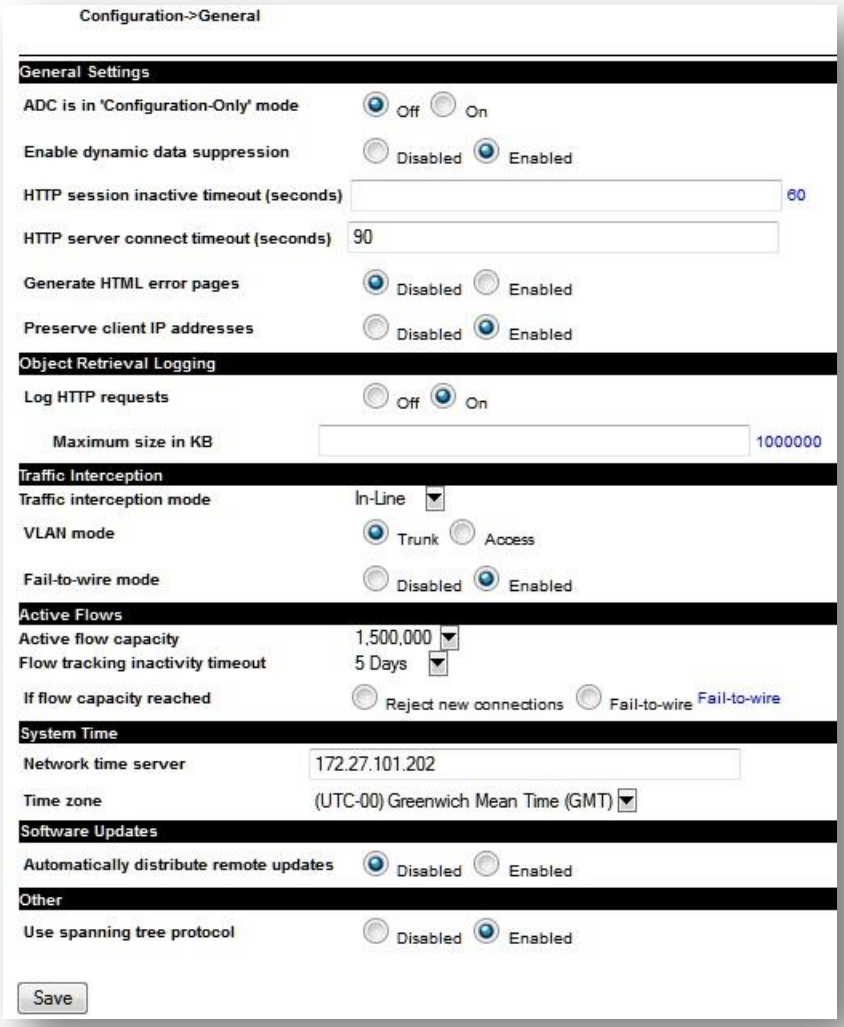

Figure 4-1 FX Series ADC General Edit Screen

FX Series ADC in 'Configuration-Only' Mode:
This setting is useful if you are in the process of configuring your FX Series ADC while the unit is networked in-line. In "Configuration-Only" mode the in-line networking card is put into "bypass" mode so that traffic is simply passed through. When you are satisfied that the FX Series ADC is properly configured you can disable this setting. This setting only applies if the 'Traffic interception mode' is set to "In-Line" or "WCCP", this setting is ignored if the 'Traffic interception mode' is set to "routed". The default setting is "Off".

4.1.1 How to Configure FX ADC in 'Configuration-Only' mode:

1. Using a browser, go to the FX's "Operations" page and select "Backup/Restore Configuration". Select the "Backup FX Acceleration and Management Configuration button"
   This button causes all FX configuration files to be stored into a zip file with the name "fxbackup-full_*hostname*_year-month-day-hour-minute.zip". Following this, a dialog will appear so that you can save this file to your desktop.

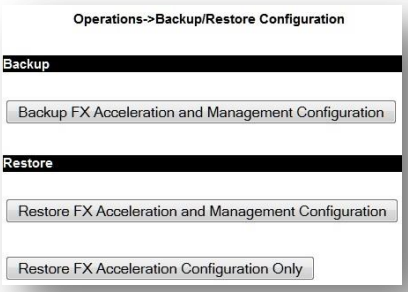

2. Save the backup file to a convenient location.
3. Put the appliance into 'Configuration-Only' mode. The setting can be found on the "General Settings".

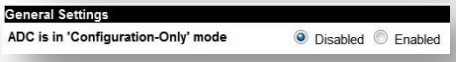

4. After making the changes to the configuration, return the configuration setting to the normal mode "disabled" and verify that there are no issues with the new configuration.
5. If issues are seen, you can return to original acceleration configuration quickly.
   a. Using a browser, go to the FX's "Configuration->General Settings" and change the "Configuration-Only" setting to "Enabled".
   b. Go to the FX's "Operations->Backup/Restore ".
   c. "Choose File" button.
   d. Following this, clicking the "Restore FX Acceleration Configuration Only" button will restore only the acceleration settings and exclude any management settings from the specified zip file. (This is also a useful method for transferring similar acceleration configurations from one FX to another without affecting management settings.)
   e. Load and apply the backup file.
   f. The acceleration appliance will automatically restart.
6. Return to the "General Settings Screen" and return change the "Configuration-Only" setting to "Disabled".

NOTE: If appropriate use the "FX Series Appliance Manager", to re-establish any high-availability keys for the cluster-mates that share a common acceleration server configuration.
Configuring the high-availability cluster keys can be accomplished by selecting "Configure Appliance", and then "Configure Passwords".

Enable Dynamic Data Suppression:
This is a global switch that applies to all traffic processed by this FX Series ADC. If "Enabled" then a cache of data and signatures and byte patterns will be maintained and when possible a signature will be sent instead of a redundant byte pattern. The default value is "Enabled".

HTTP Session Inactive Timeout (seconds):
This setting controls the maximum time that inactive browser sessions are kept open before closing them in order to minimize thread and TCP session resources. The default interval value is 60 seconds.

HTTP Server Connect Timeout (seconds):
This controls the maximum number of seconds that the FX Series ADC will wait for a TCP connection to complete to an HTTP content server before timing-out. After the timeout, a 503 HTTP error code will be returned to the browser that initiated the request. The default value is 20 seconds.

Generate HTML error pages:
This controls whether the FX Series ADC should generate an HTML page describing the problem and identifying the FX Series ADC when it encounters a problem connecting and/or receiving content from a back-end server. The default value is “Enabled”.

Preserve Client IP Addresses:
If enabled, then the FX Series ADC will send the requests to the back-end servers with a source IP address that is the same as the client that the request is on behalf of. Enabling this setting may require a Transparent Bridging configuration. This setting only applies to single-sided optimization. IP source preservation with FX Remotes is specified in the L5 application policies. The default value is “Enabled”.

4.1.2 Object Retrieval Logging

Log HTTP Requests:
Enables logging of URLs for all HTTP web object retrievals in “common log format”. The default setting is “On”.

Maximum Size in KB
This sets the maximum size of the object retrieval log file in kilobytes. When this size is reached a backup is made and the file is reset. The default setting is 1000000 (1 GB).

4.1.3 Traffic Interception

Traffic interception mode:
This is the means by which the FX-ADC will transparently intercept packets. Choose either “In-Line”, “Routed”, WCCP” or “Disabled”. The default is “In-Line”.

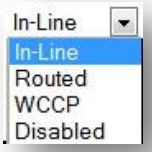

- In-Line mode, the LAN port is eth2 and the WAN port is eth3. Traffic is intercepted as a transparent bridge. If there is a service disruption then the units will “fail-to-wire”. In bridged mode, you must assign an IP address to the in-path interface which bridges the LAN and WAN interfaces. In FX nomenclature, the “WAN” interface is considered to be the interface which is connected to the satellite modems and the LAN interface is connected to the internet or to enterprise servers.
- In routed mode, traffic must be directed to these interfaces by a router. You must assign an IP address to the in-path interface which will receive traffic from the router. There is no “fail-to-wire” capability if there is a service disruption.
- In WCCP mode, traffic is redirected to the FX by a Cisco router via WCCP.
- If “Disabled” then traffic redirection is effectively shut off.

VLAN mode:
This controls how the FX will process VLAN tags. In 'Trunk' mode, the VLAN tags are already embedded in the packets when they are intercepted by the FX. In 'Access' mode, the FX will add tags to untagged traffic. The default setting is ‘Trunk’.

Note: Changing either the ‘Traffic interception mode’ or ‘VLAN mode’ settings will automatically trigger a restart of the acceleration service.

Fail-to-wire mode:
If enabled, the FX will go into bypass mode it is not accelerating traffic or if it is powered off. If disabled, it will not go into bypass mode which will prevent packets from being forwarded through the FX. If the FX is adding the VLAN tags to the traffic, it may be desirable to disable 'Fail-to-wire' mode to prevent untagged traffic from entering a network.
The default setting is “Enabled”.

Note: The FX-1010 does not support fail-to-wire mode

4.1.4 Active Flows

Active flow capacity:
Specify the maximum number of active flow tracking resources that can be allocated from the choices presented in the pull-down. If set to 'Auto-Tune', the FX will automatically set the active flow capacity based upon the hardware platform and other criteria. An 'Active Flow' is a connection (either UDP or TCP) between a client and server which flows through the FX. If no data flows through the connection for the specified 'Flow tracking inactivity timeout' then the active flow tracking resource is released. The active flow tracking resource is allocated as soon as a SYN packet is seen to mark the start of a TCP connection.

Flow tracking inactivity timeout:
This setting controls how long an active-flow tracking resource will be allocated even though no data has flowed through the FX for that connection. The active flow tracking resource is automatically released if the connection is closed. The recommended standard for this setting is five days, however in many cases in can be set lower if conservation of active flow tracking resources is important. Most TCP applications do not have an issue with this setting because generally it doesn't matter if the active tracking resource is released when there is no activity. However, one notable exception is FTP, which allocates both a control connection and a data connection. If a long FTP transfer is flowing over the data port, there will not be any data flowing over the control connection, and if the inactivity timer expires for that control connection, the FTP transfer will cease. On the other hand, HTTP applications generally use very short-lived connections in which active-flow resources are quickly released.

If flow capacity reached:
This setting allows you to specify the course of action the FX should take if the active flow capacity is reached. The default "Fail-to-wire" setting is appropriate if the FX is intercepting traffic via 'in-line' mode or 'WCCP" mode. In both of those cases, traffic will pass through the FX as though it were an Ethernet cable. However if intercepting traffic in 'routed' mode, the 'Reject new connections' choice will be the course of action regardless of the setting. In either case, the FX will emit an SNMP alert if the number of active flows reaches 99% of capacity. The 'Status->WANOP Monitor' can be used to determine if the FX 'Active flow capacity' should be adjusted.

4.1.5 System Time

Network Time Server:
This setting will specify the host address for which the FX Series ADC will attempt to synchronize its time via the "Network Time Protocol". The FX Series ADC performs this synchronization one minute following a restart and once per week thereafter.

Time Zone:
This selector allows you to specify a time zone in which the FX Series ADC resides. In most cases, the default value of UTC-0 (GMT) is desirable because this will facilitate correlating system events with troubleshooting and other logs.

4.1.6 Software Updates

Automatically Distribute FX Series Remote Updates:
If "enabled", then the FX-Remote devices will periodically check to see if a newer version of firmware is available. If so, the FX-Remote devices will automatically download and apply the firmware update. The default value is "Disabled". It is a recommended practice that this setting be enabled only when you wish to deploy updated FX firmware during off-peak hours.

4.1.7 Other

Use Spanning-Tree Protocol:

If enabled, then spanning-tree protocol (STP) will be used when operating in “in-line” mode. Otherwise STP packets will be discarded. The default value is “Enabled”.

NOTE: If this setting is changed, it is necessary to restart the acceleration software on the “Status->Real-time Monitor” screen.

5 FX Series Remote General Settings

The FX Series Remote appliance works in conjunction with a head-end FX Series ADC appliance. The FX Series (ADC) appliance resides at the data center and supports connections with multiple remote sites where FX Series Remote appliances are installed

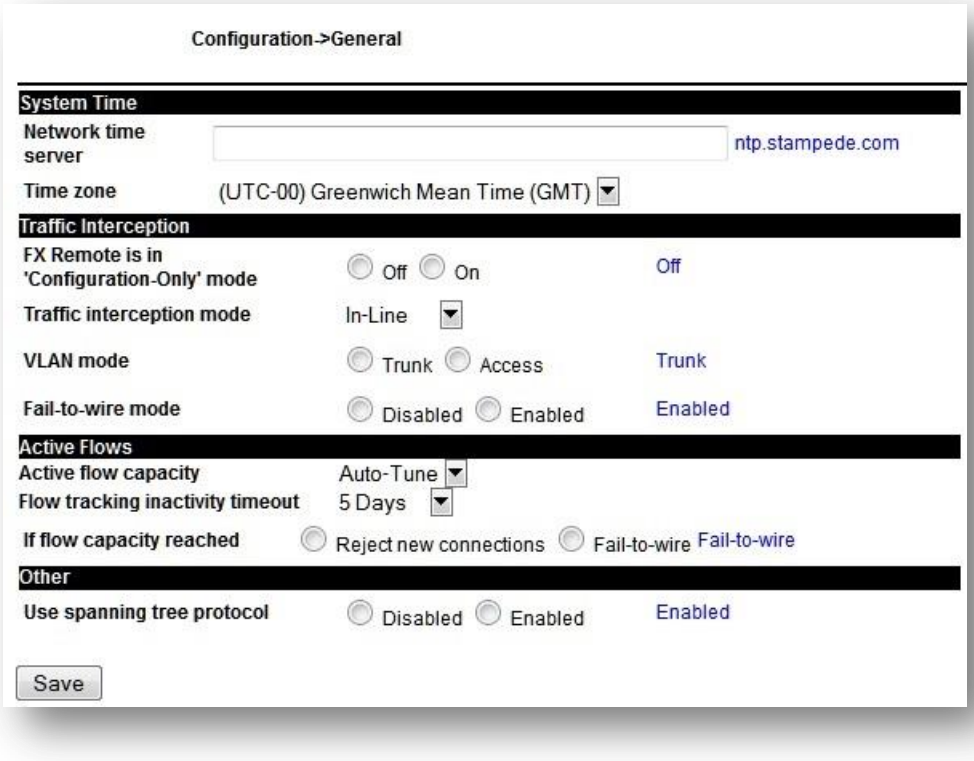

Figure 5-1 FX Series Remote General Edit Screen

5.1 System Time

Network time server:

Specify the host address for which the RWOC will attempt to synchronize its time via the "Network Time Protocol". The RWOC performs this synchronization one minute following a restart and once per week thereafter.

Time zone:

This selector allows you to specify a time zone in which the RWOC resides. In most cases, the default value of UTC-0 (GMT) is desirable because this will facilitate correlating system events with troubleshooting and other logs.

5.2 Traffic Interception

5.2.1 How to Configure FX Remote in 'Configuration-Only' mode:

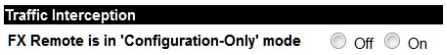

This setting is useful if you are in the process of configuring your FX Remote while the unit is networked in-line.
In "Configuration-Only" mode the in-line networking card is put into "bypass" mode so that traffic is simply passed through. When you are satisfied that the FX Remote is properly configured you can disable this setting. This setting only applies if the 'Traffic interception mode' is set to "In-Line" or "WCCP", this setting is ignored if the 'Traffic interception mode' is set to "routed". The default setting is "Off".

How to use "config-only" mode

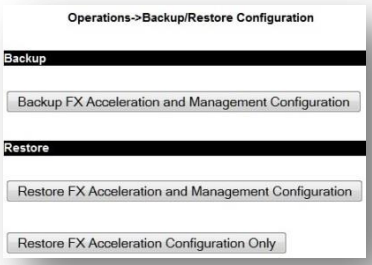

1. Using a browser, go to the FX's "Operations" page and select "Backup/Restore Configuration". Select the "Backup FX Acceleration and Management Configuration button"
   This button causes all FX configuration files to be stored into a zip file with the name "fxbackup-full_*hostname*_year-month-day.zip".
   Following this, a dialog will appear so that you can save this file to your desktop.
2. Save the backup file to a convenient location.
3. Put the appliance into "Configuration-Only" mode. The setting can be found on the "General Settings" Screen.
4. After making the changes to the configuration, return the configuration setting to the normal mode "disabled" and verify that there are no issues with the new configuration.
5. If issues are seen, you can return to original acceleration configuration quickly.
   a. Using a browser, go to the FX's "Configuration->General Settings" and change the "Configuration-Only" setting to "Enabled".
   b. Go to the FX's "Operations->Backup/Restore ".
   c. "Choose File" button.
   d. Following this, clicking the "Restore FX Acceleration Configuration Only" button will restore only the acceleration settings and exclude any management settings from the specified zip file. (This is also a useful method for transferring similar acceleration configurations from one FX to another without affecting management settings.)
   e. Load and apply the backup file.
   f. The acceleration appliance will automatically restart.
6. .Return to the "General Settings Screen" and return change the "Configuration-Only" setting to "Disabled".

NOTE: If appropriate use the "FX Series Appliance Manager", to re-establish any high-availability keys for the cluster-mates that share a common acceleration server configuration.
Configuring the high-availability cluster keys can be accomplished by selecting "Configure Appliance", and then "Configure Passwords".

Traffic Interception Mode:
This is the means by which the FX-ADC will transparently intercept packets. Choose either "In-Line", "Routed", WCCP" or "Disabled". The default is "In-Line".

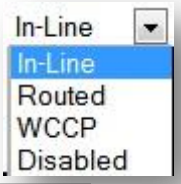

In-line Mode:
The LAN port is eth2 and the WAN port is eth3. Traffic is intercepted transparently. If there is a service disruption then the units will "fail-to-wire".

For in-line mode, you must assign an IP address to the in-path interface which bridges the LAN and WAN interfaces.
In FX nomenclature, the “WAN” interface is considered to be the interface which is connected to the satellite modems and the LAN interface is connected to the internet or to enterprise servers or clients.

Routed Mode:
Traffic must be directed to these interfaces by a router. You must assign an IP address to the in-path interface which will receive traffic from the router. There is no “fail-to-wire” capability if there is a service disruption.

WCCP mode:
Traffic is redirected to the FX by a Cisco router via WCCP.

Disabled Mode:
Traffic redirection is effectively shut off.

VLAN Mode:
Controls how the FX will process VLAN tags. In 'Trunk' mode, the VLAN tags are already embedded in the packets when they are intercepted by the FX. In 'Access' mode, the FX will add tags to untagged traffic. The default setting is ‘Trunk’.

Note: Changing either the ‘Traffic interception mode’ or ‘VLAN mode’ settings will automatically trigger a restart of the acceleration service.

Fail-to-wire Mode:
The FX will go into bypass mode it is not accelerating traffic or if it is powered off. If disabled, it will not go into bypass mode which will prevent packets from being forwarded through the FX. If the FX is adding the VLAN tags to the traffic, it may be desirable to disable 'Fail-to-wire' mode to prevent untagged traffic from entering a network. Note: The FX-1010 does not support fail-to-wire mode. The default setting is “Enabled”.

5.3 Active Flows

Active flow capacity:
Specify the maximum number of active flow tracking resources that can be allocated from the choices presented in the pull-down. If set to 'Auto-Tune', the FX will automatically set the active flow capacity based upon the hardware platform and other criteria. An 'Active Flow' is a connection (either UDP or TCP) between a client and server which flows through the FX. If no data flows through the connection for the specified 'Flow tracking inactivity timeout' then the active flow tracking resource is released. Each active flow tracking resource consumes about 300 bytes of non-swappable RAM. The active flow tracking resource is allocated as soon as a SYN packet is seen to mark the start of a TCP connection.

Flow tracking inactivity timeout:
This setting controls how long an active-flow tracking resource will be allocated even though no data has flowed through the FX for that connection. The active flow tracking resource is automatically released if the connection is closed. The recommended standard for this setting is five days, however in many cases in can be set lower if conservation of active flow tracking resources is important. Most TCP applications do not have an issue with this setting because generally it doesn't matter if the active tracking resource is released when there is no activity. However, one notable exception is FTP, which allocates both a control connection and a data connection. If a long FTP transfer is flowing over the data port, there will not be any data flowing over the control connection, and if the inactivity timer expires for that control connection, the FTP transfer will cease. On the other hand, HTTP applications generally use very short-lived connections in which active-flow resources are quickly released.

If flow capacity reached:
This setting allows you to specify the course of action the FX should take if the active flow capacity is reached. The default "Fail-to-wire" setting is appropriate if the FX is intercepting traffic via 'in-line' mode or 'WCCP" mode. In both of those cases, traffic will pass through the FX as though it were an Ethernet cable. However if intercepting traffic in 'routed' mode, the 'Reject new connections' choice will be the course of action regardless of the setting. In either case, the FX will emit an SNMP alert if the number of active flows reaches 99% of capacity. The 'Status->WANOP Monitor' can be used to determine if the FX 'Active flow capacity' should be adjusted.

5.4 Other

Use spanning-tree protocol:
If enabled, then spanning-tree protocol (STP) will be used when operating in "in-line" mode. Otherwise STP packets will be discarded. The default value is "Enabled".

Configuration Notes

Configuring Routed Mode for Two Interfaces
In routed mode, an IP address must be assigned to all in-path interfaces. This is a two-step process. First you must define two in-path interfaces each with a different IP address and gateway. Normally both of these interfaces will use VLAN 0. It is a recommended practice to define a comment for the in-path interfaces that describes both physical and routing aspects of the interface. Secondly you must define two LAN interfaces and then assign them to the in-path interfaces. On the "LAN Interface Definition" screen you would normally leave the "Untagged in-path interface" field at the default setting of "None".

6 FX Series Status

6.1 FX Series ADC Status

This section provides the reporting options and statistics on a current and aggregate basis of many parameters which are monitored on the FX Series.

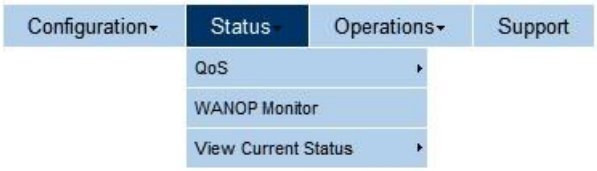

Figure 6-1 FX Series Status Menu

QOS Monitors (see Section 6.2 below)

The QoS Monitors by Queue and by Link provide real time views of vital QoS statistics, including current, average and elapsed statistics.

WANOP Monitor

Perform real-time monitoring of the acceleration activity via our JavaScript status monitor applet. In order to run this applet, JavaScript must be enabled in your browser settings.

View Current Status

View XML formatted reports of current acceleration statistics, traffic control, network interfaces, routes, current connection, and current throughput.

6.1.1 FX Series ADC WANOP Monitor

This page provides a real time view of vital ADC statistics.

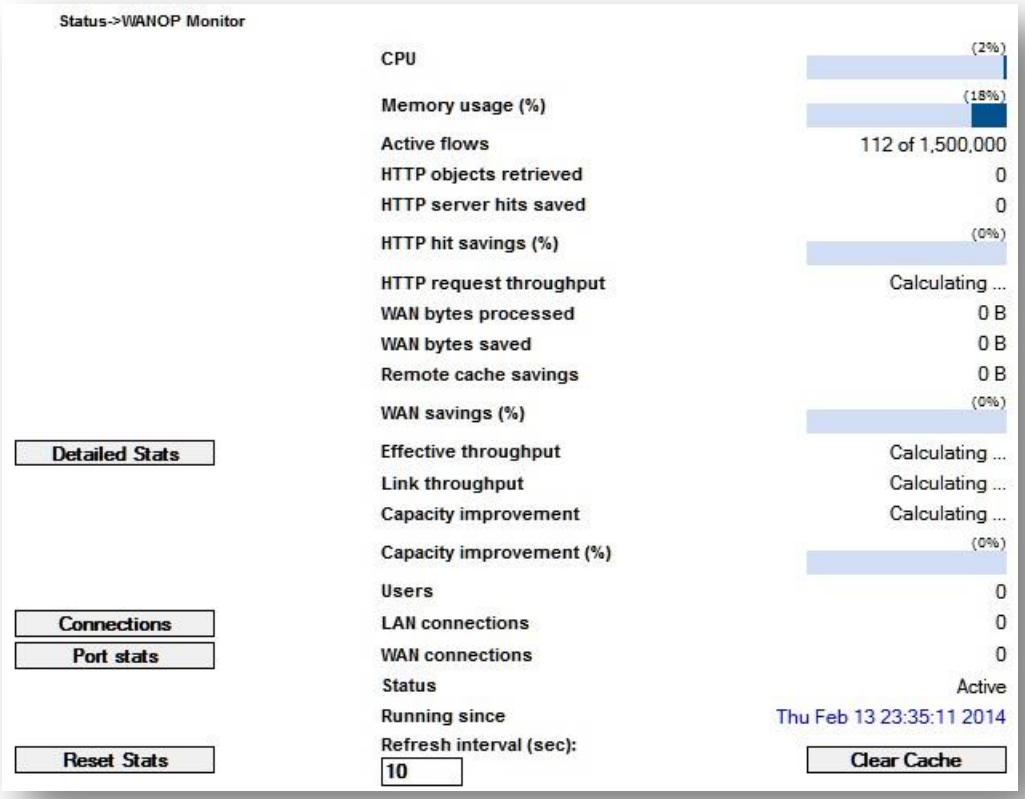

Figure 6-2 FX Series ADC WANOP Monitor Screen

Real-Time Statistics

CPU (%):

This represents the CPU utilization of the FX Series ADC for all system tasks.

Memory usage (%):

The percent of RAM being consumed by the ADC including both the WAN optimization service and the operating system is shown. When this value approaches 90%, the ADC will automatically reduce its memory resource consumption to prevent service disruption.

Active Flows:

The number of UDP and TCP connections currently flowing through the FX ADC and also the maximum capacity is shown. If the current number of active flows reaches 99% of capacity, then the FX will either go into "Fail-To-Wire" mode or not allow new connections to be established.

HTTP objects retrieved:

This shows the number of objects that the FX Series ADC retrieved from the target HTTP web servers.

HTTP server hits saved:

The number of times that the back-end HTTP server did not need to service an HTTP request because it was handled out of the FX Series ADC cache. This tally also includes savings generated by the advanced caching of the client acceleration software.

HTTP hit savings (%):
The percentage of "HTTP server hits saved" versus the total number of requests that would have been made to the back-end HTTP servers had it not been for the FX Series ADC.

HTTP request throughput:
The number of HTTP requests per second that were processed by the ADC over the last refresh interval.

WAN bytes processed:
Total number of bytes processed, both sent and received, by the FX Series ADC for traffic between clients and the FX Series ADC.

WAN bytes saved:
This represents the reduction in traffic between clients and the FX Series due to various acceleration techniques such as data compression, cache differencing, JPEG reduction etc.

Remote cache savings:
The number of bytes saved due to caching at the FX Remote.

WAN savings (%)
The "WAN bytes saved" divided by the "WAN bytes processed".

Effective throughput:
Shows the WAN bytes processed throughput based on the refresh interval.

Link throughput:
This shows the number of bits per second received over the WAN port based on the refresh interval.

Capacity improvement:
This is the bits per second improvement in the WAN link capacity due to the optimization provided by the FX ADC and Remote over the "refresh interval"

Capacity improvement %:
This is the percentage improvement of the WAN link capacity provided by the FX ADC and Remote over the last refresh interval.

Users:
This is the total number of unique IP addresses that currently have one or more TCP connections from the remote users to the FX Series ADC.

LAN Connections
This is the number of TCP connections from the FX Series ADC to backend content servers such as web servers or file servers.

WAN Connections
This is the number of TCP connections between the remote users (including FX Series Remotes) and the FX Series ADC.

Status:
This shows the current state of the acceleration service. Possible values are:

- "**Active**" - The FX is able to perform WAN optimization.
- "**Config-Only Mode**" – The FX can be configured, however traffic is not optimized unless in 'routed' traffic interception mode. If in 'in-line' traffic interception mode, the FX will go into 'bypass'. If in 'WCCP' traffic interception mode, the FX will drop out of all WCCP service groups. "Config-Only Mode" is set on the 'Configuration->General Settings' page.
- "**Invalid Or No License**" - The FaST codes for this FX do not exist or do not match this unit.

- "**Overflow – memory**" - The FX is low in memory resources. WANOP will continue for existing accelerated connections, however new connections from clients will be forwarded directly to the content servers.
- "**Overflow – sessions**" - The FX has reached its specified limit for the number of concurrent accelerated connections. WANOP will continue for existing accelerated connections, however new connections from clients will be forwarded directly to the content servers.
- "**Overflow - active flows**" - The FX has reached its limit of active flows. If in 'in-line' traffic interception mode, the FX will go into 'bypass'. If in 'WCCP' traffic interception mode 'the FX will drop out of all WCCP service groups.

Server running since
This is the date and time that the FX Series ADC was last started.

Refresh interval (sec)
This is how often the WANOP monitor will refresh the statistics. Changes take effect immediately and are remembered in subsequent accesses to the WANOP monitor.

Action Buttons

Detailed Stats:
This displays all aggregate tallies accumulated since the ADC service was last started or since they were last reset.

Connections:
This displays a list of each connection-group. If the ADC is accelerating traffic for FX-Remotes then a connection group may be comprised of multiple TCP connections that form a TurboStream. If the ADC is servicing HTTP clients in a one-sided mode then each connection group will be comprised of a single TCP connection. When the list is displayed you can drill down on an individual connection-group to examine detailed tallies.

Port stats:
Displays traffic statistics for each port definition that has been configured.

Reset Stats:
This button clears the acceleration server tallies and other stats.

Clear Cache:
This clears all objects out of all caches and then restarts the acceleration service.

6.1.2 FX Series ADC Current Statistics

This screen allows you to view various reports generated by the ADC. On most of the reports, you can hit the "F5" key and your browser will update the report. Note that the FX Series Remote has only a subset of the reports of the FX Series ADC.

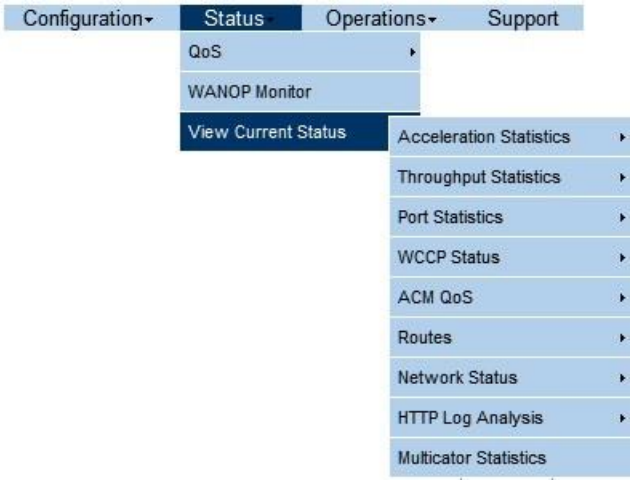

Figure 6-3 FX Series ADC Current Status Menu

Acceleration Statistics
This is a recording of the acceleration tallies that have accumulated since the last time that they were cleared or since the FX Series software was last restarted.

Aggregate Statistics:
This shows tallies of many acceleration techniques.

By L7 HTTP Policy:
This shows counters on a per L7 HTTP policy basis

By L5 Application Policy:
This shows counters on a per L5 Application policy basis

Current Connections:
This shows a report of remote clients currently connected.

Throughput Statistics
This is a recording of FX Series ADC throughput statistics based on current statistics since either the last time that the statistics were cleared or the configured "throughput statistics interval" has elapsed. The "delta" value in the upper-left corner of the report indicates the number of seconds that elapsed in the throughput interval. A snapshot of the current statistics was taken "delta" seconds ago; each counter in this snapshot is then subtracted from corresponding counter in the "current statistics". The remainders are then divided by the "delta" value to yield throughput values on a "per second" basis. It is possible that some of the values in the report are negative, meaning that a counter declined at a certain rate (for example CPU utilization may go down) over the throughput interval. The size of the throughput log file is relatively constant.

Port Statistics:
This provides summary traffic counters for each port definition

WCCP Status:
This provides a summary status of all of the service groups that the ADC is a member of. This report includes redirection maps and other important counters.

ACM QoS Status:

By Modem:

This report provides a summary status of the satellite modems if modem polling is enabled. For each modem, the current data rate and the number of successful and unsuccessful polls of both the data rate and the redundancy status is shown. If the polled data rate exceeds the licensed data rate, then the licensed data rate will be shown in the modem status table.

It also provides a statistical summary of the QoS queues and filters.

Sample below: (Note: Some reports may have pull downs to view additional information)

View Current Status->ACM QoS->By Modem

```
Tue Feb 11 17:43:05 GMT 2014

Modem Status Table Column Descriptions:

IP:         Modem IP Address
State:      Modem redundancy state (ONLINE,OFFLINE,NOT RESPONDING)
RateKbps:   Data rate from modem in Kbps
Spolls:     # Successfull data rate polls
Fpolls:     # Failed data rate polls
GoodRed:    # Successfull redundancy polls
BadRed:     # Failed redundancy polls
Link Name:  Name of the link


There are 0 'Modem Status Table' entries
              IP            State         RateKbps  Spolls  Fpolls GoodRed BadRed                    Link Name
              --            -----         --------  ------  ------ ------- ------                    ---------

#
# QoS Queues
#
There are 26 qosqs, of which 9 have packet compression enabled
                    Name  Pri Max Rate (bps) Cur Rate (bps)     Drops PktOnQ      BytesSent       PktsSent   IngressBytes    IngressPkts   Group
                    ----  --- -------------- --------------     ----- ------      ---------       --------   ------------    -----------   -----
             Signaling Q    1      20,000             208         0      0            50,760            846              0              0   LogicA
-- Internal Signals --    1     100,000,000           0         0      0                 0              0              0              0   unassigned
                   VoIP    2     2,000,000         6,632         0      0         1,718,754         27,861              0              0   LogicA
              cxm_VoIP     2     6,000,000             0         0      0                 0              0              0              0   Cyclops Xylex Management
             TommyQ35      2             0             0         0      0                 0              0              0              0   unassigned
              CAM_VoIP     2     10,000,000            0         0      0                 0              0              0              0   Cyclops Acme Management
              TommyQ5      2             0             0         0      0                 0              0              0              0   unassigned
             TommyQ28      2             0             0         0      0                 0              0              0              0   unassigned
              cxc_VoIP     2     6,000,000             0         0      0                 0              0              0              0   Cyclops Xylex Crew
             TommyQ32      2             0             0         0      0                 0              0              0              0   unassigned
              AAM_VoIP     2     10,000,000            0         0      0                 0              0              0              0   Argo Acme Management
              axm_VoIP     2     6,000,000             0         0      0                 0              0              0              0   Argo Xylex Management
             TommyQ20      3             0             0         0      0                 0              0              0              0   unassigned
              AAC_VoIP     3     10,000,000            0         0      0                 0              0              0              0   Argo Acme Crew
              CAC_VoIP     3     10,000,000            0         0      0                 0              0              0              0   Cyclops Acme Crew
      cxm_Web Browsing     3     20,000,000            0         0      0                 0              0              0              0   Cyclops Xylex Management
          Web Browsing     3     20,000,000            0         0      0                 0              0              0              0   LogicA
```

Routes:

This report shows a detailed list of all the routes that are currently in use on this FX on all route tables. This includes the management routing tables as well as the route tables associated with the in-path interfaces. Sample below:

View Current Status->Routes->ByTable

```
Tue Feb 11 15:59:06 GMT 2014

default via 192.168.1.2 dev br100  table 1100  initcwnd 14 initrwnd 14
broadcast 192.168.1.0 dev br100  table 1100  proto kernel  scope link  src 192.168.1.10
192.168.1.0/24 dev br100  table 1100  proto kernel  scope link  src 192.168.1.10  initcwnd 14 initrwnd 14
local 192.168.1.10 dev br100  table 1100  proto kernel  scope host  src 192.168.1.10
broadcast 192.168.1.255 dev br100  table 1100  proto kernel  scope link  src 192.168.1.10
default via 192.168.1.2 dev br500  table 1500  initcwnd 14 initrwnd 14
broadcast 192.168.1.0 dev br500  table 1500  proto kernel  scope link  src 192.168.1.10
192.168.1.0/24 dev br500  table 1500  proto kernel  scope link  src 192.168.1.10  initcwnd 14 initrwnd 14
local 192.168.1.10 dev br500  table 1500  proto kernel  scope host  src 192.168.1.10
broadcast 192.168.1.255 dev br500  table 1500  proto kernel  scope link  src 192.168.1.10
default via 192.168.2.2 dev br200  table 1200  initcwnd 14 initrwnd 14
broadcast 192.168.2.0 dev br200  table 1200  proto kernel  scope link  src 192.168.2.20
192.168.2.0/24 dev br200  table 1200  proto kernel  scope link  src 192.168.2.20  initcwnd 14 initrwnd 14
local 192.168.2.20 dev br200  table 1200  proto kernel  scope host  src 192.168.2.20
```

Network Status:

Network Status by Interface

This page shows various counters and status indicators of all network interfaces, including management, auxiliary, and all LAN and WAN interfaces. If tagged VLAN In-Path interfaces have been defined then stats for these will be individually be broken out. For example if you defined an in-path interface on for VLAN 100, you would see a section for “eth3.100”. If you are intercepting traffic in-line you will see a section for “br0”. The “lo” interface section represents internal loopback traffic. If you are interfacing to WCCP in ‘layer 3’ mode, you will see a section called “greX” which reflects the traffic redirected to the FX over a GRE tunnel

This screen is critical for diagnosing common network connectivity problems, such as bad cables, incompatible link speed and duplex negotiations and Ethernet frame errors.
Sample below: (Note: Some reports may have pull downs to view additional information)

FX-4010 Interface and ‘Autosense’

The FX-4010 does not autosense cable type. Therefore if you are interfacing directly to a modem device, without going through an intermediate switch, it is necessary to use a cross-over cable.

View Current Status->Network Status->By Interface

```
#
#
# Ethernet Interface Status as of Mon Feb 10 21:46:50 GMT 2014
#

TYPE  PORT      SPEED     DUPLEX AUTONEG? LINK?      STATE  MTU        MAC Address
==== =====  ========= ========== ======== ===== ========== ====  =================
MGMT  eth0   1000Mb/s       Full       on   yes         up 1500  00:25:90:3b:6a:ce
 AUX  eth1   Unknown!   Unknown!       on    no       down 1500  00:25:90:3b:6a:cf
 LAN  eth2   1000Mb/s       Full       on   yes         up 1500  00:e0:ed:18:f9:89
 WAN  eth3   1000Mb/s       Full       on   yes         up 1500  00:e0:ed:18:f9:88


br0       Link encap:Ethernet  HWaddr 00:E0:ED:18:F9:88
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:248493 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:13682226 (13.0 Mb)  TX bytes:0 (0.0 b)

br15      Link encap:Ethernet  HWaddr 00:E0:ED:18:F9:88
          inet addr:172.27.115.10  Bcast:172.27.115.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:6959 errors:0 dropped:0 overruns:0 frame:0
          TX packets:13 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:347942 (339.7 Kb)  TX bytes:2508 (2.4 Kb)

br100     Link encap:Ethernet  HWaddr 00:E0:ED:18:F9:88
          inet addr:192.168.1.10  Bcast:192.168.1.255  Mask:255.255.255.0
```

WAN Interface Status

This page shows various counters and status indicators of the WAN interface.

This page is critical for diagnosing common network connectivity problems, such as bad cables, incompatible link speed and duplex negotiations and Ethernet frame errors.

View Current Status->Network Status->Of WAN Interface

```
Mon Feb 10 22:09:35 GMT 2014
eth3      Link encap:Ethernet  HWaddr 00:E0:ED:18:F9:88
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:171915 errors:0 dropped:0 overruns:0 frame:0
          TX packets:209382 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:12068118 (11.5 Mb)  TX bytes:12989404 (12.3 Mb)

#
# Status of WAN port eth3
#
TYPE  PORT      SPEED     DUPLEX AUTONEG? LINK?      STATE  MTU        MAC Address
==== =====  ========= ========== ======== ===== ========== ====  =================
 LAN  eth3   1000Mb/s       Full       on   yes         up 1500  00:e0:ed:18:f9:88
```

HTTP Log Analysis:

- The chart shows graphically the statistics on the HTTP objects which have been requested by clients accelerated by this ADC.
- If access logging has been turned off, then the graph will not display information as is shown in the graphical example shown below: Access logging was turned off in October.
- The “Refresh” button will update the report that was last updated.
- If the “Refresh” button was never pressed then no report will be displayed.
- The “Download” button allows download of a zipped copy of the HTTP log file in its current state.
- The format for the downloaded HTTP Log is described in the information on the next page.

Usage Statistics for MLANAdc - Summary by Month

Usage Statistics for MLANAdc

Summary by Month
Generated 11-Feb-2014 16:46 GMT

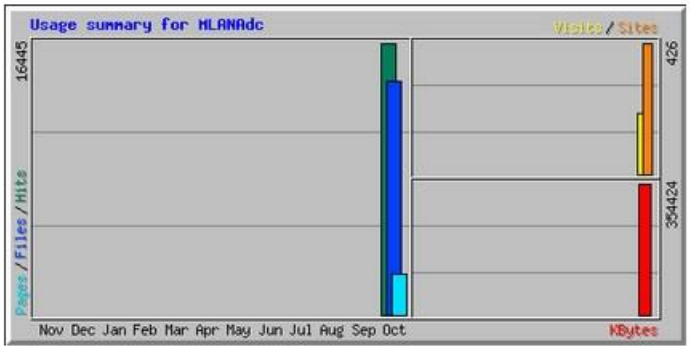

Summary by Month

Month	Daily Avg				Monthly Totals					
	Hits	Files	Pages	Visits	Sites	KBytes	Visits	Pages	Files	Hits
Oct 2013	3289	2826	499	39	426	354424	198	2495	14130	16445
Totals						**354424**	**198**	**2495**	**14130**	**16445**

HTTP Log Format

Each line of the FX HTTP log represents a separate HTTP request. Each line has the following format:

OriginatorsIP streamInfo AccelerationTechnique - [TimeStamp] "HTTPRequest" HTTPStatus Bytes

OriginatorsIP is the IP address of the client that made the request

StreamInfo is seven strings separated by commas. The first three strings are hex digits that represent flow identifiers.

The next 4 strings are:

- fbr=nnn where nnn is number of milliseconds since the client request came in until the first byte was received from web server
- lbr=nnn where nnn is number of milliseconds since the client request came in until the last byte was received from the web server
- fbs=nnn where nnn is number of milliseconds since the client request came in until the first byte was sent back to the client
- lbs=nnn where nnn is number of milliseconds since the client request came in until the last byte was sent back to the client

AccelerationTechnique is one of:

- "IMS-HIT" - The client issued an If-Modified-Since request for an object which was in the cache and fresh.
- "INM-HIT" - The client issued an If-None-Match request for an object which was in the cache and fresh.
- "HIT" - The object was returned from cache
- "CV-HIT" - The object was returned "Not Modified" from the content validation cache
- "DIFF_HHIT" – Cache difference of just the response HTTP header
- "DIFF_HIT" – Cache difference of the object
- "REFRESH-UNMODIFIED" – The requested object was cached but STALE. The IMS query for the object resulted in "304 not modified"
- "REFRESH-MODIFIED" – The requested object was cached but STALE. The IMS query returned the new content.
- "404-HIT" - The object was returned a 404 Not Found based on the freshness of the 404 cache object

"MISS" - The object was not in the cache but written/updated in cache

[TimeStamp] is the date and time of when the HTTP request was processed.

"HTTPRequest" is the full URL to the content requested

"HTTPStatus" is the HTTP status code. Bytes are the number of payload bytes sent back to the user. This does not include the number of bytes in the HTTP header.

Multicator Status:

This page shows various counters and statistics of the reliable multicast fan-out feature of the FX. If a transfer is in progress, the name of the file being transferred is shown.

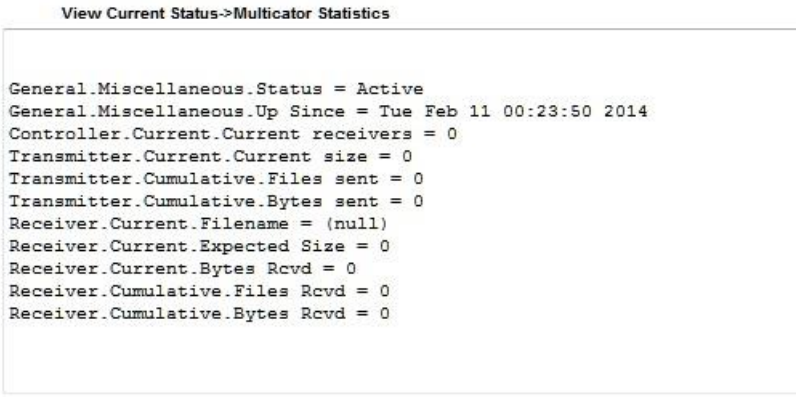

6.2 QOS Monitors

The QOS Monitors provide a real time view of vital QoS statistics, including current, average and elapsed stats. All elapsed and average stats begin with the last reset of the screen.

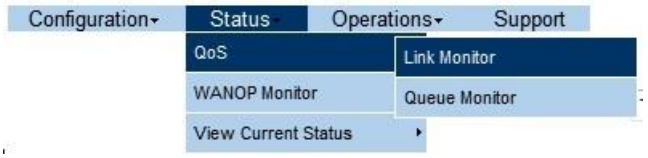

QOS Monitor Options

Pull Downs

Update Interval:
This determines the rate at which the screen is updated. There are selectable rates from 1 to 60 seconds. Every time the screen is updated, the browser gets the statistics from the FX. Keep in mind the RTT to the appliance as well as the network capacity in setting this parameter.

Average Time Period:
This selects the time for which stats are averaged in the associated field. Options are 1, 24, 72 hours and 30 days.

Display Type:
Raw shows the raw number. Human readable uses metric conventions. All data rate numbers are given as powers of 10. In this case, K = 10^3, M – 10^6, G = 10^9, and T = 10^{12}. This refers to both data rate and total number of packets sent.
When referring to the number of bytes or packets sent, the convention is powers of 2. In this case, K = 2^{10}, M = 2^{20}, G = 2^{30} and T = 2^{40}.

Reset Button:

The Reset button clears and restarts all statistics.

View QoS Queue Settings:
This is a hyperlink to the queue screen.

6.2.1 QoS Link Monitor

The QOS Link Monitor provides a real time view of vital QoS statistics, including current, average and elapsed stats. All elapsed and average stats begin with the last reset of the screen

Status->QoS Link Monitor

Update Interval (s): 1 Average Time Period: 1 Hour

Display Type: Human Reset Counters: Reset

Link				Current				1 Hour Average				
Name	Queues	Max Data Rate	ACM	WAN Rate	PPS	Comp. (%)	Drop Rate	WAN Rate	0 Day(s) Comp. (%)	2 Hour(s) 38 Min(s) Bytes Transferred	54 Sec(s) Packets Transferred	Drop Rate
Link_1	View	25.00 Mbps	Disabled	0 bps	11	0	0	0 bps	0	0 B	0	0
Link_2	View	100.00 Mbps	Disabled	0 bps	11	0	0	0 bps	0	592,137.04 TB	651,061.56 (T)	100
Aggregate Totals				0 bps	22	0	0	0 bps	0	0 B	0	0

Figure 6-4 FX Series QOS Monitor by Link

Link

All defined links are represented in the list

Name:
This is the name of the link and is a hyper link to the configuration screen for that link.

Queues:
This is a hyper link to the QOS Queue Monitor screen with only the queues associated for this link displayed.

Max Data Rate:
If ACM is enabled, this is the rate that the modem is currently running at. If ACM is not enabled, this is the configured clear sky data rate.

ACM:
This is a flag that indicates if ACM is enabled for the link.

Current

There are three parameters associated with the current WAN Rate

WAN Rate:
This shows the data rate averaged over the previous 5 seconds.

PPS:
This is the current Packet/second rate averaged over the previous 5 seconds

Compression (%)
This indicates the percentage bandwidth reduction due to header and payload compression.

Drop Rate:
This shows the drop rate averaged over the previous 5 seconds.

Average Rates:

This average is set with the average time period pull down menu shown above. Average time can be 1, 24, 72 hours or 30 days.

The fields in Average Section update at intervals that are a function of the Average Time Period. When the Average Time Period is set at 1 Hour, the fields update every minute on the minute as given by the "Since Last Reset" timer and represent the statistics for the previous hour. If the total time is less than an hour, then the fields represent the data since the last reset. For example, if it's been 37 minutes, 14 seconds since the reset, then the data represents the first 37 minutes of traffic.

When the Average Time Period is set to 24 hours, 72 hours or 30 days, the fields update every hour on the hour as given by the “since Last Reset” timer. If the total time is less than the number of hours set by the Average Time Period, then the fields represent the data since the last reset.

For example, if the Average Time Period is set to 72 hours, and it’s been 54 hours, 38 minutes and 18 seconds since the last reset, then the fields represent the first 54 hours of traffic. The rates are updated every minute and represent the data over the respective interval terminating at that time. If the ‘Average Time Period’ is set to ‘Since Reset’, the fields are updated at each update interval.

WAN Rate:
Output data rate averaged rate over the interval.

Compression (%)
This value indicates the percentage bandwidth reduction due to header and payload compression averaged over the interval.

Bytes Transferred:
Total bytes transferred during the interval.

Packets Transferred:
Total packets transferred during the interval.

Drop Rates:
This is the rate at which packets are dropped during the interval.

Aggregate Totals

This line gives the totals over all the links.

6.2.2 QOS Queue Monitor

The QoS Queue Monitor provides a real time view of vital QoS statistics, including current, average and elapsed stats for each defined queue. When navigating to the page from the top level menus, all the queues, sorted by priority, are shown. When navigating to the page from the queue tab on the link monitor screen, only the queues associated with that link are displayed. The statistics on this page are given for each queue in one of three modes: current, time average and since last reset.

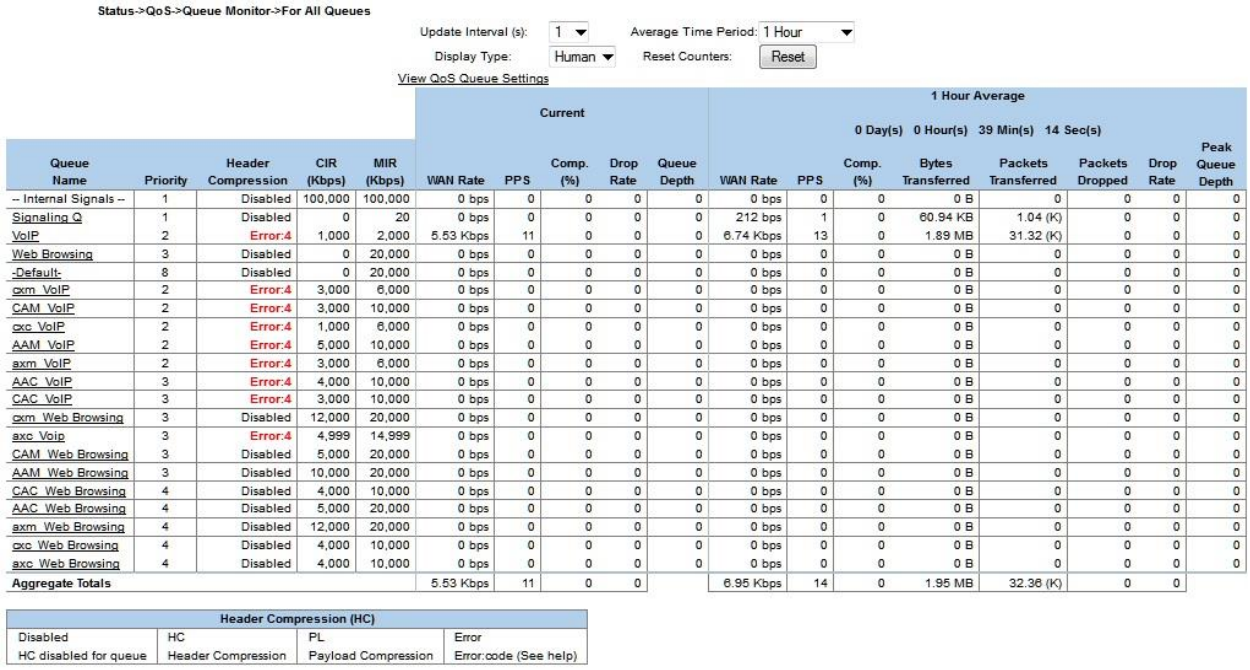

Status->QoS->Queue Monitor->For All Queues

Update Interval (s): 1 Average Time Period: 1 Hour

Display Type: Human Reset Counters: Reset

View QoS Queue Settings

					Current					1 Hour Average 0 Day(s) 0 Hour(s) 39 Min(s) 14 Sec(s)							
Queue Name	Priority	Header Compression	CIR (Kbps)	MIR (Kbps)	WAN Rate	PPS	Comp. (%)	Drop Rate	Queue Depth	WAN Rate	PPS	Comp. (%)	Bytes Transferred	Packets Transferred	Packets Dropped	Drop Rate	Peak Queue Depth
-- Internal Signals --	1	Disabled	100,000	100,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
Signaling Q	1	Disabled	0	20	0 bps	0	0	0	0	212 bps	1	0	60.94 KB	1.04 (K)	0	0	0
VoIP	2	Error:4	1,000	2,000	5.53 Kbps	11	0	0	0	6.74 Kbps	13	0	1.89 MB	31.32 (K)	0	0	0
Web Browsing	3	Disabled	0	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
-Default-	8	Disabled	0	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
cxm_VoIP	2	Error:4	3,000	6,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
CAM_VoIP	2	Error:4	3,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
cxc_VoIP	2	Error:4	1,000	6,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
AAM_VoIP	2	Error:4	5,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
axm_VoIP	2	Error:4	3,000	6,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
AAC_VoIP	3	Error:4	4,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
CAC_VoIP	3	Error:4	3,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
cxm_Web Browsing	3	Disabled	12,000	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
axc_Voip	3	Error:4	4,999	14,999	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
CAM_Web Browsing	3	Disabled	5,000	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
AAM_Web Browsing	3	Disabled	10,000	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
CAC_Web Browsing	4	Disabled	4,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
AAC_Web Browsing	4	Disabled	5,000	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
axm_Web Browsing	4	Disabled	12,000	20,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
cxc_Web Browsing	4	Disabled	4,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
axc_Web Browsing	4	Disabled	4,000	10,000	0 bps	0	0	0	0	0 bps	0	0	0 B	0	0	0	0
Aggregate Totals					5.53 Kbps	11	0	0		6.95 Kbps	14	0	1.95 MB	32.36 (K)	0	0	

Header Compression (HC)			
Disabled	HC	PL	Error
HC disabled for queue	Header Compression	Payload Compression	Error:code (See help)

Figure 6-5 FX Series QoS Monitor by Queue

Priority:

This is the assigned priority and queues are sorted in this list by priority.

Header Compression:

The parameters for this field are:

- Disabled – if no compression is enabled.
- HC – if only header compression is enabled.
- PL – if both header and payload compression are enabled.
- Error: code – Indicates an error. The codes for the errors are as follows:

  3 – This code means that the packet compressor function was unable to establish a tunnel with a peer. Possible reasons for this are:

  · The wrong MAC address was specified.

  · There is a routing problem.

  · The target peer is not in the path of the data flow.

  · The peer is not connected to the network perhaps due to faulty or disconnected Ethernet cable.

CIR: (Kbps):
This is the assigned CIR. Queues that are internal or assigned to a strict-priority link will show “n/a”.

MIR: (Kbps):
This is the assigned MIR. Queues that are internal or assigned to a strict-priority link will show “n/a”.

Current QOS Stats

WAN:
Output data rate averaged over the previous 5 seconds.

PPS:
Output packet rate averaged over the previous 5 seconds.

Compression (%)
This value indicates the percentage bandwidth reduction due to header and payload compression averaged over the previous 5 seconds.

Drop Rate:
Drop rate averaged over the previous 5 seconds.

Queue Depth:
Max queue depth over the previous 5 seconds.

Average Rate:

This average is set with the average time period pull down menu shown above. Average time can be 1, 24 or 72 hours, 30 days or since the last rest.

The fields in Average Section update at intervals that are a function of the Average Time Period. When the Average Time Period is set at 1 Hour, the fields update every minute on the minute as given by the “Since Last Reset” timer and represent the statistics for the previous hour. If the total time is less than an hour, then the fields represent the data since the last reset.

For example, if it’s been 37 minutes, 14 seconds since the reset, then the data represents the first 37 minutes of traffic. When the Average Time Period is set to either 24 or 72 hours or 30 days, the fields update every hour on the hour as given by the “since Last Reset” timer. If the total time is less than the number of hours set by the Average Time Period, then the fields represent the data since the last reset.

For example, if the Average Time Period is set to 72 hours, and it’s been 54 hours, 38 minutes and 18 seconds since the last reset, then the fields represent the first 54 hours of traffic.

The rates are updated every minute and represent the data over the respective interval terminating at that time.

WAN Rate:
Output data rate averaged over the interval.

PPS:
Output packet rate averaged over the interval.

Compression (%)
This value indicates the percentage bandwidth reduction due to header and payload compression averaged over the interval.

Bytes Transferred:
Total bytes transferred during the interval.

Packets Transferred:
Total packets transferred during the interval.

Packets Dropped:
Total packets dropped during the interval.

Drop Rates:
This indicates the rate at which packets are dropped during the interval.

Peak Queue Depth:
This is the maximum number of packets on the queue during t

Aggregate Totals

This line gives the totals over all the queues and represents the actual output rate of the WAN port of the appliance.

6.3 FX Series Remote Status

6.3.1 FX Series Remote WANOP Monitor

The WANOP Monitor provides a real time view of vital server statistics.

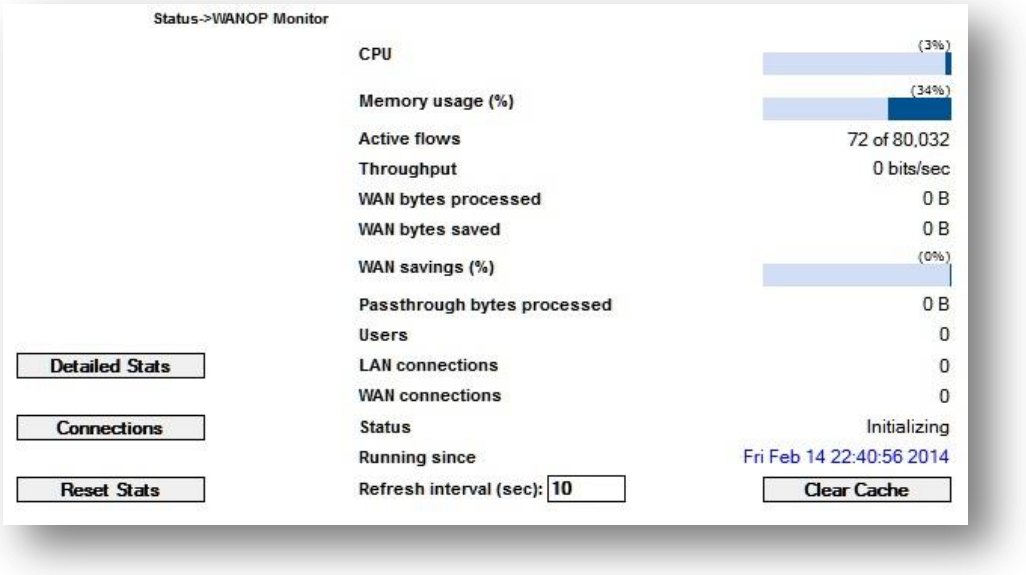

Figure 6-6 FX Series Remote Real-Time Monitor Screen

Real-Time Statistics

CPU (%):
This represents the CPU utilization of the FX for all system tasks.

Memory usage (%):
This shows the percent of RAM being consumed by the FX Series Remote including both the WAN optimization service and the operating system. When this value approaches 90%, the appliance will automatically reduce its memory resource consumption to prevent service disruption.

Active Flows:
The number of UDP and TCP connections currently flowing through the FX Remote and also the maximum capacity is shown. If the current number of active flows reaches 99% of capacity, then the FX will either go into “Fail-To-Wire” mode or not allow new connections to be established.

Throughput:
This is the throughput, in bits per second, since the last refresh interval.

WAN bytes processed:
Total number of bytes processed, both sent and received, by the acceleration server for traffic between clients and the acceleration server.

WAN bytes saved:
This represents the reduction in traffic due between clients and the acceleration server due to various acceleration techniques such as data compression, cache differencing, JPEG reduction, Dynamic Data Suppression etc.

WAN savings (%):
The "WAN bytes saved" divided by the "WAN bytes processed".

Pass through bytes processed:
This is the total quantity of bytes that have flowed through the acceleration server but no optimization operations were attempted.

Users:
This is the total quantity of unique IP addresses that have TCP connections flowing through the FX Remote.

LAN Connections:
This shows the quantity of TCP connections between clients and the FX Remote.

WAN Connections:
This shows the quantity of TCP connections between the FX Remote and FX SERIES ADCs.
Status:
This shows the current state of the acceleration service. Possible values are:

- "**Active**" - The FX is able to perform WAN optimization.
- "**Invalid Or No License**" - The FaST codes for this FX do not exist or do not match this unit.
- "**Config-Only Mode**" – The FX can be configured, however traffic is not optimized unless in 'routed' traffic interception mode. If in 'in-line' traffic interception mode, the FX will go into 'bypass'. If in 'WCCP' traffic interception mode the FX will drop out of all WCCP service groups. "Config-Only Mode" is set on the 'Configuration->General Settings' page.
- "**Overflow – memory**" - The FX is low in memory resources. WANOP will continue for existing accelerated connections, however new connections from clients will be forwarded directly to the content servers.
- "**Overflow – sessions**" - The FX has reached its specified limit for the number of concurrent accelerated connections. WANOP will continue for existing accelerated connections, however new connections from clients will be forwarded directly to the content servers.
- "**Overflow - active flows**" - The FX has reached its limit of active flows. If in 'in-line' traffic interception mode, the FX will go into 'bypass'. If in 'WCCP' traffic interception mode 'the FX will drop out of all WCCP service groups.

Running since:
This is the date and time when the acceleration server was last started.

Refresh interval (sec):
This is how often WANOP monitor will refresh the statistics. Changes take effect immediately and are remembered in subsequent accesses to the WANOP monitor.

.

Action Buttons

Detailed stats:
This displays tallies and other stats.

Connections:
This displays the active user connections to the acceleration server.

Reset Stats:
This clears the acceleration server tallies and other stats.

Clear Cache:
This clears all objects out of all caches, and then restarts the acceleration service.

6.3.2 FX Series Remote Current Status Reports

This page allows you to view various reports generated by the FX Series Remote. On most of the reports, you can press the "F5" key and your browser will update the report. See the ADC Current Status Reports section 6.1.2 above for examples of the screens not included here.

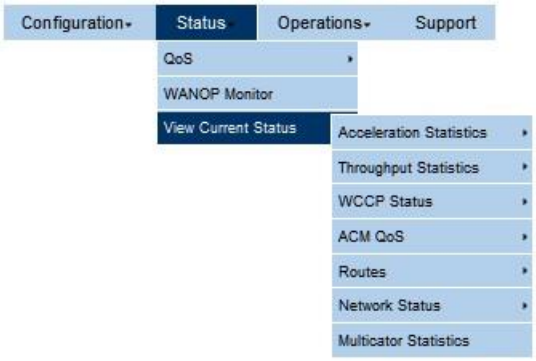

Figure 6-7 FX Series Remote Current Status Menu

Acceleration Statistics

Aggregate Statistics:
This is a recording of the acceleration tallies that have accumulated since the last time that they were cleared or since the acceleration server software was last restarted.

Current Connections:
This report shows the user names, connection type, and other data of each currently active connection to the acceleration server.

Throughput Statistics

Aggregate Throughput:
This is a recording of acceleration server throughput statistics based on current statistics since either the last time that the statistics were cleared or the configured "throughput statistics interval" has elapsed. The "delta" value in the upper-left corner of the report indicates the number of seconds that elapsed in the throughput interval.

Aggregate Throughput Calculation:
A snapshot of the current statistics was taken "delta" seconds ago; each counter in this snapshot is then subtracted from corresponding counter in the "current statistics". The remainders are then divided by the "delta" value to yield throughput values on a "per second" basis.

NOTE: It is possible that some of the values in the report are negative, meaning that a counter declined at a certain rate (for example CPU utilization may go down) over the throughput interval. The size of the throughput log file is relatively constant.

Other Status Reports

WCCP Status > by WCCP Definition:
This provides a summary status of all of the service groups that the FX Series Remote is a member. This report includes redirection maps and other important counters.

ACM QoS > by Modem:
This provides a summary status of the satellite modems if modem polling is enabled. If the polled data rate exceeds the licensed data rate, then the licensed data rate will be shown in the modem status table.

Routes > by Table:
This report shows a detailed list of all the routes that are currently in use on this FX Series Remote on all route tables.

Network Status > By Interface and > Of WAN Interface:
These reports shows detailed statistical and status information about each active network interface on this FX Series Remote.

```
Fri Feb  7 16:19:31 GMT 2014
eth3      Link encap:Ethernet  HWaddr 00:E0:ED:18:F9:88
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:3649582 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4916080 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:256187786 (244.3 Mb)  TX bytes:355968833 (339.4 Mb)

#
# Status of WAN port eth3
#
TYPE  PORT     SPEED     DUPLEX AUTONEG? LINKDETECT      STATE    MTU        MAC Address
==== =====  ========= ========== ======== ========== ========== ====== ==================
 LAN  eth3  1000Mb/s       Full       on         yes         up   1500 00:e0:ed:18:f9:88
```

This screen is critical for diagnosing common network connectivity problems, such as bad cables, incompatible link speed and duplex negotiations and Ethernet frame errors

7 FX Series Optimization Settings

7.1 Application Policies Overview

Customize the optimization techniques settings on the FX Series ADC by central policy management that will apply to your enterprise applications.

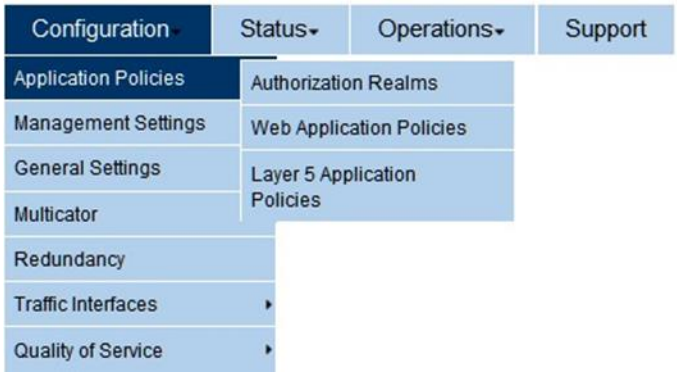

Figure 7-1 FX Series ADC Features Menu

Authorization Realms
Define objects based on source IP, browser type, and other criteria to be used as a basis for policies.

Web Application Policies
Choose this if you wish to configure the optimization properties of the HTTP-based Web applications that the FX Series ADC will accelerate.

Layer 5 Application Policies
Choose this if you wish to configure the optimization properties of the non-HTTP-based TCP/IP applications that FX Series ADC will accelerate.

7.1.1 FX Series Optimization Summary

- The FX Series is built upon a flexible and secure base of operating system and acceleration software. Utilizing an architecture that is modularized, The FX Series accelerates the movement of data, accelerates the processing of that data and ensures a safe and reliable system.
- The FX Series can be configured and implemented in several ways to accelerate enterprise application traffic. All implementations start with an appliance, which functions as a "front-end" processor to web servers. In this initial "one-sided" mode, the appliance processes traffic between HTTP clients (browsers) and web servers. The FX Series can also be implemented with optional client-side hardware, FX Series Remote appliance. The client-side appliance provides many additional optimization features in a “two-sided” mode.
- With flexible options for in-line or Cisco’s Web Cache Communication Protocol (WCCP) the FX Series devices deliver unprecedented transparent optimization. End to End assurance is maintained for all applications providing complete transparency and the ability for existing Quality of Service (QOS) and network visibility management programs to continue monitoring the health of your network.

- Business critical applications requiring 24/7 availability will always perform optimally no matter how fast your business grows. The Stampede FX Series platforms support deployments in either N+N or N+ 1 configuration. Simply add additional FX Series devices as needed for increased scalability or failover protection.
- Stampede's FX Series platforms provide a comprehensive range of flexible options for total transparent 24/7 operation within your existing or growing network infrastructure. No matter what your application acceleration or WAN optimization requirements are today or in the future, Stampede's FX Series platform solutions will handle all your business critical applications with ease

7.1.2 Single-Sided Optimizations:

GZIP:
Industry standard compression is applied to both static and dynamic HTML and XML content including attachments.

Image Optimization Transformation:
This reduces image size, while optimizing image quality.

Static Caching:
This returns un-expired static objects and pages directly from cache.

7.1.3 Two-sided Optimizations

TurboStreaming:
Optimization techniques are utilized to overcome inherent TCP inefficiencies. This removes congestion from WAN links and significantly improves response time for downloading large attachments, even those considered to be non-compressible.

Content Aware Compression

Dynamic Cache Differencing:
The FX Series ADC appliance and the Remote appliance maintain a copy of browser cache and send only the difference back to the browser when content is changed.

Persistent Connection:
The FX Series maintains a persistent TCP connection with clients with a configurable timeout based on application policy.

Content Aware Caching:
The FX Series accelerates client accesses by eliminating round-trips to servers to validate data. This is intelligently performed to avoid unnecessary network consumption and reduces server transaction processing.

7.1.4 Authorization Realms

Authorization realms provide a means for grouping users so that different policy application attributes can be assigned. For example, perhaps you want a set of users to have unrestricted Internet access and others you wish to be routed to another appliance that performs filtering, then you can define authorization realms to delineate these user groups by assigning different application policies based on authorization realm.

NOTE: Examples for using Authorization Realms is shown in Section 7.2.2

Realm Name:
This specifies the logical name to assign to this realm. This name is used to reference the definition in the application policies and client policies.

Comment: This is a description in which the administrator can delineate the rationale for this authorization realm.

Origin IP Address Ranges:
Specified using "CIDR" notation where a base IP address is followed by a '/' character which is followed by a value between 1 and 32 that denotes the number of bits used to describe the network and the remaining bits (32 – the value) are used to specify the nodes on that network. For example a setting of 192.110.1.0/24 would be equivalent to specifying a network of 192.110.1.0 with a net mask of 255.255.255.0. Separating each CIDR entry with a comma can specify multiple destinations. You may also enter one or more single IP addresses or hyphenated IP address ranges, separated by commas in the same manner. i.e. 10.2.2.5 or 10.2.2.50-10.2.2.59. In a two-sided environment with FX-Remotes, the IP address must be that of an in-path interface of the FX-Remote. The default setting is any network.

VLAN ID:
Specifies the VLAN ID for which the realm should apply. If "None" is selected then the VLAN ID is not part of the match criteria when realms are evaluated. The selector only shows VLAN IDs for which an in-path interface has been defined.

Client Type:
This field allows you to specify that client types that are associated with this authorization realm.
The valid choices are:

Native – Traffic that does not flow thru a FX Remote
FX Remote – Traffic which is accelerated by the FX Series Remote Appliance

NOTE: Press and hold the control key for multiple selections. The default is any client type.

In-Path Interface:
Allows you to designate the authorization realm to only apply to traffic that flows on a particular VLAN.

NOTE: When assessing authorization realm membership, if the appliance determines that if all of the above field criteria are met, then the user is deemed to be a member of the authorization realm.

Using Authorization Realms for Testing:
Authorization realms can be useful for testing because they provide a convenient means to stage deployment of new application policy.

- You can define an authorization realm with just one IP address – that of a test machine.
- Then you can enable an application policy and limit the deployment to that authorization realm.
- When you are satisfied that further deployment is warranted you can broaden the scope of the authorization realm for further staging or you may choose to delete the authorization realm altogether from the policy.

7.1.5 Web Application Policies

Purpose

The Web application policies control how the FX Series interacts with various web sites. This will allow users to enable, disable, and tune acceleration techniques for specific URLs. A Web application policy is defined by two parts: a Target web server (HOST) and an Application (the trailing text of the URL following the host portion).

Example:

In this example the user has requested the web site **"www.dailygalaxy.com"** which is the target web server or host. The page or application is "my_weblog/2011/09/image-of-the-day-firey-icy-beauty-of-an-infrared-neptune.html". We could enable or disable features for this web site by entering "**www.dailygalaxy.com**" as the Target Web Server in the web application policy and entering and star / asterisk "*****" as the application wildcarding all request to the defined web server. We can also gain greater granularity by adding portions of the application to the application field.

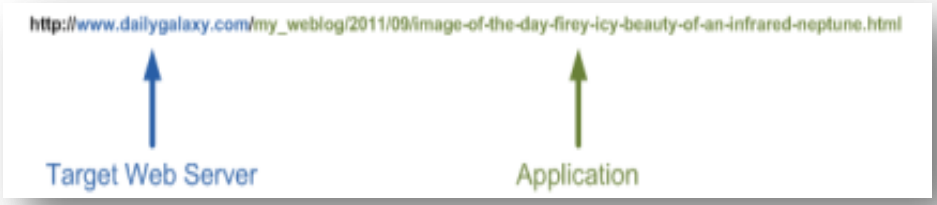

NOTE: How to set up Web Application Policies is shown in Section 7.2

A list of previously defined Application Policies is displayed in the order in which they were defined. An existing entry may be chosen by clicking on the Policy.
You can "Enable", "Disable", or "Delete" one or more Policies definitions by selecting the checkbox to the left of the Title column and clicking on the desired button. By clicking on "Add" you can add a new Application Policy which will bring up the Application Policy definition screen.

NOTE: Policy Deletions are allowed 10 minutes after all active client connections using the policy have terminated.

Target Web Server

The case-insensitive name of the web server for this policy can also be a virtual server name that corresponds to a pool defined in the "Load Balancing" configuration section. The name can be a fully qualified domain name or it can be a short name (no dots in the name). This name is relative to the browser and is often not the same name that the FX Series ADC uses to access the content when it is functioning as a forward or reverse proxy. It is possible to configure both the short name and a fully qualified name with different policies. In this case, the policy in effect will depend on whether the URL has just a short name or a fully qualified domain name.

NOTES: If a non-standard port (other than 80) is used to access the application, the port must be specified as part of the name.
An IP address may also be specified in lieu of a host name.

The server name can be configured with a leading "*". This allows an administrator to configure all servers within a domain. An example of this is "*.acme.com" which means all servers in the acme.com domain. If '*' or empty then the policy applies to all servers running the specified application. Do not enter the "**http://**" prefix when specifying this field. The administrator can configure a default server by

having a policy that is either "*" or empty. If the FX Series does not find a match for the server then it will look for this default server.

Application
The case-insensitive application name for this policy can be a partial name. For example, if "abc" is specified, then this policy will apply to a URL containing "abc/def". When the FX Series ADC matches a URL against its application policies it will apply the most specific match that it encounters. If "abc/def" is specified in addition to "abc", then this policy will be used for a URL containing "abc/def/ghi".

NOTE: Additionally, the application name of the policy can be "*" or empty. This is the default application policy for this server. If no other match is found then this policy will be used. It is possible to configure both a default server and a default application. This becomes the default policy that is used when no server and application names match.

7.1.6 Authorization realm

This is a pull-down selector that enables you to limit the scope of this policy to a specific previously defined authorization realm. Authorization realms (see Section 7.1.4 above) allow users to be grouped based on source IP address or in-path interface.
Comment: Up to 80 characters of text can be entered to help document your configuration.

7.1.7 Enable Acceleration

Acceleration can be turned on or off.

- When Enable Acceleration is "Yes", the FX Series ADC will perform the acceleration techniques that are enabled in this policy.
- When Enable Acceleration is "No", then, no acceleration techniques will be applied for applications specified by this policy.
- When a FX Series Remote is used and "Enable Acceleration is set to "No", then the FX Series Remote will pass the HTTP requests directly to the application server and bypass the FX Series ADC. The default value is "No".

The Comment field should be used to briefly document the configuration.

7.1.8 Allow Access

Access for an individual policy can be set to "Yes" to allow access for the URL associated with this policy. If se to "No", then the ADC will return an HTTP error code 403 if an attempt is made to access the URL associated with this policy. The default is "Yes"

7.1.9 Caching

Cache Data at ADC
This setting specifies if requested HTTP objects should be cached on the FX Series ADC. The "Guaranteed object freshness time" also applies to these cached objects. The default is to cache the HTTP objects ("Yes").

Cache Data at Remote
This setting specifies if requested HTTP objects should be cached on the FX Series Remotes (REM). The default is to cache the HTTP objects ("Yes").

Enable Cache Differencing:
The ADC will cache dynamic HTTP content when a FX Remote is deployed. On subsequent accesses to the same content, the ADC will send the changes rather than the entire updated content and the FX Remote will apply the changes to its cache before delivering the object to the browser. The default is to enable cache differencing ("Yes").

7.1.10 Content Validation

Guaranteed Object Freshness Time (seconds):
This setting determines the number of seconds that will elapse after storing an object in the client browser cache before the object must be verified for freshness. If the value is set to 0 then objects are verified every time the object is accessed. The default value is 10800 seconds (3 hours).

For Static Content:
The FX Series ADC optimizes data flow by keeping track of object expiration times. For native browsers, if the object has not changed then a HTTP 304 Not Modified response is sent from the FX Series ADC without requiring a round-trip between the FX Series ADC and the target web server. If the client browser is engaged with the FX Series Remote, then this operation is performed at the FX Series Remote eliminating round-trips between the browser and the FX Series ADC.

Override Expirations
This specifies whether the FX Series ADC should intelligently override the expiration date of an object based on its heuristics. This feature reduces round trips by eliminating many unnecessary content validation requests from the browser. Enabling this option for caching in some instances could produce undesirable results. The default value is "Yes".

The object must be one of the content types specified in the "For the following content types:" field**.** A comma separates each content type. The default value for this field is: "image/*, text/css, application/x-Javascript, application/x-shockwave-flash, and text/javascript"

7.1.11 Image Optimization

JPEG Transformation
This function can lower the resolution of a JPEG in order to reduce communication bandwidth requirements. The goal of the JPEG quality and smoothing values is to reduce the amount of data while maintaining a usable image. The default settings (Quality=50, Smoothing=20) are reasonable settings for most images and give a dramatic savings on the size of the jpeg images. The image is still very usable but is not as crisp as the original. Depending on the jpeg, the compression is often in the range 9:1. If image quality is paramount, setting Quality to 75 and Smoothing to 0, will still give some image size savings with little if any loss of image integrity. There is a tradeoff between quality image compressions while still yielding some savings of image sizes. The optimal setting is application dependent.

JPEG Quality
A number between 1 and 100 that specifies the tradeoff between size of the jpeg data and quality of the original image. A higher number will retain a higher quality but will not conserve as much bandwidth. The default value is 50.

The JPEG quality setting allows you to trade off compression against the quality of the reconstructed image: the higher the quality setting, the larger the JPEG bandwidth requirements, and the closer the output image will be to the original input.

JPEG Smoothing
A number between 1 and 100 that specifies the amount of dithering that will be used when generating the image. A higher value indicates more dithering and somewhat less desirable image. The default value is 20.

Minimum size threshold in K
This setting sets a minimum size of the JPEG file that is required for the JPEG optimization processing to engage. The default size is 5K.

7.1.12 Back-End Server Interface Options

These options govern how the FX Series ADC will communicate with the back-end Web application server.

Request GZIP content:
If set to "Yes" then the FX Series ADC will request that the backend servers deliver the content to the FX Series ADC in GZIP format. This saves bandwidth over the backhaul connection and is suitable for most ISP environments. If set to "No", then the back-end server will deliver the content to the FX Series ADC in non-compressed form. "No" may be a good setting for an enterprise server environment where the FX Series ADC is offloading the GZIP task. The default setting is "Yes"

Web server timeout (seconds)
This sets the number of seconds that the FX Series ADC will wait for a request to be fulfilled from the enterprise web server. After the timeout a 502 HTTP error code will be returned to the browser that initiated the request. The default value is 60 seconds.

Use reduced window scaling:
If set to "Yes" then the ADC will use a smaller window scaling factor when connecting to back-end servers. The default value is "No".

7.1.13 When Application Policies Take Effect:

Application policy changes take effect on an FX Series ADCs immediately after changes are made. Once this occurs, the FX Series Remote will be notified of the policy change at the next object retrieval request. At that time, the FX Series Remote will read all of the current policies from the FX Series ADC, and it will adhere to those policies for the applications for which that FX Series ADC is being used for optimization. The new policies will not affect the FX Series Remote interaction with other FX Series ADCs (if it is interacting with multiple FX Series ADCs concurrently).

7.1.14 Web Application Firewall

The FX Series ADC offers Buffer Overflow Prevention capabilities to an enterprise infrastructure. Detection and blockage of overflow attacks which otherwise could lead to malicious code being run on a server or other unauthorized access. This is an inherent feature of the FX Series ADC - no special configuration is required. Any buffer overflow detected in an HTTP message will cause an event to be recorded in the exception log.

7.2 How to Configure Basic Web Application Policies

This section is designed to delineate how to set up application policies. All of the detail descriptions for each item are provided in Section 7.1.5 above. Application policy changes take effect on an FX Series ADC immediately after changes are made.

7.2.1 How to Set the Policy Applicability

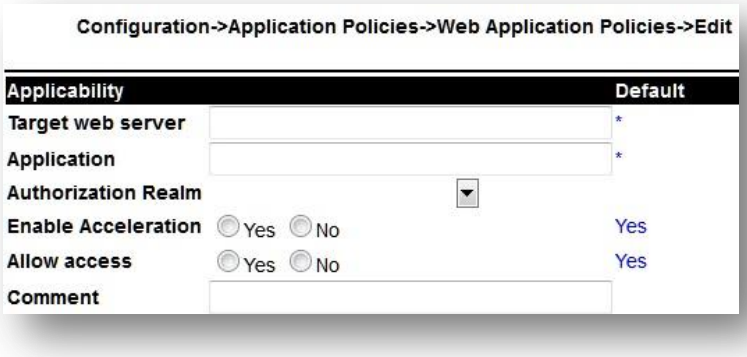

Figure 7-2 FX Series Application Policy Applicability Edit Screen

1. **Set the case-insensitive name of the Web server for this policy.**
   The server name can be configured with a leading "*".
   This allows an administrator to configure all servers within a domain

2. **Set the case-insensitive application name for this policy.**
   This can be a partial name. Additionally, the application name of the policy can be "*" or empty. This is the default application policy for this server. If no other match is found then this policy will be used.

3. **Set the authorization realm to define the users who will use this policy**
   If the default authorization realm is used, then the policy will apply to all users.
   Setting up authorization realms will be described in the next section
   Comment: This field is used to briefly describe the purpose of the policy.

4. **Enable or disable acceleration for this policy**
   Acceleration can be turned off or on. When Enable Acceleration is "Yes", the FX Series ADC will perform the acceleration techniques that are enabled in this policy. When Enable Acceleration is "No", no acceleration techniques will be applied for applications specified by this policy. The default is enabled.

5. **Allow or disallow Access for this policy**
   If access is set to 'No', then the ADC will return an HTTP error code 403. If an attempt is made to access the URL associated with this policy. The default setting is 'Yes'.

When complete click on "Add Application Policy"

Note: A "default policy" is a policy which allows access to all servers, all applications for all users. It is created with a "*" in these fields: Target Web Server, Application Policy. In addition it uses the default user authorization realm, which is set to all users. Acceleration can be enabled or disabled for this policy. If disabled, then you would want to set other policies to allow users access to specific site(s).

7.2.2 How to Set Specific Users Access

This section allows you to restrict the application policy to a specific set of users/IP addresses. All of the detail descriptions for each item are provided in Section 7.1.4 above. The standard setting preconfigured in the default settings is: *all sites available and accelerated for all users with all client types on any port.* NOTE: Authorization realms can be useful for testing because they provide a convenient means to stage deployment of new application policy by limiting the users who can access the new application policy.

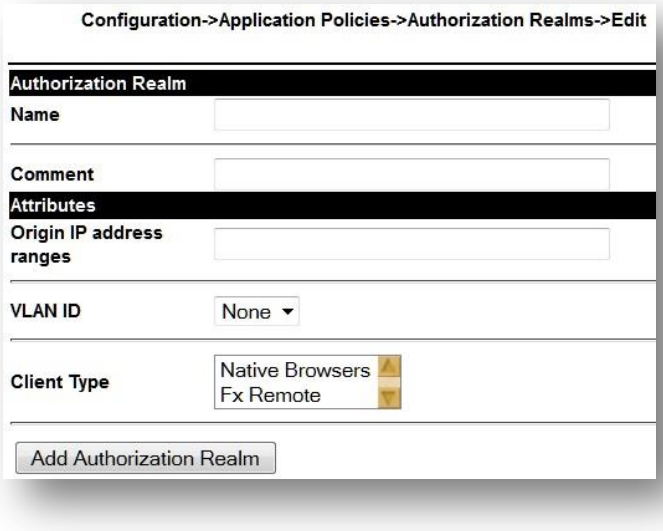

Figure 7-3 FX Series Authorization Realm Edit Screen

1. **Set the authorization realm name**
   This specifies the logical name to assign to this realm. This name is used to reference the definition in the application policies and client policies and can be up to 64 characters of text.
   Comment: This field is used to briefly describe the users that this realm applies.
2. **Set the origin IP address ranges for the users that this realm applies.**
   You may enter one or more single IP addresses or hyphenated IP address ranges, separated by commas in the same manner. i.e. 10.2.2.5 or 10.2.2.50-10.2.2.59. The default setting is any network.
3. **Set the client type that this realm applies.**
   This field allows you to specify that client types that are associated with this authorization realm Pressing the control key allows you to choose multiple selections. The default is any client type.
4. **Set In-Path Interface:**
   Allows you to designate the authorization realm to only apply to traffic that flows on a particular VLAN.
5. **When complete click on "Add Authorization Realm"**

7.2.3 How to Restrict Acceleration for Specific Sites, or Users

- To restrict acceleration for specific sites; set up application policy for each specific site with acceleration disabled in the policy
- To restrict acceleration for specific users; set an application policy for specific authorization realm with acceleration disabled in the policy

- To restrict all acceleration except for specific sites for *certain user groups*

  1) Set a * application policy with acceleration disabled in the policy for the default authorization realm.
  2) Set an application policies for the specific sites with acceleration enabled in the policy for a specific authorization realms

7.2.4 How to Set Specific Optimization Techniques

Each of these options can be enabled or disabled by clicking on the preferred choice. The inherited default for all acceleration options is "Yes" for enabled except the reduced window scaling option. All of the detail descriptions for each item are provided in Section 7.1

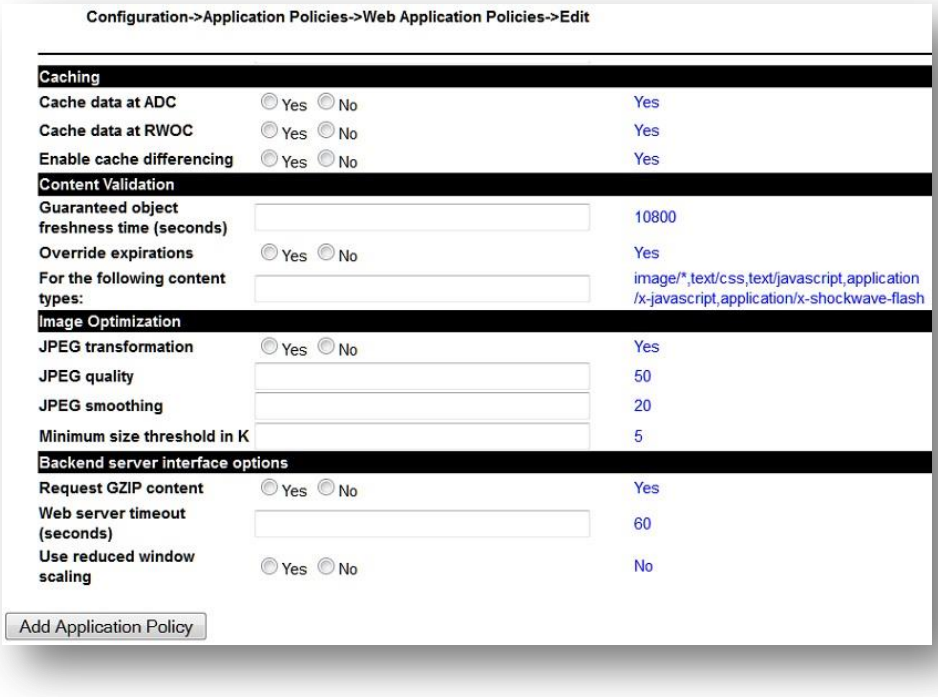

Figure 7-4 FX Series Specific Optimization Edit Screen

NOTE: Most application policies use the inherited settings. Changes for specific content types or specific JPEG quality/bandwidth requirements other than the inherited values can be made as noted below:

For the following content types
The default value for this field is "image/*, text/css, application/x-javascript, application/x-shockwave-flash, text/javascript". Additional items can be added to this line and placing all items in the field.

JPEG Quality
The default is 50. The range of the field is 1 – 100. A higher number will retain a higher quality but will not conserve as much bandwidth.

JPEG Smoothing
The default is 20. The range of the field is 1 – 100. A higher value produces a less desirable image.

7.3 Layer 5 Application Policies

"Layer 5" acceleration provides optimization of communication protocols when a FX Series Remote is installed at a branch in "in-line" mode. The FX Series Remote can then intercept and optimize the traffic with additional capabilities such as dynamic data suppression and traffic classification/prioritization.

Layer 5 application policies are defined at the ADC and then read by the FX Remote to determine how it should process the traffic that it intercepts as an in-line bridge. A list of previously defined Application Policies is displayed in the order in which they were defined. An existing entry may be chosen by clicking on the Policy.

You can "Enable", "Disable", or "Delete" one or more Policies definitions by selecting the checkbox to the left of the Title column and clicking on the desired button. By clicking on "Add" you can add a new Application Policy which will bring up the Application Policy definition screen.

NOTE: Policy Deletions are allowed 10 minutes after all active client connections using the policy have terminated.

7.3.1 How to Configure Certified Applications

This is a pull-down of pre-defined communication protocols that have been certified as compatible with FX Series application acceleration. These include CIFS, FTP, HTTP, MAPI/Exchange, NFS, POP3, SMTP, POP3, FX Configuration Traffic, Diagnostic Traffic, Other. If you select one of the pre-defined choices then the fields of the policy will be populated.

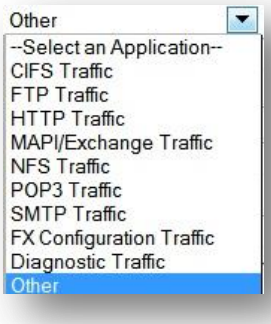

7.3.2 How to Configure Other Applications

If the traffic that you wish to accelerate is not listed then select "Other" and you can set the policy field based on the characteristics of the protocol.

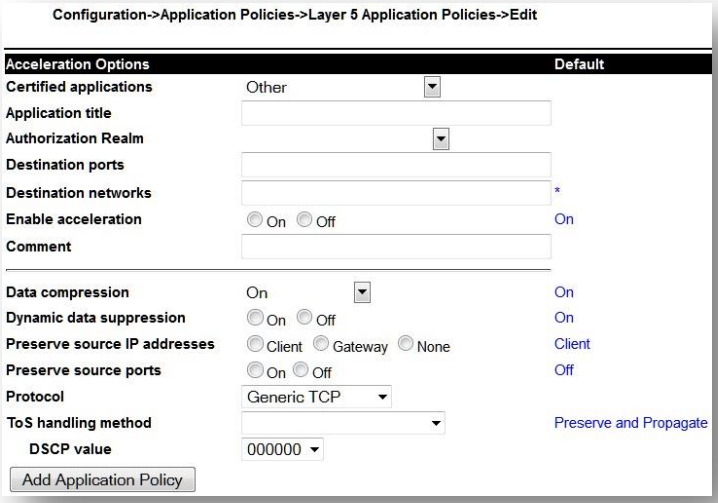

Figure 7-5 FX Series Layer 5 Policy Configuration Edit Screen

Application title:
This field applies a name to the traffic being accelerated.

Authorization Realm:
This is a pull-down selector that enables you to limit the scope of this policy to a specific previously defined authorization realm. Authorization realms allow users to be grouped based on criteria such as source IP address.

Destination ports:
This mandatory setting allows you to specify the TCP port(s) that apply to this policy. Multiple ports may be specified separated by commas or a port range may be specified with a dash. Multiple port ranges may be separated by commas. The minimum port value is 1 and the maximum value is 65535. There is no default setting.

NOTE: A "*" policy for L5 can be defined. A "*" policy is a port range of 1-65525.

Destination networks:
This setting allows you to limit the Layer 5 acceleration only to traffic that is destined for certain networks. The destination network is specified using "CIDR" notation where a base IP address is followed by a '/' character which is followed by a value between 1 and 32 that denotes the number of bits used to describe the network and the remaining bits (32 – the value) are used to specify the nodes on that network. For example a setting of 192.110.1.0/24 would be equivalent to specifying a network of 192.110.1.0 with a net mask of 255.255.255.0. Separating each CIDR entry with a comma can specify multiple destinations. The default setting is all networks.

Enable Acceleration:
If enabled then the FX-Remote will apply the acceleration techniques based upon the specifications of this policy, otherwise it will not intercept the traffic specified in this policy.
This setting is useful for defining exceptions to a broad port range defined in other policies.
The default setting is 'On'.

Comment: This field is used to describe the rationale for this configuration.

7.3.3 How to Configure Layer 5 Optimizations

Data compression:
If "On", and the remote machine is enhanced with the Layer 5 software, then bi-directional compression of all traffic is performed. If "Off" is specified, then no compression will be performed.

If "Only-to-client" is set, then data compression will only be performed on traffic that flows from the server to the client. If "Only-to-server" is set, then data compression will only be performed on traffic that flows from the client to the server. The default is on.

Dynamic data suppression:
If "On", the FX Series will maintain a byte cache of network traffic and replace repeated patterns with signatures when they are detected in the data stream. The default value is on.

Preserve source IP address:
This specifies which source IP address should be presented to the back-end server when the ADC makes requests on behalf of a remote client. If "Client" is specified then the client's IP address will be used. If "Gateway" is specified, then the IP address of the gateway will be used, this is sometimes useful when Network Address Translation makes it unfeasible to use the client's IP address.
If "None" is specified, then the ADC's IP address is used. The default setting is "Client".

Preserve source ports:
This specifies if the ADC should use the same source port as the client. If "On" then "Preserve source IP address must be set to "Client". If this is set to "Off", then an arbitrary port will be used. Some applications such as "NFS" require this to be set to "On". The default setting is "Off".

7.3.4 Layer 5 Protocols

This setting allows a protocol-specific optimization module to process the traffic. The choices are as follows:

Generic TCP:
Perform "Dynamic data suppression" QOS traffic prioritization, compression, TurboStreaming on the traffic.

HTTP:
In addition to the Generic TCP optimization, performs intelligent object caching, cache differencing, turn reduction and other techniques as specified in L7 application policies on the traffic.

CIFS:
In addition to the Generic TCP optimization, perform optimizations to overcome unneeded latencies in the CIFS application protocol.

MAPI:
In addition to the Generic TCP optimization, performs optimizations to overcome unneeded latencies in the MAPI application protocol such as used in Outlook/Exchange server communications.

SMTP/POP3:
This optimizes standard internet email traffic.

FTP:
In addition to the Generic TCP optimization, this performs optimizations to allow better DDS when communicating with an FX Series Remote.

Diagnostic:
This specifies parameters when running diagnostic tests between FX Series Remote and FX Series ADC.

FX Configuration:
This protocol entry is utilized by FX-Remotes for ADC location services. FX Remote appliances will use this entry to locate configuration information when Auto-discovery is not enabled. The default is for Auto-discovery to be enabled so normally this entry is not used. This entry is auto-generated when new in-path interfaces are created and should not be deleted. If this entry is deleted it can be easily added back by selecting 'Add' from the policy list screen and selecting it from the protocol list.

7.3.5 ToS handling method

Specifies how you want the FX to handle "Type of Service" (ToS) bits in the IP header for accelerated traffic. The default method is to "Preserve and Propagate" which means that ToS bits for request packets will be preserved through the acceleration tunnels and propagated in the responses. "Preserve" means that the ToS bits will be preserved for request traffic but not propagated in the responses. "Set" means that the FX will apply the ToS bits as specified in the "DSCP value" field in both the request and response traffic.

The DSCP value: Specifies the ToS bits based on the binary value in the "DSCP value" field. DSCP is an abbreviation for "Differentiated Services Code Points" and is described in more detail in the RFC-2724 specification**.**

7.3.6 Layer 5 Acceleration - Discussion

"Layer 5" acceleration provides optimization of communication protocols when a FX Remote is installed at a branch in "in-line" mode. The FX Remote can then intercept and optimize the traffic with additional capabilities such as dynamic data suppression and traffic classification/prioritization. Most TCP based protocols can be accelerated with this solution, however there are some exceptions:

- FTP can only be accelerated when it is operating in "passive" mode. By default the ftp function of the Internet Explorer operates in "active" mode and you must adjust an advanced browser setting in order for Layer 5 to operate. Also the standard Windows command line "ftp.exe" program only operates in "active" mode and therefore can't be accelerated by Layer 5.
- Due to the nature of some TCP/IP based protocols; The Layer 5 acceleration techniques can function however they do not provide significant benefit. For example:
  - The latest version of MS-Outlook/Exchange protocol is already highly compressed and also there is an application level acknowledgement that takes place after every 17K bytes which defeats the value of TurboStreaming. (MS-Outlook using the POP3/SMTP protocols can derive great benefit from Layer 5 acceleration.)
  - Protocols that are encrypted (or already compressed) will not benefit from the product's Layer 5 compression but may benefit from the TurboStreaming.

The following table depicts a list of TCP and UDP ports used by well-known applications.

Application	TCP / UDP	Ports
CIFS	TCP	139,445
Citrix ICA	TCP/UDP	2598,1494
CVS	TCP	2401
FTP	TCP	20,21
Lotus Notes/Domino	TCP	1352
Lotus Sametime	TCP	1533
MS Remote Desktop Connection	TCP	3389
NFS	UDP	2049
Novel GroupWise	TCP	1677
POP3	TCP	110
SIP (Session Initiation Protocol)	TCP/UDP	5060,5061
SMTP	TCP	25
Telnet	TCP	23
VNC	TCP	5900

Figure 7-6 FX Series TCP/UDP Ports Table

8 FX Series Operations Features

Figure 8-1 FX Series Operations Menu

Backup/Restore Configuration

This facility allows you to backup or restore your FX's configuration files.

Change Passwords

This function allows you to change the administrative interfaces to the FX.

Manage Licenses / Fast Codes

This function allows you to view and replace license files or Fast codes. These are used to license the specific appliance, turn on the Multicator functionality and set the Max Data Rate.

Shutdown/Restart

This page allows you to perform software restart operations or to power down or reboot your FX hardware in an orderly fashion.

Packet Capture

This function allows you to collect WireShark compatible packet captures on up to two interfaces concurrently.

Update Software

This facility allows you to upload and install the latest acceleration server images

FX Series Support

This facility provides links to FX Series Documentation, SNMP MIBS and various support links for Technical Support, Software and RMAs

8.1 Basic Operations Functions

8.1.1 How to Backup/Restore Configuration Files

This screen allows you to back up the acceleration server configuration files so that you can quickly restore your FX as part of a disaster recovery procedure. The files include hostname, time-zone, and IP addresses for the management and in-path interfaces, DNS server, settings, application policies, and SSL security certificates. The files are backed up to a single ".zip" file which is then downloaded to your desktop workstation.

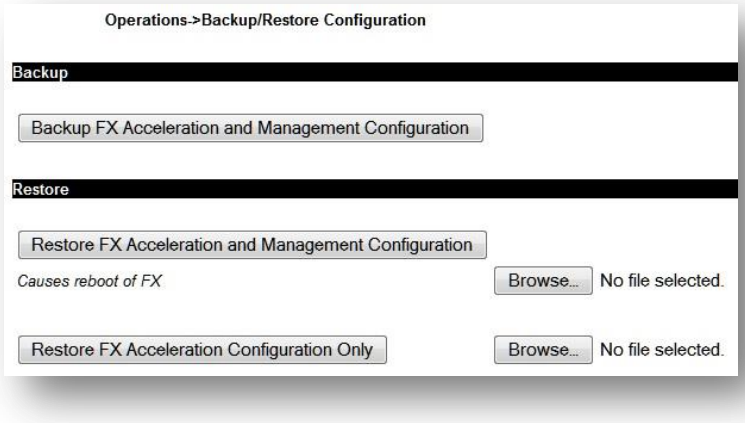

Figure 8-2 FX Series Backup and Restore Screen

Backup FX Acceleration and Management Configuration:
This button causes all FX configuration files to be stored into a zip file with the name "fxbackup-full_*hostname*_year-month-day-hour-minute.zip". Following this, a dialog will appear so that you can save this file to your desktop.

Restore FX Acceleration and Management Configuration:
This button should only be clicked after selecting a backup zip file using the "Choose File" button. Following this, clicking the "Restore Acceleration Appliance" button will restore both the acceleration and management settings from the specified zip file. The specified zip file should have been previously created using the "Backup FX Acceleration Appliance and Management" button. The FX appliance will automatically be rebooted following this restoration so that new management settings are applied.

Restore FX Acceleration Configuration Only:
This button should only be clicked after selecting a backup zip file using the "Choose File" button. Following this, clicking the "Restore Acceleration Appliance" button will restore only the acceleration settings and exclude any management settings from the specified zip file. This is useful for transferring similar acceleration configurations from one FX to another without affecting management settings. The specified zip file should have been previously created using the "Backup FX Acceleration and Management Configuration" button.

NOTE: The appliance will automatically restart after the restore.

8.1.2 How to Initiate Disaster Recovery Procedure

In the event that you need to restore your acceleration server appliance from a USB flash drive, please complete the following steps:

1. Reinstall the installation image from USB Flash drive.
2. Using the “FX Series Appliance Manager”, set your passwords for the acceleration server administrator by selection "Configure Appliance", and then "Configure Passwords".
3. Using a browser, go to the FX’s “Operations” page and select "Backup/Restore Configuration". Then, restore the configuration from the backup zip file.
4. "Backup/Restore Configuration". Then, restore the configuration from the backup zip file.
5. From the “Operations” page, select "Shutdown/Restart Appliance", then click "Restart Appliance Hardware" to restart the acceleration appliance.
6. "Restart Appliance Hardware" to restart the FX.
7. Using the “FX Series Appliance Manager”, re-establish the high-availability keys for the cluster-mates that share a common acceleration server configuration. Configuring the high-availability cluster keys can be accomplished by selecting "Configure Appliance", and then "Configure Passwords".
8. Re-apply any patches that you may have applied subsequent to your installation via the Web Administrative GUI.

8.1.3 How to Change Passwords

This function allows you to change the administrator password of the FX Series appliance. If you do not remember your password you will have to re-initialize the FX Series appliance. To change your password you must enter it twice and then press the corresponding “Change” button.

Figure 8-3 FX Series Change Passwords Screen

8.1.4 How to Manage Licenses / Fast Codes

This function allows you to install and update the license file or fast codes for the FX Series appliance. For all new installations, Fast Codes will be used for upgrading functionality and Max Data Rates

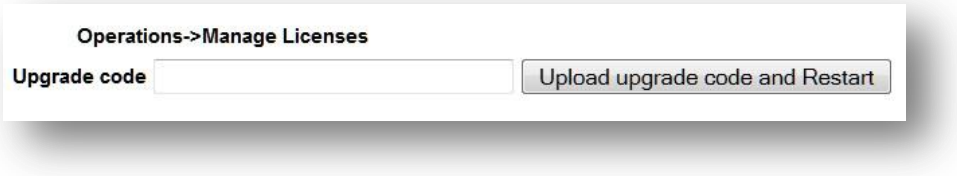

Figure 8-4 FX Series Upgrade Fast Codes Screen

Fast Codes
The Fast Code is a 14 digit key that will enable the WAN Optimization and the max data rate.

Upgrading Fast Code: Enter the character sequence provided by Comtech, and then click the 'Upload Upgrade Code and Restart' button. If successful, this action will immediately restart the acceleration service and the new upgrade license information will be displayed.

NOTE 1: There are 'trial license' fast codes for 30/60/90 days for Packet Compression and WANOP. The remaining time trial license period will be noted in the license display at the left bottom of the Web GUI screens with the firmware version and the serial number.

NOTE 2: For existing installations, where a license file is currently being used, an upgrade is available to move to Version 6.0.3 for the new features, including Multicator. License files will continue to be used for these appliances.

8.1.5 How to Shutdown/Restart

This page allows you to perform software restart operations or to power down or reboot your FX hardware in an orderly fashion.

Software

Restart Acceleration Service

This button will start the acceleration service if it is in the "stopped" state or restart it if it is in the 'active' state

Restart Acceleration Service with Full Cache Rest

This button will start the acceleration service if it is in the "stopped" state or restart it if it is in the 'active' state. Upon doing so, it will reset all the cache files.

Start or Stop Acceleration Service

This button will either start or stop the acceleration service depending upon its current state. When 'stopped', access to the appliance via the management interfaces is still possible. If the appliance is configured for "in-band" management then this button will not be displayed.

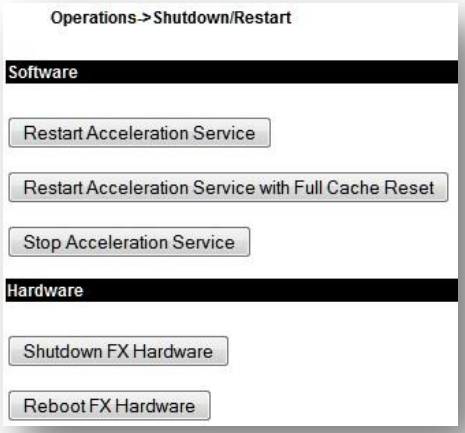

Figure 8-5 FX Series Shutdown/Restart Screen

Hardware

Shutdown FX Hardware

This button will issue a command to the native operating system to shut down the software in an orderly fashion and then power down the appliance.

Reboot FX Hardware

This button will issue a command to the native operating system to shut down the software in an orderly fashion and then reboot the appliance.

8.2 How to do Network Trouble Shooting with Packet Capture

This function allows you obtain packet captures to facilitate troubleshooting of network connectivity problems.

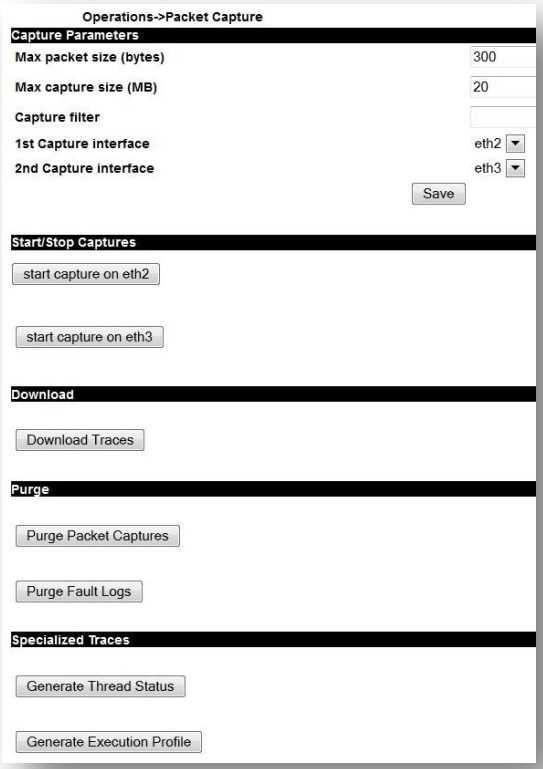

Figure 8-6 FX Series Packet Capture Screen

Capture Parameters

Max packet size (bytes):

Specifies how much of each packet to store in the capture file. The max size is whatever your MTU is set to for the interface (up to 9000 for jumbo frames) and the default size is 1500. In this case it is set to 300.

Max capture size (MB):

This specifies the maximum size of the capture file before rotating to a second file. The default size is 10 megabytes.

Capture filter:

This specifies a filter that will reduce the size of the capture file by recording only the traffic which matches the filter.

1st Capture interface:

Select the Ethernet interface that the actions (e.g. “start” “stop”) of the first capture will be associated with. In this case it is eth3.

2nd Capture interface:

Use the pull down and select the Ethernet interface that actions of the second capture will be associated with. In this case it is eth2.

Save Button:
This button must be clicked in order to apply them prior to performing the capture actions.

Capture Actions

After the common parameters have been “Saved”, you can click the “start” or “stop” buttons associated with the capture interfaces that have been defined.

Download Traces

After stopping the packet captures, click the “Download Traces” button to invoke a log collection process that will zip up the packet captures and also an assortment of event tracing logs. The captures in the “.zip” file will have names that end in either “.pcap0” or “.pcap1”.

Purge

You can reduce the size of the .zip’ download file by first purging old packet captures and fault logs.

Specialized Traces

You can add additional specialized trace logs that are downloaded to obtain the current thread state of the acceleration service and/or a 15 second execution profile of the acceleration service.

Sample Capture Filters

host 10.1.1.1

Records only traffic for which either source or destination is 10.1.1.1

tcp port 443

Records only traffic for which either source or destination is port 443

tcp port 443 and host 10.1.1.1

Records only traffic that has source/dest port 443 AND also has source/dest of 10.1.1.1

8.3 How to Update Software

This selection allows you to apply available FX software updates to your systems.

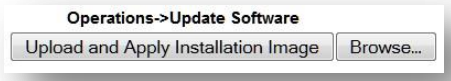

Upload and Apply Installation Image:

This button can only be clicked after selecting a file relative to your desktop workstation using the "Choose File" button. After picking the updated Acceleration Server software from the location that you saved it, click on the "Upload and Apply Installation Image" button to upload the software and install it on the system.

Download and Apply Image from ADC (FX Remote Only):

This button downloads the FX Remote firmware image that is included with the ADC and applies it to the FX Remote. This button is only shown if the FX Remote is in a state where it can download the image. This state is only achieved if:

1. WAN Optimization is licensed
2. There are in-path interfaces defined which are enabled.
3. The FX Remote can successfully connect to the ADC over an In-Path interface.

8.3.1 Software Update Discussion

Requirements for Updating FX Series Software to Version 6.2 or higher

This function allows you to apply available acceleration software updates to your FX appliance.

NOTE: You must have run the Base Platform Image upgrade (BPI3) and currently be at Version 5.78.0 or higher before you can upgrade to Version 6.2 or higher.

The Comtech FX Series version 6.1 includes features that will require a Base Platform Image (BPI) service pack upgrade to enable the functionality. It is recommended that it is installed for all WANOP enabled appliances. Note that **only** the FX 6.1 and above versions are compatible with the 3.7.9-FX27 service pack.

Determine the Current Software Version

Check version on bottom left of any Web GUI page. This indicates build 6.2.0 and service pack 3.7.9-FX27 (This appliance is updated with an earlier software version and a prior service pack)

Figure 8-7 FX Series Software Version Display Screen

If there is a more recent software package with a version for your OS, you can upgrade.

NOTE: To utilize the newest functionality you should upgrade to Version 6. 2+.
If you have not updated the BPI to v3, see the process for updating to Version 5.78.1+ with the new Base Platform Image (BPI) Appendix Section 10.3 below

- Determine availability of new software

  For the latest Version 6.2.2 + (V620 or higher) and 3.7.9-FX28a Service pack F0020713 or higher)
  Go to *www.comtechefdata.com* > support > software
  Click on "Stampede FX Series" in Table.

- Upload and Apply Server Installation Image Version 6. 2+:

  This button can only be clicked after selecting a file relative to your desktop workstation using the "Choose File" button. After picking the updated Acceleration Server software from the location that you saved it, click on the "Upload and Apply Server Installation Image" button to upload the software and install it on the system. Do the same process for the updated service patch if there is one available.

8.3.2 How to Download and Apply Image from ADC (FX Remote Only):

An FX Remote polls the FX ADC every 45 minutes to determine if a software update is available. If a newer version is available and the FX ADC is configured to "Automatically distribute remote updates" then the FX remote will automatically download the update and apply it.

Applying the firmware update, regardless of the method used to acquire the image, results in about a 30 second service disruption.

A recommended practice for deploying new FX image updates is as follows:

1. On ADC, navigate to "Configuration->General Settings". Turn off "Automatically distribute remote updates".
2. On ADC, navigate to "Operations->Update Software". Use the "Upload and Apply Installation Image" to update the firmware on the ADC.
3. After the firmware update is applied to the ADC, plan a time window when you would like firmware updates to be distributed to the FX Remotes. This time window could be an hour or two and is usually done when traffic is low.
4. At the time that you wish to distribute updates to the FX Remotes, on ADC, navigate to "Configuration->General Settings". Turn on "Automatically distribute remote updates".
5. Since each remote check every 45 minutes for an update, within an hour all FX remotes should have downloaded and applied the firmware update.
6. After all FX remotes have been updated, navigate to "Configuration->General Settings". Turn off "Automatically distribute remote updates".

To verify that the FX Remote devices have been updated, you can use SNMP to query the version number of FX Remote firmware. The OID of the version string is 11.3.6.1.4.1.6247.78.1.1.3. The following Windows console script would show the version numbers currently running at two FX Remotes using the net-snmp tools:

```
set netsnmp=c:\netsnmp
set SNMPCONFPATH=c:/netsnmp/etc/snmp;c:/netsnmp/snmp/persist
set SNMPSHAREPATH=c:/netsnmp/share/snmp
bin\snmpwalk -m all -v 2c -c public 172.27.101.206 1.3.6.1.4.1.6247.78.1.1.3
bin\snmpwalk -m all -v 2c -c public 172.27.101.244 1.3.6.1.4.1.6247.78.1.1.3
```

9 FX Series Support

9.1.1 Support Contact Information

Support Problem Report and Information Request form

This form contains summary information about the FX Series appliance which can be copied and pasted into an email to the Product Support 7x24 Contact organization. Please include a detailed description of the issue or question to facilitate our support for you.

This screen includes the internal serial number and the following configuration information:

- Date/Time of Information, FX Model, software version, and build patch version
- Options enabled
- Max throughput rate
- Interface Status

Note: the CEFD Serial Number can be found on the outside of the appliance

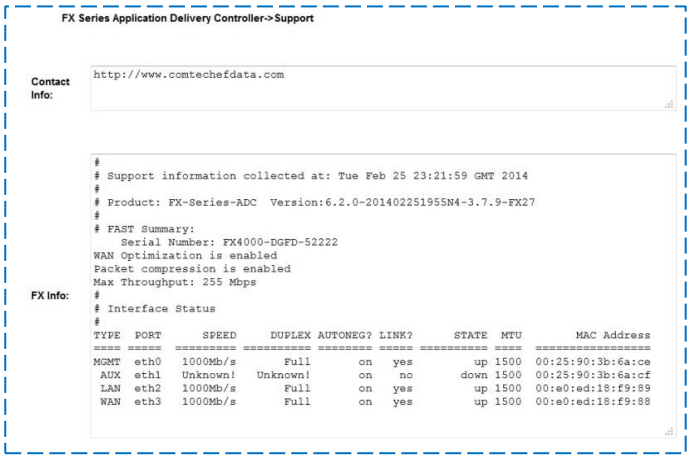

9.1.2 SNMP MIBS

These can be found on the appliance and there are two methods to connect to the Web GUI

1. To connect, log in via a web browser using its standard network addressing.
2. To connect directly to the FX Series GUI, using a PC with a web browser: Attach a cross-over cable to the eth1 interface and connect the browser via the browser interface at *http://169.254.55.55:10000*.

   - Login using the USERID and password “comtech” and “comtech”
   - Click on the “*Support*” tab from the main menu.
   - Click on SNMP MIBS to download.

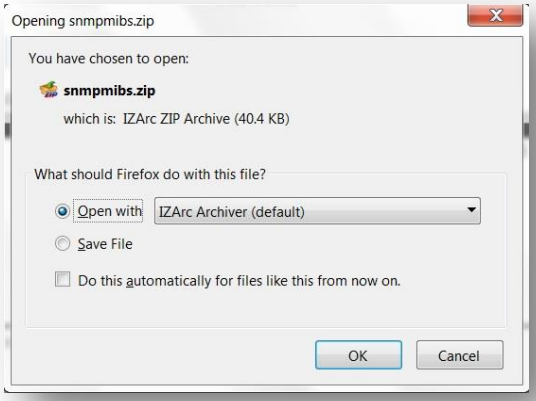

9.1.3 Product Information and Support Links

This page shows Links to FX Series Documentation. It also shows a window that contains summary information about the FX which can be copied and pasted into an email if a support issue or question arises.

These items are available at: ***http://www.comtechefdata.com/support***

- **Product Page**
  Data Sheets
  White Papers
  Case Studies and Testimonials

- **Software and Hardware**
  Software Downloads
  Product Support 7x24 Contact Info
  Warranty and RMA Requests

- **Manuals**

FX Administrative Guide

Includes FX ADC and FX Remote
In PDF and CD-ROM Formats

Note: Visit http://www.adobe.com in order download a reader for these files.

.

10 Appendix

10.1 Sample Acceleration Status Reports

Here are a few of the many reports available on the FX Series. Others have been included in the main sections as appropriate.

Acceleration and Throughput Statistics Reports

The Pull Down will allow inspection of acceleration statistics by acceleration category. Throughput Reports for ADC and Remote Aggregate Throughput are in the same format.

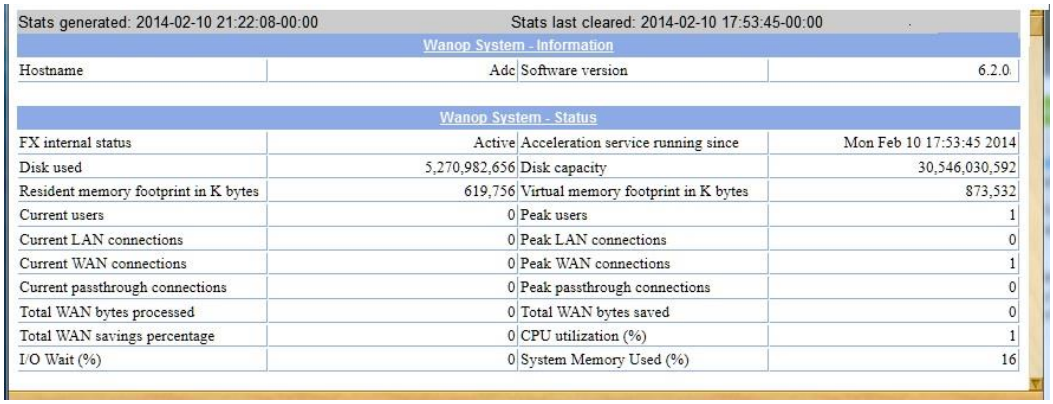

Stats generated: 2014-02-10 21:22:08-00:00 Stats last cleared: 2014-02-10 17:53:45-00:00

Wanop System - Information			
Hostname	Adc	Software version	6.2.0

Wanop System - Status			
FX internal status	Active	Acceleration service running since	Mon Feb 10 17:53:45 2014
Disk used	5,270,982,656	Disk capacity	30,546,030,592
Resident memory footprint in K bytes	619,756	Virtual memory footprint in K bytes	873,532
Current users	0	Peak users	1
Current LAN connections	0	Peak LAN connections	0
Current WAN connections	0	Peak WAN connections	1
Current passthrough connections	0	Peak passthrough connections	0
Total WAN bytes processed	0	Total WAN bytes saved	0
Total WAN savings percentage	0	CPU utilization (%)	1
I/O Wait (%)	0	System Memory Used (%)	16

Port Statistics

The Pull Down will allow inspection of acceleration statistics by port definition.

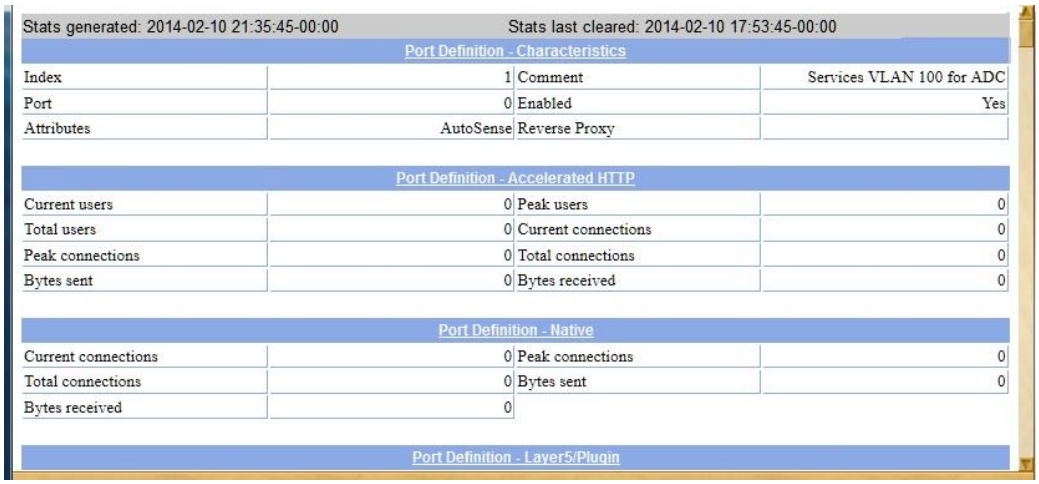

Stats generated: 2014-02-10 21:35:45-00:00 Stats last cleared: 2014-02-10 17:53:45-00:00

Port Definition - Characteristics			
Index	1	Comment	Services VLAN 100 for ADC
Port	0	Enabled	Yes
Attributes	AutoSense	Reverse Proxy	

Port Definition - Accelerated HTTP			
Current users	0	Peak users	0
Total users	0	Current connections	0
Peak connections	0	Total connections	0
Bytes sent	0	Bytes received	0

Port Definition - Native			
Current connections	0	Peak connections	0
Total connections	0	Bytes sent	0
Bytes received	0		

Port Definition - Layer5/Plugin

10.2 FX Series Console Management Functions

Here are some of the key screens for configuring many parameters on the FX Series Appliances.

The FX Series supports a basic menu-driven interface, which is accessible using the console port (eth0). It can also be accessed via SSH software once basic network connectivity is established.

Main Menu

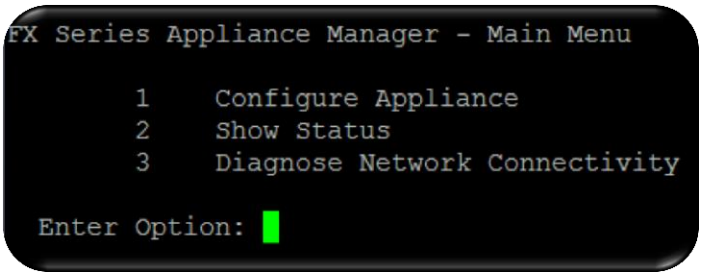

Configure Appliance

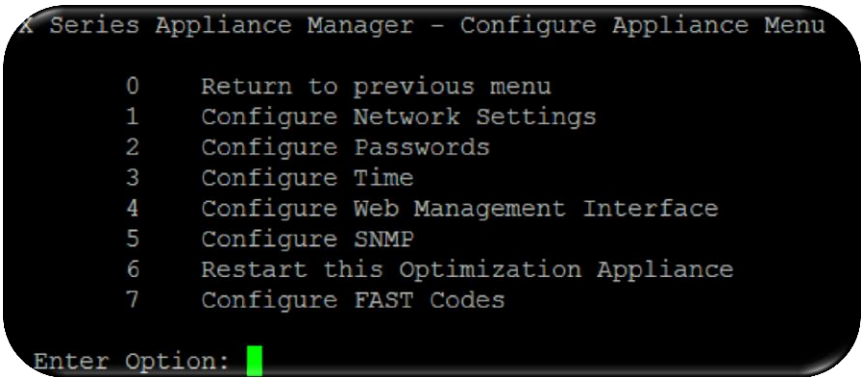

Configure Network Settings

Select “1 Configure Network Settings”

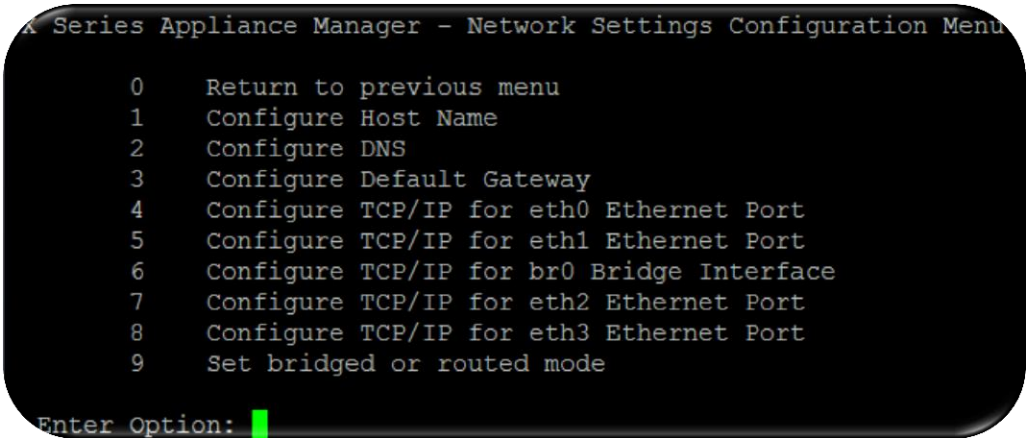

Configure SNMP Settings

Select “5 Configure SNMP”

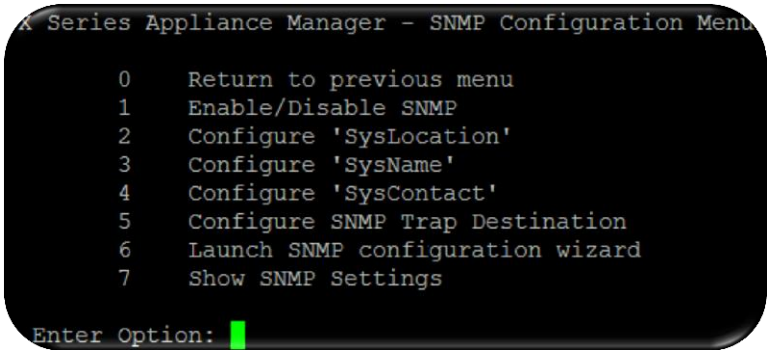

Configure FAST Codes

Select “7 Configure FAST Codes”

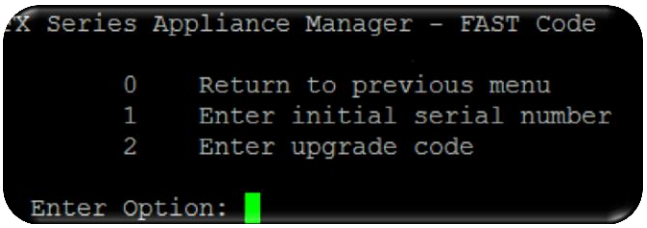

Show Status

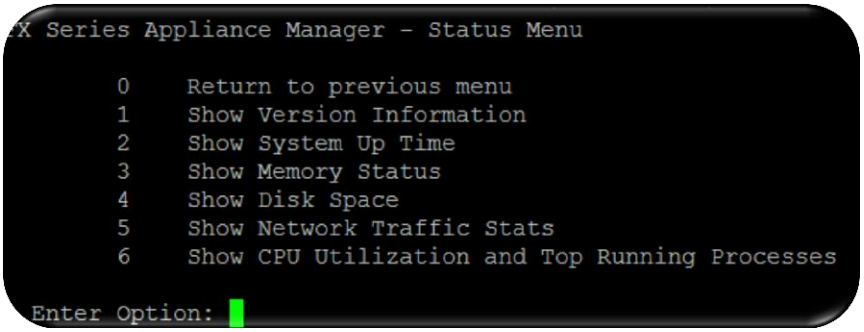

Diagnose Network Connectivity

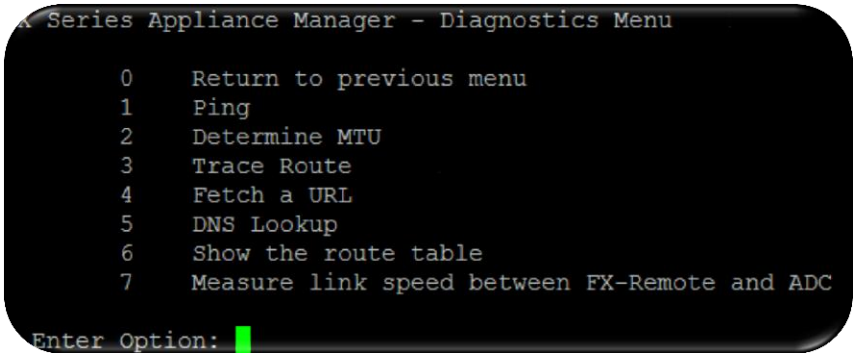

10.3 How to Update FX Series Appliance Software at 5.78.0 or earlier

This is a two-step process.

1. Upgrade the Base Platform Image and install version 5.78.0
2. Upgrade the application software to version 6.1 or later.

10.3.1 Base Platform Image (BPI) Upgrade Process

The Comtech FX Series version 6.1 or higher includes features that were not included in the previous Comtech FX Series releases. These will require a Base Platform system software upgrade.

This Upgrade kit enables upgrading from 5.78 or earlier releases to a platform capable of supporting the new FX 6.1 and higher functionality. Note that all FX 5.78 or earlier versions are compatible with this BPI upgrade. The completion of the BPI upgrade process will install version 5.78.0. The process to upgrade the appliance to Version 6.1 will follow the standard upgrade procedure.

10.3.2 Upgrade Kit and Prep

The upgrade kit consists of the following:

- This Guide
- A bootable USB drive that will automatically install the new FX Series BPI including FX Series 5.78.0 system software.

10.3.3 The Upgrade Process

1) You will need to save the current configuration file to your PC desktop.
2) The USB drive will upgrade the FX Series BPI to accommodate FX Series version 6.0.3.
3) It will ask for which appliance you are upgrading:
   a. FX Series ADC or FX Series Remote
4) Once you choose, it will install the appropriate Version 5.78 FX Series Software
5) You will then manually restore the saved configuration file.
6) You can then remove the USB drive and reboot the server.

Save Current Configuration File

Connect to the appliance using the Web GUI as described above.

1) Login in using “comtech” and your current password
2) Go to the “*Operations Screen”.*

3) Chose “*Backup/Restore Configuration”*
4) Click on “*Backup Configuration*”
5) Save configuration file. -
6) This file must be used after the BPI upgrade and before the upgrade to Version 6.0.3.

Configure “Boot Options” for USB Drive

1) With the appliance powered off, insert USB drive
2) Reboot the appliance, clicking on “*delete key*”
3) Select “*Boot Options*” tab and set USB flash drive as the first boot drive Set HHD as the 2nd boot drive
4) Choose *“(F10) Exit*” and save changes - system will reboot.

Running the upgrade process

After the system reboots, enter "Y" or "N" at the "let me ask again. Are you sure you want *to continue*" prompt.

NOTE: Entering "Y" will proceed to re-image the local hard drive.

1) The appliance will power down automatically when the imaging has completed.
2) Remove the Comtech USB drive.
3) Reboot the system and use the basic menu-driven interface as described above
4) You will be presented with "FX Series Appliance Manager" screen that includes which FX Series software (FX SERIES ADC or FX Remote) to install.
5) Select the appropriate choice to install.

Restore Current Configuration File

1) Log into the Web GUI and go to the "Operations Screen" as you did earlier
2) Chose "Backup/Restore Configuration"
3) Click on "Restore Configuration"
4) Browse to the configuration file that you saved earlier
5) Click on "Restore Configuration Files"

The FX Series BPI update is now completed and your system is on Version 5.78.0.
You are ready to now upgrade to FX Series Version 6.1. *See Section 8.3 above* for upgrade procedures.