
Foundry Switch and Router Installation and Basic Configuration Guide



FOUNDRY
NETWORKS
www.foundrynetworks.com

2100 Gold Street
P.O. Box 649100
San Jose, CA 95164-9100
Tel 408.586.1700
Fax 408.586.1900

January 2006

Copyright © 2006 Foundry Networks, Inc. All rights reserved.

No part of this work may be reproduced in any form or by any means – graphic, electronic or mechanical, including photocopying, recording, taping or storage in an information retrieval system – without prior written permission of the copyright owner.

The trademarks, logos and service marks ("Marks") displayed herein are the property of Foundry or other third parties. You are not permitted to use these Marks without the prior written consent of Foundry or such appropriate third party.

Foundry Networks, BigIron, FastIron, IronView, JetCore, NetIron, ServerIron, Turbolron, IronWare, Edgelron, IronPoint, the Iron family of marks and the Foundry Logo are trademarks or registered trademarks of Foundry Networks, Inc. in the United States and other countries.

F-Secure is a trademark of F-Secure Corporation. All other trademarks mentioned in this document are the property of their respective owners.

CHAPTER 1

GETTING STARTED.....	1-1
AUDIENCE	1-2
NOMENCLATURE	1-2
LIST OF PUBLICATIONS	1-2
WHAT'S NEW IN THIS EDITION	1-3
NEW FEATURES IN THE ENTERPRISE RELEASE 08.0.00	1-3
NEW FEATURES IN SERVICE PROVIDER 09.2.00	1-8
ENHANCEMENTS IN THE FASTIRON EDGE SWITCH RELEASE 03.4.0x	1-9
ENHANCEMENTS IN THE IRONPOINT-FASTIRON EDGE SWITCH (IP-FES)	1-10
ENHANCEMENTS IN BIGIRON MG8 AND NETIRON 40G RELEASE 02.2.01	1-10
LAYER 2 ENHANCEMENTS IN RELEASE 02.2.01	1-10
SECURITY ENHANCEMENTS IN 02.2.01	1-12
IPV6 ENHANCEMENTS IN 02.2.01	1-12

CHAPTER 2

INSTALLING A FOUNDRY LAYER 2 SWITCH

OR LAYER 3 SWITCH.....	2-1
UNPACKING A SYSTEM	2-1
PACKAGE CONTENTS	2-1
GENERAL REQUIREMENTS	2-1
SUMMARY OF INSTALLATION PROCEDURES	2-2
INSTALLATION PRECAUTIONS	2-3
GENERAL PRECAUTIONS	2-3
LIFTING PRECAUTIONS	2-3
POWER PRECAUTIONS	2-4
PREPARING THE INSTALLATION SITE	2-4
CABLING INFRASTRUCTURE	2-4
INSTALLATION LOCATION	2-4
INSTALLING OR REMOVING OPTIONAL MODULES (CHASSIS DEVICES ONLY)	2-5

INSTALLING MODULES	2-5
REMOVING MODULES	2-6
INSTALLING OR REMOVING REDUNDANT POWER SUPPLIES (CHASSIS DEVICES ONLY)	2-6
DETERMINING POWER SUPPLY STATUS	2-6
CHASSIS DEVICES – AC POWER SUPPLIES	2-6
CHASSIS DEVICES – DC POWER SUPPLIES	2-8
INSTALLING OR REMOVING A POWER SUPPLY (FASTIRON 4802 ONLY)	2-10
DETERMINING POWER SUPPLY STATUS	2-10
FASTIRON 4802 – AC POWER SUPPLIES	2-11
FASTIRON 4802 – DC POWER SUPPLIES	2-11
REPLACING FANS (4-SLOT AND 8-SLOT CHASSIS DEVICES ONLY)	2-13
REQUIRED TOOLS	2-13
DETERMINING WHICH FAN HAS FAILED	2-13
FOUR-SLOT CHASSIS	2-14
EIGHT-SLOT CHASSIS	2-15
REPLACING A FAN TRAY (15-SLOT CHASSIS DEVICES ONLY)	2-16
VERIFYING PROPER OPERATION	2-16
ATTACHING A PC OR TERMINAL	2-17
ASSIGNING PERMANENT PASSWORDS	2-19
CONFIGURING IP ADDRESSES	2-20
LAYER 3 SWITCHES	2-20
LAYER 2 SWITCHES	2-21
MOUNTING THE CHASSIS OR STACKABLE DEVICE	2-22
DESKTOP INSTALLATION	2-22
RACK MOUNT INSTALLATION – CHASSIS DEVICES	2-22
RACK MOUNT INSTALLATION – STACKABLE DEVICES	2-23
POWERING ON A SYSTEM	2-24
CONNECTING NETWORK DEVICES	2-25
CONNECTORS	2-25
CABLE LENGTH	2-25
CONNECTING TO ETHERNET OR FAST ETHERNET HUBS	2-26
CONNECTING TO WORKSTATIONS, SERVERS, OR ROUTERS	2-27
INSTALLING OR REMOVING A GBIC	2-27
TROUBLESHOOTING NETWORK CONNECTIONS	2-28
TESTING CONNECTIVITY	2-29
PINGING AN IP ADDRESS	2-29
TRACING A ROUTE	2-29
MANAGING THE DEVICE	2-29
LOGGING ON THROUGH THE CLI	2-30
SEARCHING AND FILTERING OUTPUT FROM CLI COMMANDS	2-31
LOGGING ON THROUGH THE WEB MANAGEMENT INTERFACE	2-36
LOGGING ON THROUGH IRONVIEW NETWORK MANAGER NETWORK MANAGER	2-39
CREATING AN ALIAS FOR A CLI COMMAND	2-39
SWAPPING MODULES (CHASSIS DEVICES ONLY)	2-40
REMOVING THE OLD MODULE	2-40
ADDING THE NEW MODULE	2-41

CHAPTER 3

USING REDUNDANT MANAGEMENT MODULES	3-1
CONFIGURATION CONSIDERATIONS	3-1
TEMPERATURE SENSOR	3-2
SWITCHOVER	3-2
MANAGEMENT SESSIONS	3-2
SYSLOG AND SNMP TRAPS	3-2
MAC ADDRESS CHANGES	3-3
CONFIGURING THE REDUNDANT MANAGEMENT PARAMETERS	3-3
INSTALLING REDUNDANT MANAGEMENT MODULES	3-3
DETERMINING REDUNDANT MANAGEMENT MODULE STATUS	3-8
DISPLAYING SWITCHOVER MESSAGES	3-10
FILE SYNCHRONIZATION BETWEEN THE ACTIVE AND STANDBY REDUNDANT MANAGEMENT MODULES	3-11
SWITCHING OVER TO THE STANDBY REDUNDANT MANAGEMENT MODULE	3-16
PCMCIA FLASH CARD FILE MANAGEMENT COMMANDS	3-17
PCMCIA SLOTS	3-18
SUBDIRECTORIES	3-19
FILE AND SUBDIRECTORY NAMING CONVENTIONS	3-19
WILDCARDS	3-20
FORMATTING A FLASH CARD	3-20
DETERMINING THE FLASH CARD SLOT AND SUBDIRECTORY PATH THAT CURRENTLY HAVE THE MANAGEMENT FOCUS	3-21
SWITCHING THE MANAGEMENT FOCUS	3-21
DISPLAYING A DIRECTORY OF THE FILES ON A FLASH CARD	3-22
DISPLAYING THE CONTENTS OF A FILE	3-23
DISPLAY A HEXADECIMAL DUMP OF THE DATA IN A FILE	3-24
CREATING A SUBDIRECTORY	3-24
REMOVING A SUBDIRECTORY	3-25
RENAMING A FILE	3-26
CHANGING THE READ-WRITE ATTRIBUTE OF A FILE	3-26
DELETING A FILE FROM A FLASH CARD	3-26
RECOVERING (“UNDELETING”) A FILE	3-27
APPENDING A FILE TO ANOTHER FILE	3-27
COPYING FILES	3-28
LOADING THE SOFTWARE FROM A PCMCIA FLASH CARD	3-31
SAVING CONFIGURATION CHANGES TO A PCMCIA FLASH CARD	3-32
FILE MANAGEMENT MESSAGES	3-33
USING A 3COM MANAGEMENT INTERFACE IN THE PCMCIA SLOT	3-33
USAGE NOTES	3-34
INSTALLING THE PC CARD	3-34
REMOVING THE PC CARD	3-34
CHANGING PARAMETER SETTINGS	3-34

CHAPTER 4	
USING THE VELOCITY MANAGEMENT MODULE	4-1
OVERVIEW	4-1
MANAGEMENT AND CO-PROCESSING CPUS	4-1
TEMPERATURE SENSOR	4-2
MANAGEMENT REDUNDANCY	4-2
VSP LOAD SHARING	4-3
CHANGING THE MANAGEMENT SESSION FROM THE MP TO A VSP	4-8
LOGGING IN TO A VSP	4-8
LOGGING OUT FROM THE VSP	4-8
VSP COMMANDS	4-8
DISPLAYING VM1 MODULE INFORMATION	4-9
DISPLAYING THE SOFTWARE VERSION RUNNING ON THE MODULE	4-9
DISPLAYING THE SOFTWARE VERSIONS INSTALLED ON THE MODULE	4-10
DISPLAYING GENERAL MODULE INFORMATION	4-11
DETERMINING MODULE STATUS	4-12
DETERMINING THE SLOT ALLOCATIONS FOR THE VSPs	4-14
CHAPTER 5	
USING 10 GIGABIT ETHERNET MODULES	5-1
1-POR 10 GIGABIT ETHERNET MODULE	5-1
SYSTEM REQUIREMENTS	5-1
HARDWARE ON THE 1-POR 10 GIGABIT ETHERNET MODULE	5-2
FEATURES NOT SUPPORTED ON THE 1-POR 10 GIGABIT ETHERNET MODULE	5-2
REPLACING THE OPTICS ON THE 1-POR 10 GIGABIT ETHERNET MODULE	5-2
10 GIGABIT ETHERNET MODULES WITH XENPAK OPTICS	5-2
SYSTEM REQUIREMENTS	5-3
HARDWARE ON THE XENPAK-BASED 10 GIGABIT ETHERNET MODULE	5-3
FEATURES NOT SUPPORTED ON XENPAK-BASED 10 GIGABIT ETHERNET MODULES	5-3
REMOVING AND INSTALLING XENPAK OPTICS	5-3
CLEANING THE FIBER OPTIC CONNECTORS	5-4
CABLING 10 GIGABIT ETHERNET MODULES	5-4
PORT LEDs	5-5
TROUBLESHOOTING NETWORK CONNECTIONS	5-5
LINK FAULT SIGNALING (LFS)	5-6
DETERMINING THE 10 GIGABIT ETHERNET MODULE INSTALLED IN YOUR SYSTEM	5-6
CONFIGURING LINK FAULT SIGNALLING	5-6
REMOTE FAULT NOTIFICATION (RFN) ON FIBER CONNECTIONS	5-7
CONFIGURATION NOTES	5-7
RFN ENHANCEMENTS IN 07.8.00	5-7
ENABLING REMOTE FAULT NOTIFICATION	5-8
VIEWING WHICH FIBER PORTS HAVE RFN ENABLED	5-9
UPGRADING AN FPGA ON A 10 GIGABIT ETHERNET MODULE	5-9
DISPLAYING THE INSTALLED FPGA REVISIONS	5-10

CHAPTER 6**USING POWER OVER ETHERNET MODULES 6-1**

POWER OVER ETHERNET HARDWARE DESCRIPTION	6-1
J-B24E-POE AND J-F24E-POE MODULES	6-1
RPS-POE POWER SHELF	6-3
INSTALLING THE POE HARDWARE	6-4
INSTALLING A J-B24E-POE OR J-F24E-POE MODULE IN A FOUNDRY CHASSIS DEVICE	6-4
INSTALLING THE RPS-POE SHELF	6-5
CONFIGURING THE POE SOFTWARE	6-11
ENABLING POWER OVER ETHERNET	6-11
SPECIFYING THE POWER LIMIT FOR A PORT	6-12
ASSIGNING PRIORITY TO POE PORTS	6-12
DISPLAYING POWER OVER ETHERNET INFORMATION	6-13

CHAPTER 7**USING PACKET OVER SONET MODULES 7-1**

OVERVIEW	7-1
CABLE SPECIFICATIONS	7-2
NETWORK PROCESSOR ARCHITECTURE POS MODULES	7-3
SYSTEM REQUIREMENTS	7-3
INSTALLING A POS MODULE	7-3
UPGRADING POS SOFTWARE FROM A TFTP SERVER	7-4
UPGRADING THE BOOT CODE	7-4
UPGRADING THE FLASH CODE	7-5
INTERACTIVELY UPGRADING THE POS FLASH CODE	7-6
CONFIGURING POS BOOT PARAMETERS	7-7
CHANGING THE BOOT SOURCE	7-7
COPYING A POS IMAGE FILE FROM A FLASH CARD TO A POS MODULE'S FLASH MEMORY	7-7
REBOOTING	7-7
RELOADING THE SOFTWARE ON ALL MODULES	7-8
RELOADING THE SOFTWARE ON AN INDIVIDUAL MODULE	7-8
RELOADING THE SOFTWARE FROM TFTP	7-8
CONFIGURING POS INTERFACES	7-8
ADDING AN IP ADDRESS	7-9
CHANGING THE INTERFACE STATE	7-9
CHANGING THE ENCAPSULATION TYPE	7-10
CHANGING THE CLOCK SOURCE	7-10
CHANGING THE LOOPBACK PATH	7-10
CHANGING THE MTU	7-11
CHANGING THE CRC LENGTH	7-11
DISABLING OR REENABLING KEEPALIVE MESSAGES	7-11
CHANGING THE BANDWIDTH	7-12
CHANGING THE POS FLAGS	7-12
CHANGING THE FRAME TYPE	7-13
ENABLING OR DISABLING ATM SCRAMBLING	7-13

CONFIGURING A POS INTERFACE USING THE WEB MANAGEMENT INTERFACE	7-14
CONFIGURING POS FOR FRAME RELAY	7-15
CHANGING THE ENCAPSULATION TYPE	7-15
SPECIFYING THE FRAME RELAY INTERFACE TYPE	7-15
SPECIFYING THE DLCI	7-15
SPECIFYING THE LMI TYPE	7-15
VERIFYING THE CONFIGURATION	7-16
CONFIGURING POS FOR LAYER 2 SWITCHING	7-16
LINK REDUNDANCY AND LOAD BALANCING	7-18
CONFIGURATION PROCEDURES	7-21
CONFIGURING A POS PORT FOR LAYER 2 SWITCHING	7-21
CONFIGURING STP PARAMETERS	7-22
CONFIGURING THE POS PORTS INTO A TRUNK GROUP	7-24
CONFIGURING AUTOMATIC PROTECTION SWITCHING (APS) FOR LAYER 2 POS	7-25
DISPLAYING LAYER 2 POS PORT INFORMATION	7-26
DISPLAYING POS INFORMATION	7-26
DISPLAYING THE SOFTWARE VERSION RUNNING ON THE MODULE	7-26
DISPLAYING THE SOFTWARE INSTALLED IN FLASH MEMORY	7-28
DISPLAYING GENERAL MODULE INFORMATION	7-28
DETERMINING POS MODULE STATUS	7-28
DISPLAYING INTERFACE PARAMETERS	7-30
DISPLAYING POS STATISTICS	7-33
DISPLAYING POS ALARMS AND ERROR CONDITIONS	7-34
CONFIGURING AUTOMATIC PROTECTION SWITCHING (APS)	7-35
BASIC POS APS CONFIGURATION	7-35
MULTI-GROUP APS CONFIGURATION	7-36
SINGLE-DEVICE APS CONFIGURATION	7-37
CONFIGURING OPTIONAL PARAMETERS	7-40
DISPLAYING POS APS INFORMATION	7-42
PATH TRACE (J1 BYTE) FIELD SUPPORT	7-42
XENPAK WAN PHY TRANSCEIVER	7-43
SETTING THE WAN PHY XENPAK TO WAN MODE	7-43
SETTING WAN PHY LINK FAULT SIGNALING	7-44

CHAPTER 8

USING ASYNCHRONOUS TRANSFER MODE MODULES..... 8-1

OVERVIEW	8-1
FOUNDRY ATM INTERFACE SPECIFICATIONS	8-2
VIRTUAL CHANNEL (VC) SUPPORT	8-2
CLASS OF SERVICE (CoS) SUPPORT	8-4
ACCESS CONTROL LIST SUPPORT	8-5
INSTALLING AN ATM MODULE	8-5
CONFIGURING THE CHASSIS TO RECEIVE THE MODULE	8-5
UPGRADING ATM SOFTWARE FROM A TFTP SERVER	8-5
UPGRADING THE BOOT CODE	8-6
UPGRADING THE FLASH CODE	8-6

CONFIGURING ATM BOOT PARAMETERS	8-6
CHANGING THE BOOT SOURCE	8-6
BOOTING THE MODULE FROM TFTP	8-7
RELOADING AN INDIVIDUAL ATM MODULE	8-7
COPYING AN ATM IMAGE FILE FROM A FLASH CARD TO AN ATM MODULE'S FLASH MEMORY	8-7
CONFIGURING ATM INTERFACES	8-8
CONFIGURING PORT PARAMETERS	8-10
CONFIGURING SUB-INTERFACE PARAMETERS	8-15
CONFIGURING ROUTE PARAMETERS	8-20
VERIFYING PVC CONFIGURATION	8-23
MAPPING PVCs TO VLANs	8-24
ADDING AN ATM SUB-INTERFACE TO A VLAN	8-27
REMOVING AN ATM SUB-INTERFACE FROM A VLAN	8-27
ENABLING SPANNING TREE	8-28
DISPLAYING ATM INFORMATION	8-28
DISPLAYING GENERAL MODULE INFORMATION	8-28
DETERMINING ATM MODULE STATUS	8-29
DISPLAYING INTERFACE PARAMETERS	8-30
DISPLAYING ATM INFORMATION FOR INDIVIDUAL SLOTS OR PORTS	8-32
DISPLAYING PORT STATISTICS	8-32
DISPLAYING VC STATISTICS	8-33

CHAPTER 9

CONFIGURING BASIC FEATURES..... 9-1

USING THE WEB MANAGEMENT INTERFACE FOR BASIC CONFIGURATION CHANGES	9-2
CONFIGURING BASIC SYSTEM PARAMETERS	9-3
ENTERING SYSTEM ADMINISTRATION INFORMATION	9-4
CONFIGURING SIMPLE NETWORK MANAGEMENT (SNMP) PARAMETERS	9-5
CONFIGURING AN INTERFACE AS THE SOURCE FOR ALL TELNET PACKETS	9-11
CANCELLING AN OUTBOUND TELNET SESSION	9-12
CONFIGURING AN INTERFACE AS THE SOURCE FOR ALL TFTP PACKETS	9-12
SPECIFYING A SIMPLE NETWORK TIME PROTOCOL (SNTP) SERVER	9-12
SETTING THE SYSTEM CLOCK	9-14
CHANGING THE DEFAULT GIGABIT NEGOTIATION MODE	9-16
LIMITING BROADCAST, MULTICAST, OR UNKNOWN-UNICAST RATES	9-18
CONFIGURING CLI BANNERS	9-20
CONFIGURING TERMINAL DISPLAY	9-21
CHECKING THE LENGTH OF TERMINAL DISPLAYS	9-21
CONFIGURING BASIC PORT PARAMETERS	9-21
ASSIGNING A PORT NAME	9-24
MODIFYING PORT SPEED	9-25
MODIFYING PORT MODE	9-26
DISABLING OR RE-ENABLING A PORT	9-26
DISABLING OR RE-ENABLING FLOW CONTROL	9-27
CHANGING THE 802.3X GIGABIT NEGOTIATION MODE	9-28
MODIFYING PORT PRIORITY (QoS)	9-30

CONFIGURING BASIC LAYER 2 PARAMETERS	9-30
ENABLING OR DISABLING THE SPANNING TREE PROTOCOL (STP)	9-30
ENABLING OR DISABLING LAYER 2 SWITCHING (LAYER 3 SWITCHES ONLY)	9-33
CHANGING THE MAC AGE TIME	9-35
CONFIGURING STATIC MAC ENTRIES	9-35
CONFIGURING A LOCAL MAC ADDRESS FOR LAYER 2 MANAGEMENT TRAFFIC	9-37
ENABLING PORT-BASED VLANS	9-38
DEFINING MAC ADDRESS FILTERS	9-39
DEFINING BROADCAST AND MULTICAST FILTERS	9-45
LOCKING A PORT TO RESTRICT ADDRESSES	9-47
ENABLING OR DISABLING ROUTING PROTOCOLS	9-47
DISPLAYING AND MODIFYING SYSTEM PARAMETER DEFAULT SETTINGS	9-48
USING THE TEMPERATURE SENSOR	9-52
DISPLAYING THE TEMPERATURE	9-52
DISPLAYING TEMPERATURE MESSAGES	9-53
CHANGING TEMPERATURE WARNING AND SHUTDOWN LEVELS	9-53
CHANGING THE CHASSIS POLLING INTERVAL	9-54
ASSIGNING A MIRROR PORT AND MONITOR PORTS	9-55
CONFIGURATION GUIDELINES FOR MONITORING INBOUND TRAFFIC	9-55
CONFIGURING PORT MIRRORING AND MONITORING ON NON-TERATHON DEVICES	9-57
CONFIGURING PORT MIRRORING AND MONITORING ON TERATHON DEVICES	9-58
MONITORING AN INDIVIDUAL TRUNK PORT	9-59
MONITORING 802.3AD AGGREGATE LINKS	9-60
MIRROR PORTS FOR POLICY-BASED ROUTING (PBR) TRAFFIC	9-61
DISPLAYING THE CURRENT MIRROR AND MONITOR PORT CONFIGURATION	
ON NON-TERATHON DEVICES	9-62
DISPLAYING MIRROR AND MONITOR PORT CONFIGURATION ON TERATHON DEVICES	9-62

CHAPTER 10

CONFIGURING SPANNING TREE PROTOCOL (STP)

AND IRONSPAN FEATURES 10-1

CONFIGURING STANDARD STP PARAMETERS	10-1
STP PARAMETERS AND DEFAULTS	10-2
ENABLING OR DISABLING THE SPANNING TREE PROTOCOL (STP)	10-3
CHANGING STP BRIDGE AND PORT PARAMETERS	10-5
DISPLAYING STP INFORMATION	10-8
CONFIGURING IRONSPAN FEATURES	10-19
FAST PORT SPAN	10-19
FAST UPLINK SPAN	10-21
802.1W RAPID SPANNING TREE (RSTP)	10-22
802.1W DRAFT 3	10-58
SINGLE SPANNING TREE (SSTP)	10-62
SUPERSPAN™	10-65
STP PER VLAN GROUP	10-71
PVST/PVST+ COMPATIBILITY	10-75
PVST/PVST+ COMPATIBILITY – 07.6.01 AND LATER	10-75

PVST/PVST+ COMPATIBILITY – EARLIER THAN 07.6.01 10-81

CHAPTER 11

MULTIPLE SPANNING TREE PROTOCOL

ON ENTERPRISE RELEASES 11-1

CONFIGURING MSTP 11-2

 ENABLING MSTP ON A SWITCH 11-3

 SETTING THE MSTP REVISION NUMBER AND NAME 11-3

 CONFIGURING AN MSTP INSTANCE 11-3

 DISABLING MSTP ON A PORT 11-3

 CONFIGURING PORT PRIORITY AND PATH COST FOR AN MSTP INSTANCE 11-4

 CONFIGURING BRIDGE PRIORITY FOR AN MSTP INSTANCE 11-4

 SETTING THE MSTP GLOBAL PARAMETERS 11-4

 SETTING PORTS TO BE OPERATIONAL EDGE PORTS 11-5

 SETTING POINT-TO-POINT LINK 11-5

 COMMITTING MSTP CONFIGURATION ID CHANGES 11-5

 FORCING PORTS TO TRANSMIT THE CONFIGURED BPDU VERSION 11-5

 SAMPLE CONFIGURATION 11-5

DISPLAYING MSTP INFORMATION 11-8

 DISPLAYING MSTP CONFIGURATION 11-8

 DISPLAYING MSTP INFORMATION FOR A SPECIFIED INSTANCE O 11-9

 DISPLAYING MSTP INFORMATION FOR CIST INSTANCE 0 11-11

CHAPTER 12

MULTIPLE SPANNING TREE PROTOCOL

ON BIGIRON MG8 AND NETIRON 40G 12-1

CONFIGURING MSTP 12-2

 ENABLING MSTP ON A SWITCH 12-3

 SETTING THE MSTP NAME 12-3

 SETTING THE MSTP REVISION NUMBER 12-3

 CONFIGURING AN MSTP INSTANCE 12-3

 DISABLING MSTP ON A PORT 12-3

 CONFIGURING PORT PRIORITY AND PORT PATH COST 12-4

 CONFIGURING BRIDGE PRIORITY FOR AN MSTP INSTANCE 12-4

 SETTING THE MSTP GLOBAL PARAMETERS 12-4

 SETTING PORTS TO BE OPERATIONAL EDGE PORTS 12-5

 SETTING POINT-TO-POINT LINK 12-5

 FORCING PORTS TO TRANSMIT AN MSTP BPDU 12-5

CONFIGURATION EXAMPLE 12-5

DISPLAYING MSTP INFORMATION 12-8

 DISPLAYING MSTP STATISTICS 12-8

 DISPLAYING MSTP INFORMATION FOR A SPECIFIED INSTANCE 12-10

 DISPLAYING MSTP INFORMATION FOR CIST INSTANCE 0 12-11

 DISPLAYING MSTP CONFIGURATION 12-11

CHAPTER 13
CONFIGURING TRUNK GROUPS AND DYNAMIC LINK AGGREGATION 13-1

CONFIGURING TRUNK GROUPS 13-1

 TRUNK GROUP CONNECTIVITY TO A SERVER 13-2

 TRUNK GROUP RULES 13-3

 SPECIFYING A MINIMUM NUMBER OF PORTS FOR A TRUNK GROUP 13-8

 BIGIRON MG8 AND NETIRON 40G TRUNK FORMATION RULES (RELEASE 02.2.01 AND LATER) 13-9

 BIGIRON MG8 AND NETIRON 40G TRUNK FORMATION RULES (PRIOR TO RELEASE 02.2.01) 13-9

 OTHER RULES FOR FORMING A TRUNK ON A BIGIRON MG8 OR NETIRON 40G 13-14

 TRUNK GROUP LOAD SHARING 13-15

 CONFIGURING A TRUNK GROUP 13-20

 ADDITIONAL TRUNKING OPTIONS 13-25

 SERVER TRUNK GROUP LOAD SHARING ENHANCEMENTS AND OPTIONS
 (RELEASE 07.7.00 AND HIGHER) 13-29

 ENABLING OPTIMIZED SERVER TRUNK LOAD BALANCING (VM1 ONLY) 13-31

 DISPLAYING TRUNK GROUP CONFIGURATION INFORMATION 13-34

DYNAMIC LINK AGGREGATION 13-38

 USAGE NOTES 13-38

 CONFIGURATION RULES 13-38

 802.3AD ENHANCEMENTS IN RELEASE 07.6.01 13-40

 ENABLING LINK AGGREGATION 13-43

 LINK AGGREGATION PARAMETERS 13-44

 LINK AGGREGATION ON TAGGED PORTS FOR THE FES 13-49

 DISPLAYING AND DETERMINING THE STATUS OF AGGREGATE LINKS 13-50

 CLEARING THE NEGOTIATED LINK AGGREGATIONS 13-54

CHAPTER 14
CONFIGURING UNI-DIRECTIONAL LINK DETECTION (UDLD) 14-1

CONFIGURATION CONSIDERATIONS 14-2

CONFIGURING UDLD 14-2

 CHANGING THE KEEPALIVE INTERVAL 14-2

 CHANGING THE KEEPALIVE RETRIES 14-2

 UDLD FOR TAGGED PORTS 14-2

DISPLAYING UDLD INFORMATION 14-3

 DISPLAYING INFORMATION FOR ALL PORTS 14-3

 DISPLAYING INFORMATION FOR A SINGLE PORT 14-4

CLEARING UDLD STATISTICS 14-5

CHAPTER 15
CONFIGURING METRO FEATURES 15-1

TOPOLOGY GROUPS 15-1

 MASTER VLAN AND MEMBER VLANS 15-2

 CONTROL PORTS AND FREE PORTS 15-2

 CONFIGURATION CONSIDERATIONS 15-2

CONFIGURING A TOPOLOGY GROUP	15-3
DISPLAYING TOPOLOGY GROUP INFORMATION	15-3
METRO RING PROTOCOL (MRP)	15-5
MRP RINGS WITHOUT SHARED INTERFACES (MRP PHASE 1)	15-6
MRP RINGS WITH SHARED INTERFACES (MRP PHASE 2)	15-7
RING INITIALIZATION	15-9
HOW RING BREAKS ARE DETECTED AND HEALED	15-11
MASTER VLANS AND CUSTOMER VLANS	15-13
CONFIGURING MRP	15-15
USING MRP DIAGNOSTICS	15-16
DISPLAYING MRP INFORMATION	15-18
MRP CLI EXAMPLE	15-20
MRP AND VPLS INTERNETWORKING	15-22
VIRTUAL SWITCH REDUNDANCY PROTOCOL (VSRP)	15-25
LAYER 2 AND LAYER 3 REDUNDANCY	15-27
MASTER ELECTION AND FAILOVER	15-27
VSRP-AWARE SECURITY FEATURES	15-31
VSRP PARAMETERS	15-31
CONFIGURING BASIC VSRP PARAMETERS USING THE CLI	15-34
CONFIGURING OPTIONAL VSRP PARAMETERS USING THE CLI	15-35
DISPLAYING VSRP INFORMATION USING THE CLI	15-42
CONFIGURING VSRP USING THE WEB MANAGEMENT INTERFACE	15-46
DISPLAYING VSRP STATISTICS USING THE WEB MANAGEMENT INTERFACE	15-49
VSRP FAST START	15-51
VSRP AND MRP SIGNALING	15-52

CHAPTER 16

CONFIGURING VIRTUAL LANs (VLANs)..... 16-1

OVERVIEW	16-1
TYPES OF VLANs	16-1
PROTOCOL-BASED VLANs (BIGIRON MG8 AND NETIRON 40G SOFTWARE RELEASE 02.0.00 AND LATER) ...	16-5
DEFAULT VLAN	16-6
802.1Q TAGGING	16-7
SPANNING TREE PROTOCOL (STP)	16-9
VIRTUAL ROUTING INTERFACES	16-10
VLAN AND VIRTUAL ROUTING INTERFACE GROUPS	16-11
DYNAMIC, STATIC, AND EXCLUDED PORT MEMBERSHIP	16-11
SUPER AGGREGATED VLANs	16-14
TRUNK GROUP PORTS AND VLAN MEMBERSHIP	16-14
SUMMARY OF VLAN CONFIGURATION RULES	16-14
ROUTING BETWEEN VLANs (LAYER 3 SWITCHES ONLY)	16-15
VIRTUAL ROUTING INTERFACES (LAYER 3 SWITCHES ONLY)	16-15
BRIDGING AND ROUTING THE SAME PROTOCOL SIMULTANEOUSLY	
ON THE SAME DEVICE (LAYER 3 SWITCHES ONLY)	16-15
ROUTING BETWEEN VLANs USING VIRTUAL ROUTING INTERFACES (LAYER 3 SWITCHES ONLY)	16-15

DYNAMIC PORT ASSIGNMENT (LAYER 2 SWITCHES AND LAYER 3 SWITCHES)	16-16
ASSIGNING A DIFFERENT VLAN ID TO THE DEFAULT VLAN	16-16
ASSIGNING TRUNK GROUP PORTS	16-16
CONFIGURING PORT-BASED VLANs	16-16
MODIFYING A PORT-BASED VLAN	16-20
CONFIGURING IP SUBNET, IPX NETWORK AND PROTOCOL-BASED VLANs	16-23
CONFIGURATION NOTES FOR FES DEVICES	16-23
CONFIGURATION CONSIDERATIONS FOR BIGIRON MG8 AND NETIRON 40G	16-23
CONFIGURATION EXAMPLE	16-23
CONFIGURING IP SUBNET, IPX NETWORK, AND	
PROTOCOL-BASED VLANs WITHIN PORT-BASED VLANs	16-25
CONFIGURING AN IPV6 PROTOCOL VLAN	16-28
ROUTING BETWEEN VLANs USING VIRTUAL ROUTING INTERFACES (LAYER 3 SWITCHES ONLY)	16-29
CONFIGURING APPLE TALK CABLE VLANs	16-35
CONFIGURATION GUIDELINES	16-35
CONFIGURATION EXAMPLE	16-35
CONFIGURING PROTOCOL VLANs WITH DYNAMIC PORTS	16-38
AGING OF DYNAMIC PORTS	16-38
CONFIGURATION GUIDELINES	16-38
CONFIGURING AN IP, IPX, OR APPLE TALK PROTOCOL VLAN WITH DYNAMIC PORTS	16-38
CONFIGURING AN IP SUBNET VLAN WITH DYNAMIC PORTS	16-39
CONFIGURING AN IPX NETWORK VLAN WITH DYNAMIC PORTS	16-40
CONFIGURING UPLINK PORTS WITHIN A PORT-BASED VLAN	16-40
CONFIGURING THE SAME IP SUBNET ADDRESS ON MULTIPLE PORT-BASED VLANs	16-41
USING SEPARATE ACLs ON IP FOLLOWER VIRTUAL ROUTING INTERFACES	16-44
CONFIGURING VLAN GROUPS AND VIRTUAL ROUTING INTERFACE GROUPS	16-45
CONFIGURING A VLAN GROUP	16-45
CONFIGURING A VIRTUAL ROUTING INTERFACE GROUP	16-46
DISPLAYING THE VLAN GROUP AND VIRTUAL ROUTING INTERFACE GROUP INFORMATION	16-47
ALLOCATING MEMORY FOR MORE VLANs OR VIRTUAL ROUTING INTERFACES	16-48
CONFIGURING SUPER AGGREGATED VLANs	16-50
CONFIGURATION NOTE FOR FES DEVICES	16-52
CONFIGURATION NOTE FOR BIGIRON MG8 AND NETIRON 40G	16-52
CONFIGURING AGGREGATED VLANs	16-52
APPLYING QoS MARKINGS TO SAV TRAFFIC IN THE NETWORK CORE	16-54
COMPLETE CLI EXAMPLES	16-54
CONFIGURING 802.1Q-IN-Q TAGGING	16-57
CONFIGURATION RULES	16-58
ENABLING 802.1Q-IN-Q TAGGING	16-59
EXAMPLE CONFIGURATION	16-59
CONFIGURING 802.1Q TAG-TYPE TRANSLATION	16-60
CONFIGURATION RULES	16-62
ENABLING 802.1Q TAG-TYPE TRANSLATION	16-63
CONFIGURING PRIVATE VLANs	16-64
IMPLEMENTATION NOTES	16-65
CONFIGURATION NOTES FOR FES DEVICES	16-65

CONFIGURING A PRIVATE VLAN	16-66
ENABLING BROADCAST OR UNKNOWN UNICAST TRAFFIC TO THE PRIVATE VLAN	16-68
CLI EXAMPLE FOR FIGURE 22.24	16-68
DUAL-MODE VLAN PORTS	16-69
JETCORE MODULE HARDWARE FLOODING FOR LAYER 2 MULTICAST AND BROADCAST PACKETS	16-72
JETCORE BROADCAST SUPPRESSION	16-73
UNICAST FLOODING ON VLAN PORTS	
(BIGIRON MG8 AND NETIRON 40G SOFTWARE RELEASE 02.0.00 AND LATER)	16-73
CONFIGURATION CONSIDERATIONS	16-74
CONFIGURING VLAN TRANSLATION (BIGIRON MG8 AND NETIRON 40G SOFTWARE RELEASE 02.0.00 AND LATER)	
16-74	
CONFIGURATION CONSIDERATIONS	16-75
CLI COMMAND FOR VLAN TRANSLATION	16-75
CONFIGURATION EXAMPLE	16-75
CONFIGURING INNER VLAN TRANSLATION WITH SUPER AGGREGATED VLANs (BIGIRON MG8 AND NETIRON 40G	
SOFTWARE RELEASE 02.0.00 AND LATER)	16-76
CONFIGURATION CONSIDERATIONS	16-76
CLI COMMAND TO CONFIGURE AN INTERFACE FOR VLAN TRANSLATION ON A	
SUPER AGGREGATED VLAN	16-77
CONFIGURATION EXAMPLE	16-77
CONFIGURING MAC VLANs (STACKABLE FASTIRON BACKBONE LAYER 2 SWITCH ONLY)	16-78
CONFIGURING A MAC VLAN LIST	16-79
LOADING A MAC VLAN LIST	16-80
SPECIFYING A DEFAULT VLAN FOR MAC ADDRESSES THAT ARE NOT IN THE MAC VLAN LIST	16-80
CLEARING MAC VLAN ENTRIES FROM THE MAC TABLE	16-81
CONFIGURING VLAN BYTE ACCOUNTING	16-81
CONFIGURING VLANs USING THE WEB MANAGEMENT INTERFACE	16-83
CONFIGURING A PORT-BASED VLAN	16-83
CONFIGURING A PROTOCOL-BASED VLAN	16-85
CONFIGURING AN IP SUBNET VLAN	16-86
CONFIGURING AN IPX NETWORK VLAN	16-87
CONFIGURING AN APPLE TALK CABLE VLAN	16-89
DISPLAYING VLAN INFORMATION	16-90
DISPLAYING SYSTEM-WIDE VLAN INFORMATION	16-90
DISPLAYING VLAN INFORMATION FOR SPECIFIC PORTS	16-91
USING SHOW RUN TO DISPLAY VLAN INFORMATION	16-92

CHAPTER 17

CONFIGURING IP MULTICAST TRAFFIC REDUCTION 17-1

ENABLING IP MULTICAST TRAFFIC REDUCTION	17-1
CHANGING THE IGMP MODE	17-2
DISABLING IGMP ON INDIVIDUAL PORTS	17-3
MODIFYING THE QUERY INTERVAL	17-4
MODIFYING THE AGE INTERVAL	17-4
FILTERING MULTICAST GROUPS	17-5
PIM SM TRAFFIC SNOOPING	17-5

APPLICATION EXAMPLES	17-6
CONFIGURATION REQUIREMENTS	17-7
ENABLING PIM SM TRAFFIC SNOOPING	17-8
DISPLAYING IP MULTICAST INFORMATION	17-8
DISPLAYING MULTICAST INFORMATION ON LAYER 2 SWITCHES	17-8
DISPLAYING MULTICAST INFORMATION ON LAYER 3 SWITCHES	17-17
DISPLAYING IP MULTICAST STATISTICS	17-20
CLEARING IP MULTICAST STATISTICS	17-21
CLEARING IGMP GROUP FLOWS	17-21

CHAPTER 18

CONFIGURING

GARP VLAN REGISTRATION PROTOCOL (GVRP) 18-1

APPLICATION EXAMPLES	18-1
DYNAMIC CORE AND FIXED EDGE	18-2
DYNAMIC CORE AND DYNAMIC EDGE	18-3
FIXED CORE AND DYNAMIC EDGE	18-4
FIXED CORE AND FIXED EDGE	18-4
VLAN NAMES	18-4
CONFIGURATION CONSIDERATIONS	18-4
CONFIGURING GVRP	18-5
CHANGING THE GVRP BASE VLAN ID	18-5
INCREASING THE MAXIMUM CONFIGURABLE VALUE OF THE LEAVEALL TIMER	18-6
ENABLING GVRP	18-6
DISABLING VLAN ADVERTISING	18-6
DISABLING VLAN LEARNING	18-7
CHANGING THE GVRP TIMERS	18-7
CONVERTING A VLAN CREATED BY GVRP INTO A STATICALLY-CONFIGURED VLAN	18-8
DISPLAYING GVRP INFORMATION	18-9
DISPLAYING GVRP CONFIGURATION INFORMATION	18-9
DISPLAYING GVRP VLAN INFORMATION	18-12
DISPLAYING GVRP STATISTICS	18-14
DISPLAYING CPU UTILIZATION STATISTICS	18-15
DISPLAYING GVRP DIAGNOSTIC INFORMATION	18-17
CLEARING GVRP STATISTICS	18-17
CLI EXAMPLES	18-17
DYNAMIC CORE AND FIXED EDGE	18-17
DYNAMIC CORE AND DYNAMIC EDGE	18-18
FIXED CORE AND DYNAMIC EDGE	18-19
FIXED CORE AND FIXED EDGE	18-19

CHAPTER 19

CONFIGURING BASE LAYER 3 19-1

ADDING A STATIC IP ROUTE	19-1
ADDING A STATIC ARP ENTRY	19-2

CONFIGURING RIP	19-2
ENABLING RIP	19-2
ENABLING REDISTRIBUTION OF IP STATIC ROUTES INTO RIP	19-3
ENABLING REDISTRIBUTION	19-4
ENABLING LEARNING OF DEFAULT ROUTES	19-4
CHANGING THE ROUTE LOOP PREVENTION METHOD	19-4
ADDITIONAL FEATURES	19-5

CHAPTER 20

ENABLING THE FOUNDRY DISCOVERY PROTOCOL (FDP) AND READING

CISCO DISCOVERY PROTOCOL (CDP) PACKETS..... 20-1

USING FDP	20-1
CONFIGURING FDP	20-1
DISPLAYING FDP INFORMATION	20-2
CLEARING FDP AND CDP INFORMATION	20-5
READING CDP PACKETS	20-5
ENABLING INTERCEPTION OF CDP PACKETS GLOBALLY	20-6
ENABLING INTERCEPTION OF CDP PACKETS ON AN INTERFACE	20-6
DISPLAYING CDP INFORMATION	20-6
CLEARING CDP INFORMATION	20-8

APPENDIX A

USING SYSLOG.....A-1

OVERVIEW	A-1
DISPLAYING SYSLOG MESSAGES	A-2
CONFIGURING THE SYSLOG SERVICE	A-3
DISPLAYING THE SYSLOG CONFIGURATION	A-3
DISPLAYING AND CONFIGURING SYSLOG BUFFER PARAMETERS USING THE WEB MANAGEMENT INTERFACE	A-7
DISABLING OR RE-ENABLING SYSLOG	A-9
SPECIFYING A SYSLOG SERVER	A-9
SPECIFYING AN ADDITIONAL SYSLOG SERVER	A-9
DISABLING LOGGING OF A MESSAGE LEVEL	A-10
CHANGING THE NUMBER OF ENTRIES THE LOCAL BUFFER CAN HOLD	A-10
CHANGING THE LOG FACILITY	A-11
DISPLAYING THE INTERFACE NAME IN SYSLOG MESSAGES	A-12
CLEARING THE SYSLOG MESSAGES FROM THE LOCAL BUFFER	A-12
DISPLAYING TCP/UDP PORT NUMBERS IN SYSLOG MESSAGES	A-12
SYSLOG MESSAGES	A-13
SYSLOG MESSAGES FOR WEB AUTHENTICATION	A-39

APPENDIX B

HARDWARE SPECIFICATIONS.....B-1

CHASSIS DEVICES	B-1
-----------------------	-----

BIGIRON	B-1
NETIRON ROUTER	B-3
FASTIRON II FAMILY	B-4
REDUNDANT MANAGEMENT MODULES	B-4
STACKABLE DEVICES	B-5
FASTIRON 4802	B-5
FASTIRON EDGE SWITCH	B-5
NETIRON LAYER 3 SWITCH	B-5
NETIRON 4802	B-5
CONTROL FEATURES	B-5
CONTROL PANELS	B-5
PORTS	B-8
LEDs	B-10
RESET BUTTON	B-14
POWER SPECIFICATIONS	B-14
POWER SPECIFICATIONS FOR CHASSIS DEVICES	B-14
POWER SPECIFICATIONS FOR STACKABLE DEVICES	B-18
PHYSICAL DIMENSIONS	B-19
OPERATING ENVIRONMENT	B-19
STORAGE ENVIRONMENT	B-19
ELECTROMAGNETIC EMISSIONS	B-19
SAFETY AGENCY APPROVALS	B-20
LASER	B-20

APPENDIX C

SOFTWARE SPECIFICATIONS	C-1
IEEE COMPLIANCE	C-1
RFC SUPPORT	C-2
ISO/IEC SPECIFICATIONS	C-5
INTERNET DRAFTS	C-5

APPENDIX D

JETCORE CHASSIS MODULES	D-1
DETERMINING YOUR DEVICE TYPE	D-1
JETCORE CHASSIS MODULES	D-1
THE JETCORE MANAGEMENT MODULE	D-3
HARDWARE OVERVIEW	D-3
J-B2GMR4 AND J-F2GMR4 JETCORE MANAGEMENT MODULES	D-4
J-BxGMR4 AND J-FIXGMR4 JETCORE MANAGEMENT MODULES	D-5
J-F2404GMR4 JETCORE MANAGEMENT MODULE	D-6
JETCORE GIGABIT ETHERNET FORWARDING MODULES	D-9
J-BxG AND J-FIXG 8-PORT FORWARDING MODULES	D-9
J-B16GC AND J-F16GC 16-PORT FORWARDING MODULES	D-9
JETCORE 10/100 ETHERNET FORWARDING MODULES	D-9
24-PORT 100BASEFX FORWARDING MODULE	D-9
J-B48E AND J-FI48E 48-PORT ENTERPRISE FORWARDING MODULES	D-10

J-B48E-A AND J-F48E-A 48-PORT FORWARDING MODULES D-11
J-B48T AND J-FI48T 48-PORT TELCO FORWARDING MODULES D-11
J-B48T-A AND J-F48T-A 48-PORT FORWARDING MODULES D-13
CONFIGURATION CONSIDERATIONS D-13

APPENDIX E

CAUTIONS AND WARNINGS..... E-1
CAUTIONS E-1
WARNINGS E-8

Chapter 1

Getting Started

This guide contains hardware installation and basic software configuration procedures. \

The General hardware Installation procedures discussed in this guide apply to the following devices:

- BigIron 4000/8000/15000 Chassis devices with IronCore or JetCore management modules
- FastIron II, FastIron II Plus, and FastIron III with M2 or higher management modules
- FastIron 400/800/1500 Chassis devices with JetCore modules
- FastIron 4802 Stackable device
- FastIron 4802 Stackable device
- NetIron 400/800/1500 Chassis devices with IronCore or JetCore management modules
- NetIron 4802 Stackable device

See the “List of Publications” on page 1-2 for documents containing hardware installation procedures for other Foundry devices.

The basic software configuration information in this guide applies to the following devices:

- Enterprise IronWare software releases, which apply to the following products:
 - NetIron 400/800/1500 Chassis devices with IronCore or JetCore management modules
 - BigIron 4000/8000/15000 Chassis devices with IronCore or JetCore management modules
 - FastIron II, FastIron II Plus, and FastIron III with M2 or higher management modules
 - FastIron 400/800/1500 Chassis devices with JetCore modules
 - FastIron 4802 Stackable device
- Service Provider IronWare software releases, which apply to the following products:
 - NetIron 400/800/1500 Chassis devices with IronCore or JetCore management modules
 - BigIron 4000/8000/15000 Chassis devices with IronCore or JetCore management modules
 - NetIron 4802 Stackable device
 - FastIron 4802 Stackable device

NOTE: You cannot use this software on FastIron Chassis devices.

- Terathon devices that include the following:

- BigIron MG8
- NetIron 40G
- NetIron IMR 640
- FastIron Edge Switch
- IronPoint-FastIron Edge Switch (IP-FES) Release 01.3.00 through 01.4.01
- ServerIron product family

Audience

This manual is designed for system administrators with a working knowledge of Layer 2 and Layer 3 switching and routing.

If you are using a Foundry Layer 3 Switch, you should be familiar with the following protocols if applicable to your network – IP, RIP, OSPF, IS-IS, BGP4, MBGP, MPLS, IGMP, PIM, DVMRP, IPX, AppleTalk, FSRP, VRRP, and VRRPE.

Nomenclature

This guide uses the following typographical conventions to show information:

Italic highlights the title of another publication and occasionally emphasizes a word or phrase.

Bold highlights a CLI command.

Bold Italic highlights a term that is being defined.

Underline highlights a link on the Web management interface.

Capitals highlights field names and buttons that appear in the Web management interface.

NOTE: A note emphasizes an important fact or calls your attention to a dependency.

WARNING: A warning calls your attention to a possible hazard that can cause injury or death.

CAUTION: A caution calls your attention to a possible hazard that can damage equipment.

List of Publications

The following Foundry Networks documents are available.

- *Foundry Switch and Router Installation and Basic Configuration Guide* – provides configuration guidelines for Layer 2 and Layer 3 devices and installation procedures for the Foundry devices with IronCore and JetCore modules.
- *Foundry Security Guide* – provides procedures for securing management access to Foundry devices and for protecting against Denial of Service (DoS) attacks.
- *Foundry Enterprise Configuration and Management Guide* – provides configuration information for enterprise routing protocols including IP, RIP, IP multicast, OSPF, BGP4, VRRP and VRRPE.
- *Foundry NetIron Service Provider Configuration and Management Guide* – provides configuration information for IS-IS and MPLS for Foundry devices that support IS-IS and MPLS, except for the NetIron IMR 640.
- *Foundry NetIron IMR 640 Service Provider Configuration and Management Guide* – provides configuration information for IS-IS and MPLS for for the NetIron IMR 640.
- *Foundry Switch and Router Command Line Interface Reference* – provides a list and syntax information for all the Layer 2 Switch and Layer 3 Switch CLI commands.

- *Foundry Diagnostic Guide* – provides descriptions of diagnostic commands that can help you diagnose and solve issues on Layer 2 Switches and Layer 3 Switches.
- *Foundry BigIron Mg8 Switch Installation and Basic Configuration Guide* – provides installation procedures for the BigIron MG8. This guide also presents the management modules available in the device.
- *Foundry NetIron 40G Switch Installation and Basic Configuration Guide* – provides installation procedures for the BigIron MG8. This guide also presents the management modules available in the device.
- *NetIron IMR 640 Installation and Basic Configuration Guide* – provides procedures for installing modules into and connecting your DC power source(s) to the NetIron IMR 640 chassis, cabling the Ethernet interface ports, and performing a basic configuration of the software.
- *Foundry Management Information Base Reference* – presents the Simple Network Management Protocol (SNMP) Management Information Base (MIB) objects that are supported in the Foundry devices.
- *Foundry IPv6 Configuration Guide* – provide configuration information for IPv6 features.
- *Foundry IronPoint Wireless LAN Configuration Guide* – presents the features for the IronPoint wireless LAN (WLAN).

To order additional copies of these manuals, do one of the following:

- Call 1.877.TURBOCALL (887.2622) in the United States or 1.408.586.1881 outside the United States.
- Send email to info@foundrynet.com.

What's New in this Edition

This edition describes the features in the following software releases:

- Enterprise IronWare software release 08.0.00.
- Service Provider IronWare software releases up to 09.2.00.
- Releases up 02.2.01 for the BigIron MG8 and NetIron 40G software.
- FastIron Edge Switch software release 03.4.01, 03.4.00, 03.3.01a, and 03.2.00
- FastIron Edge Switch X-Series release 02.2.00 and 02.1.01
- FastIron SuperX release 02.0.01, 02.1.00, 02.2.00, and 02.2.01
- IronPoint-FastIron Edge Switch (IP-FES) Release 01.3.00 through 01.4.01

New features in the Enterprise Release 08.0.00

The following tables present the new in the Enterprise software release 08.0.00. The J and I columns in the tables indicate the platforms on which each feature is supported. A “✓” in the J column indicates the feature is supported on JetCore devices. A “✓” in the I column indicates the feature is supported on IronCore devices.

Layer 2 Enhancements in 08.0.00

Enhancement	Description	J	I	See...
Dynamic configuration of Cisco Voice over IP (VoIP) phones via voice VLANs	This release enables the Foundry device to automatically re-configure a VoIP phone when the phone is physically moved from one port to another.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>

Enhancement	Description	J	I	See...
JetCore Layer 2 Access Control Lists (ACLs)	You can configure Layer 2 ACLs, which filter incoming traffic based on Layer 2 MAC header fields in the Ethernet/IEEE 802.3 frame, including source MAC address and mask, destination MAC address and mask, VLAN ID, and Ethernet type.	✓		<i>Foundry Enterprise Configuration and Management Guide</i>
Layer 2 ACL-based rate limiting	This feature is an extension to the existing IP ACL-based rate limiting on JetCore devices. Whereas the existing feature provides the facility to limit the rate for IP traffic that matches the permit conditions in standard or extended IP ACLs; the new feature enables you to limit traffic rates using the Layer 2 parameters defined in the associated JetCore Layer 2 ACL table.	✓		<i>Foundry Enterprise Configuration and Management Guide</i>
802.1s support	Multiple Spanning Tree Protocol (MSTP) as defined in IEEE 802.1s-2002 allows you to configure multiple STP instances.	✓	✓	"Multiple Spanning Tree Protocol on Enterprise Releases" on page 11-1

Layer 3 Enhancements in 08.0.00

Enhancement	Description	J	I	See...
Support for IPv6	This release supports IPv6 on JetCore devices running the B2P image only. IPv6 functionality works as documented in the <i>Foundry IPv6 Configuration Guide</i> , with the following exceptions: <ul style="list-style-type: none"> IPv6 PIM-SSM is not supported in release 08.0.00 IPv6 MLD is not supported in release 08.0.00 	✓		<i>Foundry IPv6 Configuration Guide</i>
Clearing OSPF information from the Foundry device	This release includes new CLI commands that allow you to clear specific kinds of information from the Foundry device's OSPF link state database and OSPF routing table. You do not need to remove statements from the Foundry device's configuration or reload the software for the commands to take effect.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>

Enhancement	Description	J	I	See...
OSPF redistribution filter rebinding	In previous releases, if you modified OSPF redistribution filters in the Foundry device's configuration, you then had to remove and reapply the redistribution rip and redistribution static statements in order for OSPF to start redistributing routes based on the filters. Starting in release 08.0.00, a new CLI command rebinds the redistribution filters, so you no longer have to remove and reapply the redistribution rip and redistribution static statements.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Configuring an OSPF non-broadcast interface	You can configure an interface on a Foundry device to send OSPF traffic to its neighbor as unicast packets rather than broadcast packets.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Configuring a unicast route with multiple outgoing ports	In this release, the Foundry device can route incoming Layer 3 unicast IP packets to two or more statically defined outgoing Layer 3 interfaces. Previous releases support sending traffic received from a VLAN or a Layer 3 routed interface to multiple Layer 2 addresses.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Advertising an IBGP next hop as a null0 route as a defense against denial of service attacks	You can create a static route to forward traffic destined to a host or network targeted in a denial of service attack to the null0 interface, and then use IBGP to advertise the null0 route to other routers in the network, so that the other routers drop the traffic for the targeted host or network.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
IP load sharing (ECMP) for RIPv2 routes	In release 08.0.00, IP load sharing is supported for RIPv2 routes. IP load sharing is also known as "Equal-Cost Multi-Path (ECMP)" load sharing or just "ECMP".	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
VRRPE slow start timer	The VRRPE slow start timer causes a specified amount of time to elapse between the time a VRRPE Master router comes back up and when it takes over from a Backup router. This interval allows time for OSPF convergence when the Master is restored.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>

Multicast Enhancements in 08.0.00

Enhancement	Description	J	I	See...
IGMP V3 Snooping	Release 08.0.00 supports IGMP V3 snooping on devices running both switch and router code.	✓		<i>Foundry Enterprise Configuration and Management Guide</i>
Increased size of the MSDP Source Active cache	The size of the cache used to store MSDP Source Active messages has been increased from 4k to 8k.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Specifying a Designated Router election priority for PIM	You can assign a DR election priority to each PIM router in a multi-access network. The router with the highest DR election priority is elected the DR.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Multicast MIB support	Support for the following standard multicast MIBs has been added to this release: <ul style="list-style-type: none"> • RFC 2932: IPv4 Multicast Routing MIB • RFC 2933: IGMP MIB • RFC 2934: PIM for IPv4 	✓	✓	<i>Management Information Base Guide</i>

Security Enhancements in 08.0.00

Enhancement	Description	J	I	See ...
Configuration bootguard feature	The configuration bootguard feature causes the Foundry device to erase its startup-config file in the event that the device is booted using the default configuration, and a Super User password had been stored on the device previously. This can prevent an unauthorized user from booting the Foundry device using the default-config and then using the Foundry device to gain access to other devices in the network.	✓	✓	<i>Foundry Security Guide</i>
Using multi-device port authentication and 802.1X security on the same port	You can configure the Foundry device to use multi-device port authentication and 802.1X security on the same port.	✓	✓	<i>Foundry Security Guide</i>

Enhancement	Description	J	I	See ...
CPU ACL	Starting in release 08.0.00, you can configure the Foundry device to protect itself from denial of service attacks for all protocols and applications received on the device. To do so, extended ACLs applied to ingress ports filter on the protocols specified in ACL clauses. For example, you can use extended ACLs to filter on protocols such as SNMP, TFTP, BGP, and ICMP.	✓		<i>Foundry Security Guide</i>

System-Level Enhancements in 08.0.00

Enhancement	Description	J	I	See ...
Creating an alias for a CLI command	You can create aliases for CLI commands. An alias serves as a shorthand version of a longer CLI command; for example, shoro can be defined as an alias for the show ip route CLI command.	✓	✓	“Creating an Alias for a CLI Command” on page 2-39
Directing debugging output to multiple destinations	You can direct debugging output (output from debug commands) to multiple destinations. This allows debugging output to be displayed on multiple user sessions concurrently. In previous releases, debugging output could be directed only to a single destination (the Syslog buffer, or a specified Telnet or SSH session).	✓	✓	<i>Foundry Diagnostic Guide</i>
sFlow enhancements	Release 08.0.00 includes the following enhancements to Foundry’s support for sFlow: <ul style="list-style-type: none"> • Support for sFlow version 5 • Support for selecting which and how often data destined to the CPU (for example, Telnet sessions) is sent to the sFlow collector 	✓		<i>Foundry Enterprise Configuration and Management Guide</i>
Jumbo packet counter in show statistics output	The output of the show statistics command has been enhanced to display information about the number of packets received that were longer than 1518 octets.	✓		<i>Foundry Enterprise Configuration and Management Guide</i>
Rate limiting for IP interface traffic on JetCore devices	You can set a limit for the rate at which traffic directed to IP interfaces on the Foundry device is sent to the management CPU for processing. This can prevent the management CPU from being overwhelmed when a high rate of traffic is directed to IP interfaces.	✓		<i>Foundry Enterprise Configuration and Management Guide</i>
Displaying VLANs in numerical order	The output of the show run command displays the VLANs configured on the Foundry device in numerical order (for example, VLAN 100, VLAN 200, and so on).	✓	✓	“Using show run to Display VLAN Information” on page 22-92

Enhancement	Description	J	I	See ...
Specifying a host name in an ACL statement	You can specify a host name within the configuration of an access control list. Previous releases supported only IP addresses in ACL statements.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Multiple domain list support	You can create a list of domain names that can be used by the DNS Resolver to resolve host names with their IP addresses. You can also verify the host name for an IP address or the IP address for a host name.	✓	✓	<i>Foundry Enterprise Configuration and Management Guide</i>
Multicast MIB support	Support for the following standard multicast MIBs has been added to this release: <ul style="list-style-type: none"> • RFC 2932: IPv4 Multicast Routing MIB • RFC 2933: IGMP MIB • RFC 2934: PIM for IPv4 	✓	✓	<i>MIB Reference</i>

New Features in Service Provider 09.2.00

Layer 2 Enhancement in 09.2.00

Enhancement	Description	J	I	See...
MRP Phase 2	With MRP Phase 2, MRP rings can be configured to share the same interfaces as long as the interfaces belong to the same VLAN	✓	✓	"MRP Rings with Shared Interfaces (MRP Phase 2)" on page 20-7
MRP and VPLS internetworking	MRP protocol can now work with VPLS.	✓	✓	"VSRP and MRP Signaling" on page 20-52

System Level Enhancements in 09.2.00

Enhancement	Description	J	I	See...
Drastic CPU protection scheme enhancement	When the hardware flooding option is enabled, hardware flooding is used to protect the CPU during normal traffic. Once the CPU usage goes above a certain percentage, drastic CPU protection scheme can then implemented. Once the CPU usage is below a defined trigger point, hardware flooding CPU protection scheme is used.	✓		<i>Foundry Security Guide</i>
New trap message for port priority changes	A trap message is generated when a port's priority is changed.	✓	✓	<i>MIB Reference</i>

Enhancements in the FastIron Edge Switch Release 03.4.0x

Security Enhancements in FES 03.4.00

Enhancement	Description	See...
802.1X termination-action for RADIUS clients	Software release 03.4.00 and 802.1X support the <i>termination-action</i> attribute, which specifies the action to be taken by the the Foundry device after a service has completed.	<i>Foundry Security Guide</i>
802.1X session timeout for RADIUS clients	Software release 03.4.00 and 802.1X support the <i>session timeout</i> attribute for RADIUS clients.	<i>Foundry Security Guide</i>

Enhancements in FES 03.4.01

Enhancement	Description	See...
Support for IPv6 host features	<p>This release supports IPv6 host and management functions. These release notes describe the following:</p> <ul style="list-style-type: none"> IPv6 features and CLI commands supported in this release IPv6 Implementation differences between the Layer 2 and Layer 3 code IPv6 commands that are new and specific to this release <p>For information about configuring IPv6 features on a Foundry device that are not covered in these release notes, see the <i>Foundry IPv6 Configuration Guide</i>.</p>	<i>Foundry IPv6 Configuration Guide</i>
Enhanced AAA support for console commands	Release 03.4.01 enhances authentication, authorization, and accounting (AAA) support for CLI commands entered at the console.	<i>Foundry Security Guide</i>

Enhancements in the IronPoint-FastIron Edge Switch (IP-FES)

The following enhancements have been added to the IP-FES, the switch that supports wireless LAN (WLAN) features.

Enhancement	Description	See...
New device	The IP-FES has been introduced to support WLAN features that are configured on an edge switch. Either a WLAN switch can be purchased or an existing FastIron Edge Switch can be upgraded to an IP-FES. Requirements and procedures for the upgrade are in the <i>Foundry IronPoint Wireless LAN Configuration Guide</i> and in the instructions that come with the upgrade kit.	<i>Foundry IronPoint Wireless LAN Configuration Guide</i>
Automatic Port Deactivation	This feature disables a WLAN switch's interface if a link down condition is seen on the interface.	
Automatic Discovery and Configuration (ADC)	This feature allows you to assign an IP address, subnet mask, default gateway to an access point when the access point is connected to a WLAN switch's interface.	
Wireless Mobility	This feature allows wireless clients to roam from one access point to another without losing connection to the network.	
VPN Passthrough	This feature filters VPN traffic to prevent unwanted traffic from accessing the network through an access point.	
Web Authentication	Web Authentication authenticates a user's login credentials via a Web browser before the device can access the network.	

Enhancements in BigIron MG8 and NetIron 40G Release 02.2.01

Layer 2 Enhancements in Release 02.2.01

Enhancement	Description	See...
802.1s Multiple Spanning Tree Protocol (MSTP)	With this release, you can configure multiple STP instances using MSTP protocol, as defined in IEEE 802.1s-2002.	"Multiple Spanning Tree Protocol on BigIron MG8 and NetIron 40G" on page 14-1
New rule for creating trunk groups	In this release, trunks can be formed from any number of ports, as long as they contain at least 2 ports and no more than 8 ports.	"BigIron MG8 and NetIron 40G Trunk Formation Rules (Release 02.2.01 and Later)" on page 17-9
Enhancement to the lACP system-priority command	The lACP system-priority command has been moved from the interface configuration level to the global configuration level.	"Configuring Link Aggregation System Priority on BigIron MG8 and NetIron 40G" on page 17-49

Enhancement	Description	See...
VLAN Byte Accounting	With this release, you can configure a VLAN to account for the number of bytes received by all the member ports.	“Configuring VLAN Byte Accounting” on page 22-81

Layer 2 Enhancements in Release 02.2.01

Enhancement	Description	See...
Graceful Restart	With this release, you can enable Graceful Restart for OSPF and BGP.	<i>Foundry Enterprise Configuration and Management Guide</i>
BGP Null0 Routing	With this release, BGP can use null0 to resolve the next hop and install null0 BGP routes to the routing table	<i>Foundry Enterprise Configuration and Management Guide</i>
Configuring Egress Priority Merging	This feature allows you to preserve the incoming 802.1p priority.	<i>Foundry Enterprise Configuration and Management Guide</i>
IP receive ACLs	You can use IPv4 ACLs to filter the packets intended for the management process to protect the management module from being overloaded with heavy traffic that was sent to one of the Layer 3 Switch IP interfaces.	<i>Foundry Enterprise Configuration and Management Guide</i>
OSPF point-to-point	OSPF point-to-point eliminates the need for Designated and Backup Designated routers, allowing for faster convergence of the network.	<i>Foundry Enterprise Configuration and Management Guide</i>
IP option attack prevention	Packets with IP options in their header are automatically dropped. Enabling the ip ip-option-process command allows the device to process packets that use IP options.	<i>Foundry Enterprise Configuration and Management Guide</i>
Static Route Tagging	Static routes can be configured with tag values.	<i>Foundry Enterprise Configuration and Management Guide</i>
Neighbor Local AS	Neighbor Local Autonomous System (AS) feature allows a router that is a member of one AS to appear to be a member of another AS.	<i>Foundry Enterprise Configuration and Management Guide</i>
Multi-protocol Border Gateway Protocol (MBGP)	With this release, you can configure MBGP, an extension to BGP that allows a router to support separate unicast and multicast topologies.	<i>Foundry Enterprise Configuration and Management Guide</i>
MSDP Mesh Groups	With this release, the BigIron MG8 and NetIron 40G support MSDP Mesh Groups. NOTE: The MSDP cache size is 32 K.	<i>Foundry Enterprise Configuration and Management Guide</i>
Full AS Path information in sFlow	In this release, sFlow packets now contain full AP Path information.	<i>Foundry Enterprise Configuration and Management Guide</i>

Enhancement	Description	See...
MTU enhancements for IPv4 and IPv6	In this release, you can configure IPv4 MTU to be greater than 1500 bytes. MTU value for IPv6 is still limited to 1500 bytes.	<i>Foundry Enterprise Configuration and Management Guide and Foundry IPv6 Configuration Guide</i>
ISIS	IPv6 ISIS are now supported on the NetIron 40G. Also, IPv4 ISIS has been enhanced with the addition of address-family configuration level.	<i>Foundry NetIron Service Provider Configuration and Management Guide</i>
IP fragmentation protection	Fragmented IP packets with undersized fragments and overlapping fragments are dropped.	N/A

Security Enhancements in 02.2.01

Enhancement	Description	See...
802.1x port security	This release allows you to enable 802.1X port security and multi-device port authentication on the same interface.	<i>Foundry Security Guide</i>
Multi-device port authentication	Multi-device port authentication is now supported on the BigIron MG8 and NetIron 40G.	<i>Foundry Security Guide</i>
Port Security MAC Deny	With this release, you can configure deny mac addresses on a global level or on a per port level.	<i>Foundry Security Guide</i>
Enhancement to an SSH command	The ip ssh rsa-authentication no command has been changed is ip ssh key-authentication no .	<i>Foundry Security Guide</i>

IPv6 Enhancements in 02.2.01

Enhancement	Description	See...
IPv6 Over IPv4 Tunnels	This release allows you to configure IPv6 over IPv4 tunnels through either of the following mechanisms: <ul style="list-style-type: none"> • Manually configured tunnels • Automatic 6to4 tunnels • Automatic IPv4-compatible IPv6 tunnels 	<i>Foundry IPv6 Configuration Guide</i>

Chapter 2

Installing a Foundry Layer 2 Switch or Layer 3 Switch

This chapter describes how to install Foundry Layer 2 Switches and Layer 3 Switches and attach them to your network. For information about basic software configuration, see “Configuring Basic Features” on page 9-1.

WARNING: The procedures in this manual are for qualified service personnel.

Unpacking a System

The Foundry systems ship with all of the following items. Please review the list below and verify the contents. If any items are missing, please contact the place of purchase.

Package Contents

- Foundry Networks Layer 2 Switch or Layer 3 Switch
- 115V AC power cable (for AC sourced devices)
- Rack mount brackets and mounting screws
- CD-ROM containing software images and the user documentation (including this guide)
- Warranty card

General Requirements

To manage the system, you need the following items for serial connection to the switch or router:

- A management station, such as a PC running a terminal emulation application.
- A straight-through EIA/TIA DB-9 serial cable (F/F). The serial cable can be ordered separately from Foundry Networks. If you prefer to build your own cable, see the pinout information in “Attaching a PC or Terminal” on page 2-17.

You use the serial connection to perform basic configuration tasks including assigning an IP address and network mask to the system. This information is required for managing the system using the Web management interface or IronView Network Manager or using the CLI through Telnet.

WARNING: Do not use the handles on the power supply units to lift or carry a Chassis device.

Summary of Installation Procedures

Follow the steps listed below to install your Layer 2 Switch or Layer 3 Switch. Details for each of the steps highlighted below are provided later in this chapter.

1. Ensure that the physical environment that will host the device has the proper cabling and ventilation. See “Preparing the Installation Site” on page 2-4.
2. Chassis devices only – If needed, insert or remove chassis modules. There are many optional modules designed for the module slots on the Chassis devices. Depending on where you plan to install a device, it might be easier to install the modules first. However, the modules are “hot swappable”, and can be installed or removed after the device is mounted and powered-on. See “Installing or Removing Optional Modules (Chassis Devices Only)” on page 2-5.

NOTE: If you are installing redundant management modules (Management 2 or higher), see “Using Redundant Management Modules” on page 3-1 for complete installation, configuration, and management instructions for the modules.

3. Chassis devices or the FastIron 4802 only – Optionally insert or remove redundant power supplies. The 4-slot Chassis devices and the FastIron 4802 can hold one or two power supplies. The 8-slot and 15-slot Chassis devices can hold up to four power supplies. If you need to install a power supply, it may be easier to install it before mounting the device, although the power supplies are “hot swappable”, and can be installed or removed after the device is mounted and powered-on. See “Installing or Removing Redundant Power Supplies (Chassis Devices Only)” on page 2-6 or “Installing or Removing a Power Supply (FastIron 4802 only)” on page 2-10.

CAUTION: Remove the power cord from a power supply before you install it in or remove it from the device. Otherwise, the power supply or the device could be damaged as a result. (The device can be running while a power supply is being installed or removed, but the power supply itself should not be connected to a power source.)

4. Chassis devices only – Optionally replace cooling fans. Generally, this procedure is not required during installation but is included in case you ever need to replace a fan after the device is placed in operation. See “Replacing Fans (4-Slot and 8-Slot Chassis Devices Only)” on page 2-13 or “Replacing a Fan Tray (15-Slot Chassis Devices Only)” on page 2-16.
5. Verify that the system and module LEDs are registering the proper LED state after power-on of the system. See “Verifying Proper Operation” on page 2-16.
6. A terminal or PC serial port connection is all that is required to support configuration on the device. See “Attaching a PC or Terminal” on page 2-17.0
7. No default password is assigned to the Command Line Interface (CLI). For additional access security, assign a password. See “Assigning Permanent Passwords” on page 2-19.
8. Before attaching equipment to the device, you need to configure an interface IP address to the subnet on which it will be located. Initial IP address configuration is performed using the CLI with a direct serial connection. Subsequent IP address configuration can be performed using the Web management interface. See “Configuring IP Addresses” on page 2-20.
9. Foundry devices can be installed on a desktop or in an equipment rack. See “Mounting the Chassis or Stackable Device” on page 2-22.
10. Once the device is physically installed, plug the device into a nearby power source that adheres to the regulatory requirements outlined in this manual. See “Powering On a System” on page 2-24.
11. Once you power on the device and assign IP addresses, the system is ready to accept network equipment. See “Connecting Network Devices” on page 2-25.
12. Test IP connectivity to other devices by pinging them and tracing routes. See “Testing Connectivity” on page 2-29.

13. Continue configuring the device using the CLI or the Web management interface. See "Managing the Device" on page 2-29.

NOTE: You also can use IronView Network Manager to manage the device. See the *Foundry IronView Network Management User's Guide* for information.

14. Secure access to the device. See the *Foundry Security Guide*.

Installation Precautions

Follow these precautions when installing a Foundry device.

General Precautions

WARNING: All fiber-optic interfaces use Class 1 Lasers.

CAUTION: Do not install the device in an environment where the operating ambient temperature might exceed 40° C (104° F).

CAUTION: Make sure the air flow around the front, sides, and back of the device is not restricted.

CAUTION: To provide additional safety and proper airflow to the device, make sure that slot cover plates are installed on all chassis slots that do not have either a module or power supply installed.

CAUTION: Never leave tools inside the chassis.

WARNING: Metal edges on the power supply units may be sharp.

Lifting Precautions

WARNING: Do not use the handles on the power supply units to lift or carry Chassis devices.

WARNING: Do not lift the 15-slot chassis using the lifting handles unless the chassis is empty. Remove the power supplies and interface modules before lifting the chassis.

WARNING: You can lift the 4-slot and 8-slot Chassis devices when they contain modules and power supplies. However, fully populated chassis are heavy. **TWO OR MORE PEOPLE ARE REQUIRED WHEN LIFTING, HANDLING, OR MOUNTING THESE DEVICES.**

WARNING: Make sure the rack or cabinet housing the device is adequately secured to prevent it from becoming unstable or falling over.

WARNING: Mount the devices you install in a rack or cabinet as low as possible. Place the heaviest device at the bottom and progressively place lighter devices above.

Power Precautions

CAUTION: Use at least two separate branch circuits for the power. This provides redundancy in case one of the circuits fails.

WARNING: Disconnect the power cord from all power sources to completely remove power from the device.

WARNING: Make sure that the power source circuits are properly grounded, then use the power cord supplied with the device to connect it to the power source.

WARNING: If the installation requires a different power cord than the one supplied with the device, make sure you use a power cord displaying the mark of the safety agency that defines the regulations for power cords in your country. The mark is your assurance that the power cord can be used safely with the device.

CAUTION: Ensure that the device does not overload the power circuits, wiring, and over-current protection. To determine the possibility of overloading the supply circuits, add the ampere (amp) ratings of all devices installed on the same circuit as the device. Compare this total with the rating limit for the circuit. The maximum ampere ratings are usually printed on the devices near the input power connectors.

CAUTION: All devices with DC power supplies are intended for installation in restricted access areas only. A restricted access area is where access can be gained only by service personnel through the use of a special tool, lock and key, or other means of security, and is controlled by the authority responsible for the location.

CAUTION: For the DC input circuit to a 15-slot Chassis device (DC power supply part number RPS4DC), make sure there is a 30-amp circuit breaker on the input to the power supply.

CAUTION: For the DC input circuit to a FastIron 4802 (DC power supply part number RPS5DC), make sure there is a 10-amp circuit breaker when installed in the end system.

CAUTION: For DC power supplies in 15-slot Chassis devices (part number RPS4DC), use a grounding wire of at least 10 American Wire Gauge (AWG). For DC power supplies in 4-slot and 8-slot Chassis devices (part number RPS3DC), use at least 14 AWG. For DC power supplies in the FastIron 4802 (part number RPS5DC), use at least 14 AWG.

Preparing the Installation Site

Cabling Infrastructure

Ensure that the proper cabling is installed in the site. See "Hardware Specifications" on page B-1 or www.foundrynetworks.com for a summary of supported cabling types and their specifications.

Installation Location

Before installing the device, plan its location and orientation relative to other devices and equipment. Allow at least 3" of space at the front of the device for the twisted-pair, fiber-optic, and power cabling. Also, allow a minimum of 3" of space between the sides and the back of the device and walls or other obstructions.

Installing or Removing Optional Modules (Chassis Devices Only)

NOTE: If you are installing redundant management modules (Management 2 or higher), see “Using Redundant Management Modules” on page 3-1 for complete installation, configuration, and management instructions for the modules.

Installing Modules

To install a module in the chassis, do the following:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

WARNING: To avoid risk of shock, do not attach the clip end to the air flow panel of the power supply.

2. Remove the blank face plate from the slot in which the module will be installed. Place the blank face plate in a safe place for future use.
3. Remove the module from its packaging.
4. Insert the module into the chassis slot and slide the card along the card guide until the card ejectors on the front of the module touch the chassis.


NOTE: Modules for the 8-slot and 15-slot Chassis devices slide in vertically with port number 1 at the top. Modules for the 4-slot Chassis devices slide in horizontally with port number 1 on the left.

5. Push the ejectors toward the center of the module until they are flush with the front panel of the module. The module will be fully seated in the backplane.
6. Tighten the two screws at either end of the module.

CAUTION: To provide additional safety and proper airflow to the device, make sure that slot cover plates are installed on all chassis slots that do not have either a module or power supply installed.

NOTE: If installing a module into a slot *previously occupied by a different type of module*, you must use the CLI to configure the new module (use the CLI command **module** <slot-num> <module-type>) and then use the **write memory** command to save the configuration and the **reload** command to reset the device. See “Swapping Modules (Chassis devices only)” on page 2-40. If the slot has never contained a module or you are swapping in exactly the same type of module, you do not need to enter these commands.

Figure 2.1 Installing a module



Removing Modules

To remove a module from the chassis, do the following:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

WARNING: To avoid risk of shock, do not attach the clip end to the air flow panel of the power supply.

2. Loosen the two screws on the ends of the module.
3. Pull the card ejectors towards you, and away from the module front panel. The card will unseat from the backplane.
4. Pull the module out of the chassis and place in an anti-static bag for storage.
5. Cover the slot with the blank face plate that shipped with the chassis.

CAUTION: To provide additional safety and proper airflow to the device, make sure that slot cover plates are installed on all chassis slots that do not have either a module or power supply installed.

NOTE: Modules can be installed and removed when the unit is powered on (hot swap). You do not need to power the system down, and you do not need to change the slot's configuration unless you plan to insert a different type of module. However, you do need to disable the module before removing it. See "Swapping Modules (Chassis devices only)" on page 2-40.

Installing or Removing Redundant Power Supplies (Chassis Devices Only)

Determining Power Supply Status

If you are replacing a power supply that has failed and you are not sure which supply has failed, enter the following command at any CLI command prompt:

```
BigIron# show chassis
```

This command displays status information for the fans and the power supplies. The power supplies are numbered in the display. The power supply numbers correspond to the following positions. These positions assume you are facing the front of the chassis, not the rear.

Table 2.1: Power Supply Positions in Foundry Chassis Devices

Product	Power Supply 1 Position	Power Supply 2 Position	Power Supply 3 Position	Power Supply 4 Position
4-slot Chassis device	left side	right side	n/a	n/a
8-slot Chassis device	bottom	second from bottom	second from top	top
15-slot Chassis device	left side	second left	second right	right side

Chassis Devices – AC Power Supplies

Use the following procedures for AC power supplies in 4-slot, 8-slot or 15-slot Chassis devices. Use power supply model RPS3 for 4-slot or 8-slot Chassis devices. Use power supply model RPS4 for 15-slot Chassis devices.

Installing an AC Power Supply

To install a power supply in the chassis, do the following:

1. Use a screwdriver to remove the blank power supply face plate. This will expose the empty power supply slot.
2. Remove the power supply from its packaging.


WARNING: Metal edges on the power supply may be sharp.

3. Hold the bar on the front panel of the power supply and insert the power supply into the empty power supply slot. Use the module guides provided on either side of the compartment.

CAUTION: Carefully follow the mechanical guides on each side of the power supply slot and make sure the power supply is properly inserted in the guides. Never insert the power supply upside down.

4. Continue to slide the power supply towards the back of the chassis until the two metal rods and the connector make contact with the back connector. Then push the power supply until the front panel of the power supply is flush with the rest of the chassis.
5. Use a screwdriver to tighten the two screws on either side of the power supply.
6. Connect the power cord to the front of the power supply.
7. Connect the power plug into an outlet.

Figure 2.2 Installing a power supply (4-slot chassis shown)



Removing an AC Power Supply

To remove a power supply module from the chassis, do the following:

WARNING: Power supplies are hot swappable. However, Foundry Networks recommends that you disconnect the power supply from AC power before installing or removing the supply. The device can be running while a power supply is being installed or removed, but the power supply itself *should not be connected* to a power source. Otherwise, you could be injured or the power supply or other parts of the device could be damaged.

1. Unplug the power supply AC power cord from the outlet.
2. Disconnect the power cord from the power supply.
3. Use a screwdriver to loosen the screws on either side of the power supply.
4. Hold the bar on the front panel of the power supply and pull outward. This will disconnect the power supply from the backplane.

WARNING: Metal edges on the power supply may be sharp.

5. Continue to pull the power supply until it is removed from the chassis.
6. Place the power supply in an anti-static bag for storage.
7. Cover the power supply slot with the blank power supply cover that came with the device.
8. Use a screwdriver to tighten the screws.

Chassis Devices – DC Power Supplies

Use the following procedures for DC power supplies in 4-slot, 8-slot or 15-slot Chassis devices. Use power supply model RPS3DC for 4-slot or 8-slot Chassis devices. Use power supply model RPS4DC for 15-slot Chassis devices.


WARNING: Before beginning the installation, see the precautions in “Power Precautions” on page 2-4.

4-Slot and 8-Slot Chassis Devices

The following procedures describe how to install or remove a DC power supply in a 4-slot or 8-slot Chassis device.

WARNING: Before beginning the installation, see the precautions in “Power Precautions” on page 2-4.

Figure 2.3 DC power supply for 4-slot or 8-slot Chassis device (part number RPS3DC)



Installing a DC Power Supply

1. Prepare the positive, negative, and ground wires by stripping about 1/4" of insulation off the end of each one. (Use at least 14 AWG wire.)
2. Loosen the three screws used to hold the wires in the connector. These are the wires under the following markings:

— + ⊕

NOTE: If you cannot easily reach the screws and wire openings due to the connector guard, remove the connector guard from the power supply. The guard is fastened to the supply by two Phillips-head screws and is mounted over the connector.

3. Slip the ground wire into the opening under the ⊕ marking until the wire is fully in place, then tighten the screw to hold the wire in place.
4. Repeat for the negative (—) and positive (+) wires.
5. Pull gently on each wire to make sure they are securely fastened in the connector.

6. Re-attach the connector guard over the connector.
7. Insert the power supply into the chassis. Seat the supply firmly so that the faceplate of the supply is flush with the chassis surface.

WARNING: Metal edges on the power supply may be sharp.

8. Tighten the two thumb screws to secure the power supply to the chassis.
9. After the power supply is properly inserted, connect the power source to the wires to activate the circuit.

Removing a DC Power Supply

1. Turn off the DC power source or disconnect it from the power supply.
2. Loosen the three screws used to hold the wires in the connector, then pull out the wires.

NOTE: If you cannot easily reach the screws and wire openings due to the connector guard, remove the connector guard from the power supply. The guard is fastened to the supply by two Phillips-head screws and is mounted over the connector. If you remove the connector guard, re-attach it once you are finished removing the wires to prevent the guard from being lost.

3. Loosen the two thumb screws that secure the power supply to the chassis.
4. Pull the power supply completely out of the chassis.


WARNING: Metal edges on the power supply may be sharp.

15-Slot Chassis Devices


The following procedures describe how to install or remove a DC power supply in a 15-slot Chassis device.

WARNING: Before beginning the installation, see the precautions in "Power Precautions" on page 2-4.

Figure 2.4 DC power supply for 15-slot Chassis device (part number RPS4DC)



Installing a DC Power Supply

1. Prepare the positive, negative, and ground wires by stripping about 1/4" of insulation off the end of each one. (Use 10 AWG wire.)
2. Slip one of the lugs shipped with the supply over the stripped end of the ground wire. All three lugs are the same size.
3. Use the crimper to crimp the lug snugly onto the wire. Gently pull the lug away from the wire to verify that the lug is securely fastened.
4. Repeat for the negative and positive wires.
5. Slip the lug of the ground wire under the head of the ground screw (under the  marking) so that the bent part of the lug is facing toward you, away from the power supply.
6. Tighten the screw to hold the lug securely in place.
7. Repeat for the negative (—) and positive (+) wires.
8. Slide the clear plastic wire cover onto the wire connector housing, over the three screw heads.
9. Insert the power supply into the chassis. Make sure the supply is oriented so that the wire connector is near the top. Seat the supply firmly so that the faceplate of the supply is flush with the chassis surface.

WARNING: Metal edges on the power supply may be sharp.

10. Tighten the two thumb screws to secure the power supply to the chassis.
11. After the power supply is properly inserted, connect the power source to the wires to activate the circuit.

Removing a DC Power Supply

1. Turn off the DC power source or disconnect it from the power supply.
2. Slide the clear plastic wire cover off of the wire connector housing, to expose the three screw heads.
3. Loosen the three screws used to hold the wires, then pull the wires out.
4. Slide the clear plastic wire cover back over the wire connector housing, to prevent the cover from being lost.
5. Loosen the two thumb screws that secure the power supply to the chassis.
6. Remove the power supply completely from the chassis.

WARNING: Metal edges on the power supply may be sharp.

Installing or Removing a Power Supply (FastIron 4802 only)

Use the following procedures to insert or remove a power supply in a FastIron 4802.

WARNING: Power supplies are hot swappable. However, Foundry Networks recommends that you disconnect the power supply from AC power before installing or removing the supply. The device can be running while a power supply is being installed or removed, but the power supply itself should not be connected to a power source. Otherwise, you could be injured or the power supply or other parts of the device could be damaged.

Determining Power Supply Status

If you are replacing a power supply that has failed and you are not sure which supply has failed, enter the following command at any CLI command prompt:

```
SW-FI4802-PREM# show chassis
```

This command displays status information for the fans and the power supplies. The power supplies are numbered from left to right. These numbers assume you are facing the front of the chassis, not the rear.

FastIron 4802 – AC Power Supplies

Use the following procedures for AC power supplies in the FastIron 4802. Use power supply model RPS5.

Installing an AC Power Supply

To install a power supply in the FastIron 4802, do the following:

1. If the empty power supply bay has a cover plate, press the two latches near the edges of the supply inward to unlock the plate, then remove the plate.
2. Remove the power supply from its packaging.

WARNING: Metal edges on the power supply may be sharp.

3. With one hand, hold the bar on the front panel of the power supply. With the other hand, support the underside of the power supply, and insert the power supply into the empty power supply slot. Press until the supply is completely in the slot, so that the connectors on the back of the supply are fully engaged with the pins on the power backplane.

CAUTION: Make sure you insert the power supply right-side up. It is possible to insert the supply upside down, although the supply will not engage with the power backplane when upside down. The power supply is right-side up when the power connector is on the left and the fan vent is on the right.

4. Press the two latches near the edges of the supply outward to lock the supply in place.
5. Connect the power cord to the power supply.
6. Connect the plug end of the power cord into an outlet.

Removing an AC Power Supply

1. Unplug the power supply from the power source.
2. Disconnect the power cord from the power supply.
3. Press the two latches near the edges of the supply inward to unlock the supply.
4. Hold the bar on the front panel of the power supply and pull outward. This will disconnect the power supply from the backplane.
5. Continue to pull the power supply until it is removed from the device.
6. Place the power supply in an anti-static bag for storage.
7. Insert a new supply, or place the cover plate over the empty power supply bay and press the two latches near the edges of the supply outward to lock the plate into place.


WARNING: Metal edges on the power supply may be sharp.

FastIron 4802 – DC Power Supplies

Use the following procedures for DC power supplies in the FastIron 4802. Use power supply model RPS5DC.

WARNING: Before beginning the installation, see the precautions in “Power Precautions” on page 2-4.

Figure 2.5 DC power supply for FastIron 4802 (part number RPS5DC)



Installing a DC Power Supply

1. Prepare the positive, negative, and ground wires by stripping about 1/4" of insulation off the end of each one. (Use 14 AWG wire.)
2. Loosen the three screws used to hold the wires in the connector. These are the wires under the following markings:

— + ⊕

3. Slip the ground wire into the opening under the ⊕ marking until the wire is fully in place, then tighten the screw to hold the wire in place.
4. Repeat for the negative (—) and positive (+) wires.
5. Pull gently on each wire to make sure they are securely fastened in the connector.
6. With one hand, hold the bar on the front panel of the power supply. With the other hand, support the underside of the power supply, and insert the power supply into the empty power supply slot. Press until the supply is completely in the slot, so that the connectors on the back of the supply are fully engaged with the pins on the power backplane.

WARNING: Metal edges on the power supply may be sharp.

CAUTION: Make sure you insert the power supply right-side up. It is possible to insert the supply upside down, although the supply will not engage with the power backplane when upside down. The power supply is right-side up when the power connector is on the left and the fan vent is on the right.

7. Press the two latches near the edges of the supply outward to lock the supply in place.
8. After the power supply is properly inserted, connect the power source to the wires to activate the circuit.

Removing a DC Power Supply

1. Turn off the DC power source or disconnect it from the power supply.
2. Loosen the three screws used to hold the wires in the connector, then pull out the wires.
3. Press the two latches near the edges of the supply inward to unlock the supply.
4. Hold the bar on the front panel of the power supply and pull outward. This will disconnect the power supply from the backplane.
5. Continue to pull the power supply until it is removed from the device.
6. Place the power supply in an anti-static bag for storage.
7. Insert a new supply, or place the cover plate over the empty power supply bay and press the two latches near the edges of the supply outward to lock the plate into place.

WARNING: Metal edges on the power supply may be sharp.

Replacing Fans (4-Slot and 8-Slot Chassis Devices Only)

The 4-slot and 8-slot Chassis devices contain field-upgradable fans. The fans are upgradable on an individual basis. You need to replace only the fan that has failed.

The 4-slot Chassis devices contain four fans:

- Two fans are mounted to the inside of the rear chassis panel.
- Two fans are mounted to a removable tray on the upper left side of the chassis. (The fans are on the right side if you are facing the rear of the chassis.)

The 8-slot Chassis devices contain six fans:

- Two fans are mounted to the inside of the rear chassis panel.
- Four fans are mounted to two removable trays in the top of the chassis, above the highest module slot. The fans are on the right if you are facing the front of the chassis.

Each fan in a four-slot or eight-slot chassis is connected to the chassis backplane by a three-hole connector. Make a note of the connector each fan uses. The software recognizes the fan position based on the connector.

NOTE: When you connect a fan cable to a fan connector on the backplane or fan tray, make sure the red wire in the connector is on the right side (for horizontally oriented connectors) or facing down (for vertically oriented connectors). If you accidentally reverse the wires, the fan will not operate.

Also, make sure the fan cable connector is seated over all three pins on the backplane connector.

Required Tools

You need the following tools for this procedure:

- Phillips-head screwdriver
- Flat-head screwdriver
- Pair of wire cutters

Determining Which Fan Has Failed

If you are not sure which fan has failed, enter the following command at any CLI command prompt:

```
BigIron# show chassis
```

This command displays status information for the fans and the power supplies. The fans are numbered in the display. The fan numbers correspond to the following fan positions. These positions assume you are facing the front of the chassis, not the rear.

Table 2.2: Fan Positions in Foundry Chassis Devices

Product	Fan 1 Position	Fan 2 Position	Fan 3 Position	Fan 4 Position
4-slot Chassis device	Fan tray on left side; back fan	Fan tray on left side; front fan	Rear fan, left side	Rear fan, right side
8-slot Chassis device	Rear fan, top	Rear fan, bottom	top fan tray, left side	top fan tray, right side

NOTE: The software monitors the fans in the top of the 8-slot chassis in pairs, not individually. Thus, fan position 3 indicates the left fan tray and fan position 4 indicates the right fan tray.

Four-Slot Chassis

To replace a fan in a 4-slot chassis:

1. Power down the chassis and remove the power cables from the chassis power supplies.
2. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

WARNING: To avoid risk of shock, do not attach the clip end to the air flow panel of the power supply.

3. Remove all 18 Phillips-head screws from the rear panel of the chassis.

NOTE: The fans on the rear panel are connected to the chassis backplane by wire cables. Be careful when you remove the rear panel to avoid accidentally damaging the cables or connectors.

4. Unplug the fan cables from the backplane and set the rear panel on a workbench. If you do not need to replace a fan in the fan tray mounted on the side of the chassis, skip to step 8; otherwise, go to step 5.
5. Loosen the two flat-head screws that fasten the side fan tray to the chassis.
6. Carefully pull the side fan tray out of the chassis and set the tray on a workbench.

NOTE: The fastener push-ons that fasten the fans to the fan tray may catch on the chassis. In this case, gently move the fan tray from side to side as you pull the tray back to free it from the chassis.

7. Unplug the fan cables from the backplane and set the fan tray on the workbench.
8. Use the wire cutters to cut the tie wraps fastening the wires of the two fans together.
9. Gently use the wire cutters or similar tool to remove the four plastic fastener push-ons that fasten the failed fan to the rear panel or fan tray.

NOTE: Be careful when removing the fastener push-ons. They are reusable.

10. Remove the fan.
11. Align the new fan over the fastener holes on the rear panel or fan tray, then insert the fastener push-ons to fasten the new fan in place.
12. Fasten new tie wraps around the wires to keep them neatly together and away from other components.
13. If you replaced a fan on the rear panel and did not remove the fan tray mounted on the side of the chassis, skip to step 15; otherwise, go to step 14.
14. Gently reinsert the fan tray into the chassis and partially tighten both screws. Then tighten the upper screw, then the lower screw.

NOTE: Make sure you tighten the upper screw first to properly align the tray in the chassis.

15. Plug the fan cables into the three-pin connectors on the backplane.

CAUTION: When you connect a fan cable to a fan connector on the backplane, make sure the red wire in the connector is on the right side of the connector. If you accidentally reverse the wires, the fan will not operate.

Also, make sure the fan cable connector is seated over all three pins on the backplane connector.

16. Align the rear panel over the rear screw holes.
17. Screw the 18 Phillips-head screws back in.

18. Verify that all chassis modules and power supplies are fully seated and all cover plates and panels are fully fastened.
19. Reconnect the power and power on the chassis.
20. Access the CLI and enter the **show chassis** command to verify that all fans are now operating normally.

Eight-Slot Chassis

To replace a fan in an 8-slot chassis:

1. Remove the power cables from the chassis power supplies.
2. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

WARNING: To avoid risk of shock, do not attach the clip end to the air flow panel of the power supply.

3. Remove all 34 Phillips-head screws from the rear panel of the chassis.

NOTE: The fans on the rear panel are connected to the chassis backplane by wire cables. Be careful when you remove the rear panel to avoid accidentally damaging the cables or connectors.

4. Unplug the fan cables from the backplane and set the rear panel on the workbench. If you do not need to replace a fan in one of the fan trays mounted on the top of the chassis, skip to step 8; otherwise, go to step 5.
5. Loosen the two flat-head screws that fasten the fan tray containing the failed fan to the chassis.
6. Unplug the fan cables from the backplane.
7. Carefully pull the fan tray out of the chassis and set the tray on a workbench.

NOTE: The fastener push-ons that fasten the fans to the fan rack may catch on the chassis. In this case, gently move the fan rack from side to side as you pull the rack back to free it from the chassis.

8. Use the wire cutters to cut the tie wraps fastening the wires of the two fans together.
9. Gently use the wire cutters or similar tool to remove the four plastic fastener push-ons that fasten the failed fan to the rear panel or fan tray.

NOTE: Be careful when removing the fastener push-ons. They are reusable.

10. Remove the fan.
11. Align the new fan over the fastener holes on the rear panel or fan tray, then insert the fastener push-ons to fasten the new fan in place.
12. Fasten new tie wraps around the wires to keep them neatly together and away from other components.
13. If you replaced a fan on the rear panel and did not remove the fan tray in the side of the chassis, skip to step 15; otherwise, go to step 14.
14. Gently reinsert the fan tray into the chassis and tighten both screws.
15. Plug the fan cables onto the three-pin connectors on the backplane.

CAUTION: When you connect a fan cable to a fan connector on the backplane, make sure the red wire in the connector is on the right side (for horizontally oriented connectors) or facing down (for vertically oriented connectors). If you accidentally reverse the wires, the fan will not operate.

Also, make sure the fan cable connector is seated over all three pins on the backplane connector.

16. Align the rear panel over the rear screw holes.

17. Screw the 34 Phillips-head screws back in.
18. Verify that all chassis modules and power supplies are fully seated and all cover plates and panels are fully fastened.
19. Reconnect the power cables and power on the chassis.
20. Access the CLI and enter the **show chassis** command to verify that all fans are now operating normally.

Replacing a Fan Tray (15-Slot Chassis Devices Only)

The 15-slot Chassis devices contain field-upgradable, hot-swappable fans. If a fan fails, you can remove the fan tray and replace it with a new fan tray without powering off the chassis device.

NOTE: To avoid overheating, do not leave the chassis powered on for more than a few minutes without a fan tray installed.

To replace a fan in a 15-slot chassis:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

WARNING: To avoid risk of shock, do not attach the clip end to the air flow panel of the power supply.

2. Loosen the two screws on the fan tray. The fan tray is located above the power supply bays and below the air filter tray.
3. Carefully pull the fan tray out of the chassis and set the tray on a workbench or other static-free area.
4. Insert the new fan tray into the fan tray slot and push it in until the face plate is flush with the chassis.
5. Tighten the two screws on the fan tray.
6. Access the CLI and enter the **show chassis** command to verify that all fans are operating normally.

Verifying Proper Operation

After you have installed any modules or redundant power supplies, but before mounting the device in its network location, verify that the device is working properly by plugging it into a power source and verifying that it passes its self test.

If your device has more than one power supply installed, repeat this procedure for each power supply.

1. Connect the power cord supplied with the device to the power connector on the power supply on the front of the device.
2. Insert the other end into a properly grounded electrical outlet.
3. Verify that the LED on each power supply is a solid green.

NOTE: The devices do not have power switches. They power on when you connect a power cord to the device and to a power source.

If your installation requires a different power cord than that supplied with the device, make sure you obtain a power cord displaying the mark of the safety agency that defines the regulations for power cords in your country. The mark is your assurance that the power cord can be used safely with the device.

4. Verify proper operation by observing the LEDs:
 - Chassis devices – Make sure the LED on each power supply is a solid green. Also make sure that some of the port LEDs on each module momentarily light up. The LEDs indicate that the device is performing diagnostics. After the diagnostics are complete, the LEDs will be dark except for the ones that are attached by cables to other devices. If the links on these cables are good and the connected device is

powered on, the link LEDs will light.

NOTE: If all of the LEDs on a module do not light up during the diagnostics, this does not indicate an error. Only some of the LEDs are lighted during the diagnostics.

- Fixed-port devices (Stackable devices) – All the port LEDs should flash momentarily, usually in sequence, while the device performs diagnostics. After the diagnostics are complete, the LEDs will be dark except for the ones that are attached by cables to other devices. If the links on these cables are good and the connected device is powered on, the link LEDs will light.

For more details on specific LED conditions after system start-up, see “Hardware Specifications” on page B-1.

Attaching a PC or Terminal

To assign an IP address, you must have access to the **Command Line Interface (CLI)**. The CLI is a text-based interface that can be accessed through a direct serial connection to the device and through Telnet connections. The CLI is described in detail in the *Foundry Switch and Router Command Line Interface Reference*.

You need to assign an IP address using the CLI. You can access the CLI by attaching a serial cable to the Console port. After you assign an IP address, you can access the system through Telnet, the Web management interface, or IronView Network Manager.

To attach a management station using the serial port:

1. Connect a PC or terminal to the serial port of the system using a straight-through cable. The serial port has a male DB-9 connector.

NOTE: You need to run a terminal emulation program on the PC.

2. Open the terminal emulation program and set the session parameters as follows:
 - Baud: 9600 bps
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None

When you establish the serial connection to the system, press Enter to display one of the following CLI prompts in the terminal emulation window:

- BigIron>
- FastIron>
- FastIronII>
- FI4802>
- FI4802-PREM>
- NetIron>
- ServerIron>
- TurboIron>

NOTE: If you install Layer 2 Switch code on a Layer 3 Switch, the command prompt begins with “SW-” to indicate the software change. This is true even if you change the system name. If you install Base Layer 3 code, the prompt begins with “BR-”.

If you see one of these prompts, you are now connected to the system and can proceed to “Assigning Permanent Passwords” on page 2-19.

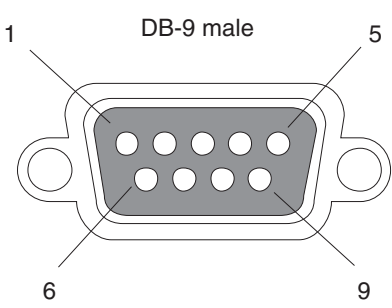
You can customize the prompt by changing the system name. See “Entering System Administration Information” on page 9-4.

If you do not see one of these prompts:

1. Make sure the cable is securely connected to your PC and to the Foundry system.
2. Check the settings in your terminal emulation program. In addition to the session settings listed above, make sure the terminal emulation session is running on the same serial port you attached to the Foundry system.

The EIA/TIA 232 serial communication port serves as a connection point for management by a PC or SNMP workstation. Foundry switches and Layer 3 Switches come with a standard male DB-9 connector, shown in Figure 2.6.

Figure 2.6 Serial port pin and signalling details

Pin Assignment	Pin Number	Switch Signal
 <p style="text-align: center;">DB-9 male</p>	1	Reserved
	2	TXD (output)
	3	RXD (input)
	4	Reserved
	5	GND
	6	Reserved
	7	CTS (input)
	8	RTS (output)
	9	Reserved

Most PC serial ports also require a cable with a female DB-9 connector.

Terminal connections will vary, requiring either a DB-9 or DB-25 connector, male or female.

Serial cable options between a Foundry switch or router and a PC or terminal are shown in Figure 2.7.

NOTE: As indicated in Figure 2.6 and Figure 2.7, some of the wires should not be connected. If you do connect the wires that are labeled “Reserved”, you might get unexpected results with some terminals.

Figure 2.7 Serial port pin assignments showing cable connection options to a terminal or PC

DB-9 to DB-9 Female Switch			Terminal or PC	DB-9 to DB-25 Female Switch			Terminal or PC
1	Reserved		1	1	Reserved		8
2	→		2	2	→		3
3	←		3	3	←		2
4	Reserved		4	4	Reserved		20
5	→		5	5	→		7
6	Reserved		6	6	Reserved		6
7	←		7	7	←		4
8	→		8	8	→		5
9	Reserved		9	9	Reserved		22

Assigning Permanent Passwords

The CLI contains the following access levels:

- EXEC at the User level – The level you enter when you first start a CLI session. At this level, you can view some system information but you cannot configure system or port parameters.
- EXEC at the Privileged level – This level is also called the Enable level and can be secured by a password. You can perform tasks such as manage files on the flash module, save the system configuration to flash, and clear caches at this level.
- CONFIG – The configuration level. This level lets you configure the system's IP address and configure switching and routing features. To access the CONFIG mode, you must already be logged into the Privileged level of the EXEC mode.

By default, there are no CLI passwords. To secure CLI access, you must assign passwords. See the *Foundry Security Guide*.

NOTE: You cannot assign a password using the Web management interface. You can assign passwords using the IronView Network Manager if an Enable password for a Super User is already configured on the device.

You can set the following levels of Enable passwords:

- Super User – Allows complete read-and-write access to the system. This is generally for system administrators and is the only password level that allows you to configure passwords.

NOTE: You must set a super user password before you can set other types of passwords.

- Port Configuration – Allows read-and-write access for specific ports but not for global (system-wide) parameters.
- Read Only – Allows access to the Privileged EXEC mode and CONFIG mode but only with read access.

USING THE CLI

To set passwords:

1. At the opening CLI prompt, enter the following command to change to the Privileged level of the EXEC mode:

```
BigIron> enable
```

2. Access the CONFIG level of the CLI by entering the following command:

```
BigIron# configure terminal
BigIron(config)#
```

3. Enter the following command to set the super-user password:

```
BigIron(config)# enable super-user-password <text>
```

NOTE: You must set the super-user password before you can set other types of passwords.

4. Enter the following commands to set the port configuration and read-only passwords:

```
BigIron(config)# enable port-config-password <text>
BigIron(config)# enable read-only-password <text>
```

NOTE: If you forget your super-user password, see the Release Notes.

Syntax: enable super-user-password | read-only-password | port-config-password <text>

Passwords can be up to 32 characters long.

Configuring IP Addresses

You must configure at least one IP address using the serial connection to the CLI before you can manage the system using the other management interfaces. In addition, Foundry routers require an IP subnet address for the subnet in which you plan to place them in your network.

Foundry devices support both classical IP network masks (Class A, B, and C subnet masks, and so on) and Classless Interdomain Routing (CIDR) network prefix masks.

- To enter a classical network mask, enter the mask in IP address format. For example, enter “209.157.22.99 255.255.255.0” for an IP address with a Class-C subnet mask.
- To enter a prefix number for a network mask, enter a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter “209.157.22.99/24” for an IP address that has a network mask with 24 significant (“mask”) bits.

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the display to the prefix format. See the “Configuring IP” chapter in the *Foundry Enterprise Configuration and Management Guide*.

NOTE: If your network uses a Bootstrap Protocol (BootP) server or a Dynamic Host Configuration Protocol (DHCP) server, you can allow the Foundry device to obtain IP addresses for the hosts on the network.

Layer 3 Switches

Before attaching equipment to a Foundry router, you must assign an interface IP address to the subnet on which the router will be located. You must use the serial connection to assign the first IP address. For subsequent addresses, you also can use the CLI through Telnet or the Web management interface.

By default, you can configure up to 24 IP interfaces on each port, virtual routing interface, and loopback interface. On Stackable Layer 3 Switches, you can increase this amount to up to 64 IP subnet addresses per port by increasing the size of the subnet-per-interface table. See “Displaying and Modifying System Parameter Default Settings” on page 9-48.

The following procedure shows how to add an IP address and mask to a router port.

1. At the opening CLI prompt, enter **enable**.

```
BigIron> enable
```

2. Enter the following command at the Privileged EXEC level prompt (for example, BigIron#), then press Enter. This command erases the factory test configuration if still present:

```
BigIron# erase startup-config
```

CAUTION: Use the **erase startup-config** command only for new systems. If you enter this command on a system you have already configured, the command erases the configuration. If you accidentally do erase the configuration on a configured system, enter the **write memory** command to save the running configuration to the startup-config file.

3. Access the configuration level of the CLI by entering the following command:

```
BigIron# configure terminal          Privileged EXEC Level
BigIron(config)#                   Global CONFIG Level
```

4. Configure the IP addresses and mask addresses for the interfaces on the router.

```
BigIron(config)# int e 1/5
BigIron(config-if-1/5)# ip address 192.22.3.44 255.255.255.0
```

NOTE: You can use the syntax **ip address <ip-addr>/<mask-bits>** if you know the subnet mask length. In the above example, you could enter **ip address 192.22.3.44/24**.

Syntax: enable [<password>]

Syntax: configure terminal

Syntax: [no] ip address <ip-addr> <ip-mask> [secondary]

or

Syntax: [no] ip address <ip-addr>/<mask-bits> [secondary]

Use the **secondary** parameter if you have already configured an IP address within the same subnet on the interface.

Layer 2 Switches

To configure an IP Address to a Foundry switch:

1. At the opening CLI prompt, enter **enable**.

```
FastIronII> enable
```

2. Enter the following command at the Privileged EXEC level prompt (for example, FastIronII#), then press Enter. This command erases the factory test configuration if still present:

```
FastIronII# erase startup-config
```

CAUTION: Use the **erase startup-config** command only for new systems. If you enter this command on a system you have already configured, the command erases the configuration. If you accidentally do erase the configuration on a configured system, enter the **write memory** command to save the running configuration to the startup-config file.

3. Access the configuration level of the CLI by entering the following command:

```
FastIronII# configure terminal          Privileged EXEC Level
FastIronII(config)#                   Global CONFIG Level
```

4. Configure the IP address and mask for the switch.

```
FastIronII(config)# ip address 192.22.3.44 255.255.255.0
```

5. Set a default gateway address for the switch.

```
FastIronII(config)# ip default-gateway 192.22.3.1
```

NOTE: You do not need to assign a default gateway address for single subnet networks.

Syntax: enable [<password>]

Syntax: configure terminal

Syntax: [no] ip address <ip-addr> <ip-mask>

or

Syntax: [no] ip address <ip-addr>/<mask-bits>

Syntax: ip default-gateway <ip-addr>

Mounting the Chassis or Stackable Device

You can install Foundry systems on a desktop or in an equipment rack.

WARNING: The Chassis devices are very heavy, especially when fully populated with modules and power supplies. TWO OR MORE PEOPLE ARE REQUIRED WHEN LIFTING, HANDLING, OR MOUNTING THESE DEVICES.

WARNING: Do not use the handles on the power supply units to lift or carry Chassis devices.

WARNING: Make sure the rack or cabinet housing the device is adequately secured to prevent it from becoming unstable or falling over.

WARNING: Mount the devices you install in a rack or cabinet as low as possible. Place the heaviest device at the bottom and progressively place lighter devices above.

Desktop Installation


1. Set the device on a flat desktop, table, or shelf. Make sure that adequate ventilation is provided for the system – a 3-inch clearance is recommended on each side.
2. Go to “Testing Connectivity” on page 2-29.

Rack Mount Installation – Chassis Devices

1. Remove the rack mount kit from the shipping carton. The kit should include two L-shaped mounting brackets and mounting screws.

NOTE: You need a #2 Phillips-head screwdriver for installation.

2. Attach the mounting brackets to the sides of the device as illustrated in Figure 2.8.
3. Attach the system in the rack as illustrated in Figure 2.8.
4. Go to “Powering On a System” on page 2-24.

Figure 2.8 Installing a Chassis device in a rack mount


Rack Mount Installation – Stackable Devices

NOTE: You need a #2 Phillips-head screwdriver for installation.

1. Remove the rack mount kit from the shipping carton. The kit contains two L-shaped mounting brackets and mounting screws.
2. Attach the mounting brackets to the sides of the device as illustrated in Figure 2.9.
3. Attach the device in the rack as illustrated in Figure 2.9.
4. Proceed to “Testing Connectivity” on page 2-29.

NOTE: If you are installing a Chassis device, see “Installing or Removing Optional Modules (Chassis Devices Only)” on page 2-5 and “Installing or Removing Redundant Power Supplies (Chassis Devices Only)” on page 2-6 before proceeding to “Testing Connectivity” on page 2-29.

Figure 2.9 Installing a Stackable device in a rack mount



Powering On a System

After you complete the physical installation of the system, you can power on the system.

1. Ensure that all modules and power supplies are fully and properly inserted and no module slots or power supply slots are uncovered.

CAUTION: Never leave tools inside the chassis.

2. Remove the power cord from the shipping package.
3. Attach the AC power cable to the AC connector on the rear panel. For Chassis devices, the AC connector is located on the front of the Chassis device, embedded within each power supply.
4. Insert the power cable plug into a 115V/120V outlet.

NOTE: When you power on a Chassis device that requires multiple power supplies, make sure you apply power to all the supplies (or at least the minimum number of supplies required for your configuration) at the same time. Otherwise, the device either will not boot at all, or will boot and then repeatedly display a warning message stating that you need to add more power supplies.

NOTE: Foundry devices are designed to provide uninterrupted service even when you insert or remove modules. Therefore, the systems do not have separate on/off power switches. To turn the system off, simply unplug the power cord(s).

NOTE: The socket should be installed near the equipment and should be easily accessible.

NOTE: If the outlet is not rated 115/120V, stop and get the appropriate cable for the outlet.

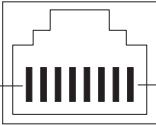
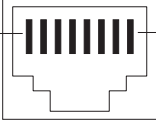
Connecting Network Devices

Foundry devices can support connections to other vendors' routers, switches, and hubs as well other Foundry devices.

Connectors

- 10BaseT/100BaseTX ports come with RJ45 jacks for standard unshielded twisted pair (UTP/Category 5) cable connections.
- 100BaseFX ports come equipped with MT-RJ connectors.
- 1000BaseSX ports come equipped with SC connectors.
- 1000BaseLX ports come equipped with SC connectors.
- 1000BaseLH ports come equipped with SC connectors.
- 1000BaseT ports come equipped with RJ-45 connectors.

Figure 2.10 Pin assignment and signalling for 10/100BaseTX and 1000BaseT ports

Pin Assignment	10BaseT		100BaseTX and 1000BaseT	
	Pin Number	MDI-X ports	Pin Number	MDI-X ports
	1	RD+	1	RD+
	2	RD-	2	RD-
	3	TD+	3	TD+
	4	Not used	4	CMT
	5	Not used	5	CMT
	6	TD-	6	TD-
	7	Not used	7	CMT
	8	Not used	8	CMT

Cable Length

- 100BaseTX: Cable length should not exceed 100 meters.
- 1000BaseTX: Cable length should not exceed 100 meters.
- 100BaseFX: Cable length should not exceed 2 kilometers.
- 1000BaseSX: Cable length should not exceed 550 meters when operating with multi-mode cabling.
- 1000BaseLX:
 - Cable length of 2 – 550 meters is supported on 62.5 μm multi-mode fiber (MMF) cabling.
 - Cable length of 2 – 550 meters is supported on 50 μm multi-mode fiber (MMF) cabling.
 - Cable length of 2 – 5000 meters is supported on 9 μm single-mode fiber (SMF) cabling.
- 1000BaseLH: Cable length should not exceed 70 kilometers for LHA or 150 kilometers for LHB.

Table 2.3: Cable length summary table



	Fiber Type	Core Diameter (microns)	Modal Bandwidth (MHz*km)	Minimum Range (meters)
1000BaseSX	MMF	62.5	160	2 – 200 ^a
	MMF	62.5	200	2 – 275 ^b
	MMF	50	400	2 – 500
	MMF	50	500	2 – 550 ^c
1000BaseLX	MMF	62.5	500	2 – 550
	MMF	50	400	2 – 550
	MMF	50	500	2 – 550
	SMF	9	n/a	2 – 5000
1000BaseLHA	SMF	9	n/a	2 – 70000 (70km)
1000BaseLHB	SMF	9	n/a	2 – 150000 (150km)

- a. The TIA 568 building wiring standard specifies 160/500 MHz*km MMF (Multi-mode Fiber).
- b. The international ISO/IEC 11801 building wiring standard specifies 200/500 MHz*km MMF.
- c. The ANSI Fibre Channel specification specifies 500/500 MHz*km 50 micron MMF and 500/500 MHz*km fiber has been proposed for addition to ISO/IEC 11801.

NOTE: Cable installation and network configuration will affect overall transmission capability. The numbers provided above represent the accepted recommendations of the various standards. For network-specific recommendations, consult your local Foundry reseller or system engineer.

Connecting to Ethernet or Fast Ethernet Hubs

For connections to Ethernet hubs, a 10/100BaseTX or 1000BaseT switch, or another Foundry device, a crossover cable is required (Figure 2.11 or Figure 2.12). If the hub is equipped with an uplink port, it will require a straight-through cable instead of a crossover cable.

Figure 2.11 UTP crossover cable**Figure 2.12 Cat-5 crossover cable for 1000BaseT**

NOTE: The 802.3ab standard calls for automatic negotiation of the connection between two 1000BaseT ports. Consequently, a crossover cable may not be required; a straight-through cable may work as well.

Connecting to Workstations, Servers, or Routers

Straight-through UTP cabling is required for direct UTP attachment to workstations, servers, or routers using network interface cards (NICs).

Fiber cabling with SC connectors is required for direct attachment to Gigabit NICs or switches and routers.

Installing or Removing a GBIC

Some modules use Gigabit Interface Converters (GBICs) or miniature GBICs (mini-GBICs), which are individually insertable and removable port connectors. To insert or remove a GBIC or mini-GBIC, use the following procedures.

WARNING: All fiber-optic interfaces use Class 1 Lasers.

NOTE: See "Installation Precautions" on page 2-3 for other hardware installation precautions.

NOTE: The procedures for GBICs and mini-GBICs are different. Use the procedure that applies to the type of GBIC on your module.

Installing or Removing a Standard GBIC

To install a GBIC:

1. Put on an electrostatic discharge (ESD) wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Remove the GBIC from its protective packaging.
3. Gently insert the GBIC into the slot on the front panel of the module until the GBIC clicks into place. The GBICs are keyed to prevent incorrect insertion.
4. Remove the protective covering from the port connectors and store the covering for future use.
5. Insert the interface cable.

To remove a GBIC:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Disconnect the interface cable from the GBIC.
3. Insert the protective covering into the port connectors.
4. Squeeze and hold the tabs on each side of the GBIC, then gently pull the GBIC out of the module.
5. Store the GBIC in a safe, static-free place.

Installing or Removing a Mini-GBIC

To install a mini-GBIC:

1. Put on an electrostatic discharge (ESD) wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Remove the mini-GBIC from its protective packaging.
3. Gently insert the mini-GBIC into the slot on the front panel of the module until the mini-GBIC clicks into place. The mini-GBICs are keyed to prevent incorrect insertion. A tab on the bottom of the mini-GBIC locks the mini-GBIC to the front panel of the module.
4. Remove the protective covering from the port connectors and store the covering for future use.
5. Insert the interface cable.

To remove a mini-GBIC:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Disconnect the interface cable from the mini-GBIC.
3. Insert the protective covering into the port connectors.
4. Pull the sliding tab on the bottom of the mini-GBIC forward, away from the front panel of the module. Pulling this tab unlocks the mini-GBIC from the front panel.
5. Pull the mini-GBIC out of the module.
6. Store the mini-GBIC in a safe, static-free place.

Troubleshooting Network Connections

- For the indicated port, verify that both ends of the cabling (at the device and the connected device) are snug.
- Verify the connected device and device are both powered on and operating correctly.

- Verify that you have used the correct cable type for the connection:
 - For twisted-pair connections to an end node, use straight-through cabling.
 - For fiber-optic connections, verify that the transmit port on the device is connected to the receive port on the connected device, and that the receive port on device is connected to the transmit port on the connected device.
- Verify that the port has not been disabled through a configuration change. You can use the CLI. If you have configured an IP address on the device, you also can use the Web management interface or IronView Network Manager.
- If the other procedures don't resolve the problem, try using a different port or a different cable.

Testing Connectivity

After you install the network cables, you can test network connectivity to other devices by pinging those devices. You also can perform trace routes.

Pinging an IP Address

To verify that a Foundry device can reach another device through the network, enter a command such as the following at any level of the CLI on the Foundry device:

```
BigIron> ping 192.33.4.7
```

Syntax: ping <ip addr> | <hostname> [source <ip addr>] [count <num>] [timeout <msec>] [ttl <num>] [size <byte>] [quiet] [numeric] [no-fragment] [verify] [data <1-to-4 byte hex>] [brief]

See the *Foundry Switch and Router Command Line Interface Reference* for information about the parameters.

NOTE: If you address the ping to the IP broadcast address, the device lists the first four responses to the ping.

Tracing a Route

To determine the path through which a Foundry device can reach another device, enter a command such as the following at any level of the CLI on the Foundry device:

```
BigIron> traceroute 192.33.4.7
```

Syntax: traceroute <host-ip-addr> icmp | udp [dest-UDP-port] [maxttl <value>] [minttl <value>] [numeric] [timeout <value>] [source-ip <ip addr>]

NOTE: **icmp** and **udp** options are supported in releases 07.8.00 and later.

The CLI displays trace route information for each hop as soon as the information is received. Traceroute requests display all responses to a given TTL. In addition, if there are multiple equal-cost routes to the destination, the Foundry device displays up to three responses by default.

See the *Foundry Switch and Router Command Line Interface Reference* for information about the command syntax.

Managing the Device

You can manage a Foundry device using any of the following applications:

- Command Line Interface (CLI) – a text-based interface accessible through a direct serial connection or a Telnet session.
- Web management interface – A GUI-based management interface accessible through an HTTP (web browser) connection.
- IronView Network Manager – An optional SNMP-based standalone GUI application.

Logging on Through the CLI

Once an IP address is assigned to a Layer 2 Switch or ServerIron or to an interface on the Layer 3 Switch, you can access the CLI either through the direct serial connection to the device or through a local or remote Telnet session.

You can initiate a local Telnet or SNMP connection by attaching a straight-through RJ-45 cable to a port and specifying the assigned management station IP address.

The commands in the CLI are organized into the following levels:

- User EXEC – Lets you display information and perform basic tasks such as pings and traceroutes.
- Privileged EXEC – Lets you use the same commands as those at the User EXEC level plus configuration commands that do not require saving the changes to the system-config file.
- CONFIG – Lets you make configuration changes to the device. To save the changes across reboots, you need to save them to the system-config file. The CONFIG level contains sub-levels for individual ports, for VLANs, for routing protocols, and other configuration areas.

NOTE: By default, any user who can open a serial or Telnet connection to the Foundry device can access all these CLI levels. To secure access, you can configure Enable passwords or local user accounts, or you can configure the device to use a RADIUS or TACACS/TACACS+ server for authentication. See the *Foundry Security Guide*.

On-Line Help

To display a list of available commands or command options, enter “?” or press Tab. If you have not entered part of a command at the command prompt, all the commands supported at the current CLI level are listed. If you enter part of a command, then enter “?” or press Tab, the CLI lists the options you can enter at this point in the command string.

If you enter an invalid command followed by ?, a message appears indicating the command was unrecognized. For example:

```
BigIron(config)# router ip
Unrecognized command
```

Command Completion

The CLI supports command completion, so you do not need to enter the entire name of a command or option. As long as you enter enough characters of the command or option name to avoid ambiguity with other commands or options, the CLI understands what you are typing.

Scroll Control

By default, the CLI uses a page mode to paginate displays that are longer than the number of rows in your terminal emulation window. For example, if you display a list of all the commands at the global CONFIG level but your terminal emulation window does not have enough rows to display them all at once, the page mode stops the display and lists your choices for continuing the display.

Here is an example:

```
aaa
all-client
appletalk
arp
boot

some lines omitted for brevity..

ipx
lock-address
logging
mac
```

--More--, next page: Space, next line:
Return key, quit: Control-c

The software provides the following scrolling options:

- Press the Space bar to display the next page (one screen at time).
- Press the Return or Enter key to display the next line (one line at a time).
- Press Ctrl-C or Q to cancel the display.

Line Editing Commands

The CLI supports the following line editing commands. To enter a line-editing command, use the CTRL-key combination for the command by pressing and holding the CTRL key, then pressing the letter associated with the command.

Table 2.4: CLI Line Editing Commands

Ctrl-Key Combination	Description
Ctrl-A	Moves to the first character on the command line.
Ctrl-B	Moves the cursor back one character.
Ctrl-C	Escapes and terminates command prompts and ongoing tasks (such as lengthy displays), and displays a fresh command prompt.
Ctrl-D	Deletes the character at the cursor.
Ctrl-E	Moves to the end of the current command line.
Ctrl-F	Moves the cursor forward one character.
Ctrl-K	Deletes all characters from the cursor to the end of the command line.
Ctrl-L; Ctrl-R	Repeats the current command line on a new line.
Ctrl-N	Enters the next command line in the history buffer.
Ctrl-P	Enters the previous command line in the history buffer.
Ctrl-U; Ctrl-X	Deletes all characters from the cursor to the beginning of the command line.
Ctrl-W	Deletes the last word you typed.
Ctrl-Z	Moves from any CONFIG level of the CLI to the Privileged EXEC level; at the Privileged EXEC level, moves to the User EXEC level.

For a complete list of CLI commands and syntax information for each command, see the *Foundry Switch and Router Command Line Interface Reference*.

Searching and Filtering Output from CLI Commands

You can filter CLI output from **show** commands and at the --More-- prompt. You can search for individual characters, strings, or construct complex regular expressions to filter the output.

Searching and Filtering Output from show commands

You can filter output from **show** commands to display lines containing a specified string, lines that do not contain a specified string, or output starting with a line containing a specified string. The search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct

complex regular expressions. See “Using Special Characters in Regular Expressions” on page 2-34 for information on special characters used with regular expressions.

Displaying Lines Containing a Specified String

The following command filters the output of the **show interface** command for port 3/11 so it displays only lines containing the word “Internet”. This command can be used to display the IP address of the interface.

```
BigIron# show interface e 3/11 | include Internet
  Internet address is 192.168.1.11/24, MTU 1518 bytes, encapsulation ethernet
```

Syntax: <show-command> | include <regular-expression>

NOTE: The vertical bar (|) is part of the command.

Note that the regular expression specified as the search string is case sensitive. In the example above, a search string of “Internet” would match the line containing the IP address, but a search string of “internet” would not.

Displaying Lines That Do Not Contain a Specified String

The following command filters the output of the **show who** command so it displays only lines that do not contain the word “closed”. This command can be used to display open connections to the Foundry device.

```
BigIron# show who | exclude closed
Console connections:
  established
  you are connecting to this session
  2 seconds in idle
Telnet connections (inbound):
  1    established, client ip address 192.168.9.37
      27 seconds in idle
Telnet connection (outbound):
SSH connections:
```

Syntax: <show-command> | exclude <regular-expression>

Displaying Lines Starting with a Specified String

The following command filters the output of the **show who** command so it displays output starting with the first line that contains the word “SSH”. This command can be used to display information about SSH connections to the Foundry device.

```
BigIron# show who | begin SSH
SSH connections:
  1    established, client ip address 192.168.9.210
      7 seconds in idle
  2    closed
  3    closed
  4    closed
  5    closed
```

Syntax: <show-command> | begin <regular-expression>

Searching and Filtering Output at the --More-- Prompt

The --More-- prompt is displayed when output extends beyond a single page. From this prompt, you can press the Space bar to display the next page, the Return or Enter key to display the next line, or Ctrl-C or Q to cancel the display. In addition, you can search and filter output from this prompt. For example:

```
BigIron# ?
  append                Append one file to another
  appletalk-ping        Ping AppleTalk node
  atm                   ATM commands
  attrib                Change flash card file attribute
  boot                  Boot system from bootp/tftp server/flash image
  cd                    Change flash card working slot or current directory
  chdir                 Change flash card working slot or current directory
  clear                 Clear table/statistics/keys
  clock                 Set clock
  configure              Enter configuration mode
  copy                  Copy between flash, flash card, tftp, config/code
  debug                 Enable debugging functions (see also 'undebug')
  delete                Delete flash card files
  dir                   List flash card files
  disable               Disable a module before removing it
  enable                Enable a disabled module
  erase                 Erase image/configuration from flash
  exit                  Exit Privileged mode
  fastboot              Select fast-reload option
  format                Format flash card
  gignpa                Gigabit processor commands
  hd                    Display hex dump of flash card file
  kill                  Kill active CLI session
--More--, next page: Space, next line: Return key, quit: Control-c
```

At the --More-- prompt, you can press the forward slash key (/) and then enter a search string. The Foundry device displays output starting from the first line that contains the search string, similar to the **begin** option for **show** commands. For example:

```
--More--, next page: Space, next line: Return key, quit: Control-c
/telnet
```

The results of the search are displayed:

```
searching...
  telnet                Telnet by name or IP address
  temperature           temperature sensor commands
  terminal               display syslog
  traceroute            TraceRoute to IP node
  undebug                Disable debugging functions (see also 'debug')
  undelete              Undelete flash card files
  whois                 WHOIS lookup
  write                 Write running configuration to flash or terminal
```

To display lines containing only a specified search string (similar to the **include** option for **show** commands) press the plus sign key (+) at the --More-- prompt and then enter the search string.

```
--More--, next page: Space, next line: Return key, quit: Control-c
+telnet
```

The filtered results are displayed:

```
filtering...
telnet                Telnet by name or IP address
```

To display lines that do not contain a specified search string (similar to the **exclude** option for **show** commands) press the minus sign key (-) at the --More-- prompt and then enter the search string.

```
--More--, next page: Space, next line: Return key, quit: Control-c
-telnet
```

The filtered results are displayed:

```
filtering...
sync-standby         Synchronize active and standby module
temperature          temperature sensor commands
terminal             display syslog
traceroute           TraceRoute to IP node
undebg               Disable debugging functions (see also 'debug')
undetele             Undelete flash card files
whois                WHOIS lookup
write                Write running configuration to flash or terminal
```

As with the commands for filtering output from **show** commands, the search string is a regular expression consisting of a single character or string of characters. You can use special characters to construct complex regular expressions. See the next section for information on special characters used with regular expressions.

Using Special Characters in Regular Expressions

You use a regular expression to specify a single character or multiple characters as a search string. In addition, you can include special characters that influence the way the software matches the output against the search string. These special characters are listed in the following table.

Table 2.5: Special Characters for Regular Expressions

Character	Operation
.	The period matches on any single character, including a blank space. For example, the following regular expression matches "aaz", "abz", "acz", and so on, but not just "az": a.z
*	The asterisk matches on zero or more sequential instances of a pattern. For example, the following regular expression matches output that contains the string "abc", followed by zero or more Xs: abcX*
+	The plus sign matches on one or more sequential instances of a pattern. For example, the following regular expression matches output that contains "de", followed by a sequence of "g"s, such as "deg", "degg", "deggg", and so on: deg+

Table 2.5: Special Characters for Regular Expressions (Continued)

Character	Operation
?	<p>The question mark matches on zero occurrences or one occurrence of a pattern.</p> <p>For example, the following regular expression matches output that contains "dg" or "deg": de?g</p> <p>Note: Normally when you type a question mark, the CLI lists the commands or options at that CLI level that begin with the character or string you entered. However, if you enter Ctrl-V and then type a question mark, the question mark is inserted into the command line, allowing you to use it as part of a regular expression.</p>
^	<p>A caret (when not used within brackets) matches on the beginning of an input string.</p> <p>For example, the following regular expression matches output that begins with "deg": ^deg</p>
\$	<p>A dollar sign matches on the end of an input string.</p> <p>For example, the following regular expression matches output that ends with "deg": deg\$</p>
_	<p>An underscore matches on one or more of the following:</p> <ul style="list-style-type: none"> • , (comma) • { (left curly brace) • } (right curly brace) • ((left parenthesis) •) (right parenthesis) • The beginning of the input string • The end of the input string • A blank space <p>For example, the following regular expression matches on "100" but not on "1002", "2100", and so on. _100_</p>
[]	<p>Square brackets enclose a range of single-character patterns.</p> <p>For example, the following regular expression matches output that contains "1", "2", "3", "4", or "5": [1-5]</p> <p>You can use the following expression symbols within the brackets. These symbols are allowed only inside the brackets.</p> <ul style="list-style-type: none"> • ^ – The caret matches on any characters except the ones in the brackets. For example, the following regular expression matches output that does not contain "1", "2", "3", "4", or "5": [^1-5] • - The hyphen separates the beginning and ending of a range of characters. A match occurs if any of the characters within the range is present. See the example above.

Table 2.5: Special Characters for Regular Expressions (Continued)


Character	Operation
	A vertical bar separates two alternative values or sets of values. The output can match one or the other value. For example, the following regular expression matches output that contains either “abc” or “defg”: abc defg
()	Parentheses allow you to create complex expressions. For example, the following complex expression matches on “abc”, “abcabc”, or “defg”, but not on “abcdefgdefg”: ((abc+) ((defg)?

If you want to filter for a special character instead of using the special character as described in the table above, enter “\” (backslash) in front of the character. For example, to filter on output containing an asterisk, enter the asterisk portion of the regular expression as “*”.

```
BigIron# show ip route bgp | include \*
```

Logging On Through the Web Management Interface

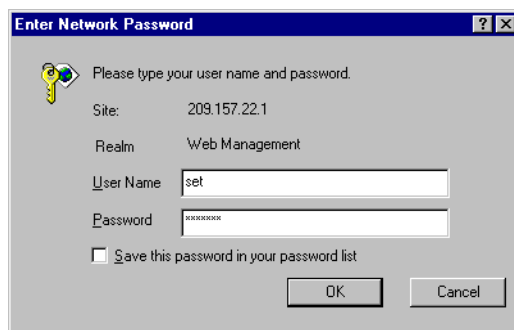
To use the Web management interface, open a web browser and enter the IP address of the Foundry device in Location or Address field. The web browser contacts the Foundry device and displays a picture of the device’s front panel dialog, as shown in Figure 2.13.



NOTE: If you are unable to connect with the device through a Web browser due to a proxy problem, it may be necessary to set your Web browser to direct Internet access instead of using a proxy. For information on how to change a proxy setting, refer to the on-line help provided with your Web browser.

To log in, click on the [Login](#) link. The following dialog is displayed.

Figure 2.13 Web management interface login dialog



By default, you can use the user name “get” and the default read-only password “public” for read-only access. However, for read-write access, you must enter “set” for the user name, and enter a read-write community string you have configured on the device for the password. There is no default read-write community string. You must add one using the CLI. See the *Foundry Security Guide*.

As an alternative to using the SNMP community strings to log in, you can configure the Foundry device to secure Web management access using local user accounts or Access Control Lists (ACLs). See the *Foundry Security Guide*.

Specifying a TCP Port for Web Management Interface

Beginning with software release 07.7.00, you can also specify the TCP port that will be used to access a device’s Web management interface by entering a command such as the following:

```
BigIron(config)# web-management tcp-port 168
```

Syntax: [no] web-management tcp-port <port-number>

The **tcp-port** <port-number> option specifies the port to be used to access the device’s Web management interface.

If IronView Network Manager is being used to manage the device, its Element Manager will query the device for the Web management port before it sends HTTP packets to the device.

Navigating the Web Management Interface

When you log into a device, the System configuration panel is displayed. This panel allows you to enable or disable major system features. You can return to this panel from any other panel by selecting the [Home](#) link.

The [Site Map](#) link gives you a view of all available options on a single screen.

The left pane of the Web management interface window contains a “tree view,” similar to the one found in Windows Explorer. Configuration options are grouped into folders in the tree view. These folders, when expanded, reveal additional options. To expand a folder, click on the plus sign to the left of the folder icon.

You can configure the appearance of the Web management interface by using one of the following methods.

USING THE CLI

Using the CLI, you can modify the appearance of the Web management interface with the **web-management** command.

To cause the Web management interface to display the List view by default:

```
BigIron(config)# web-management list-menu
```

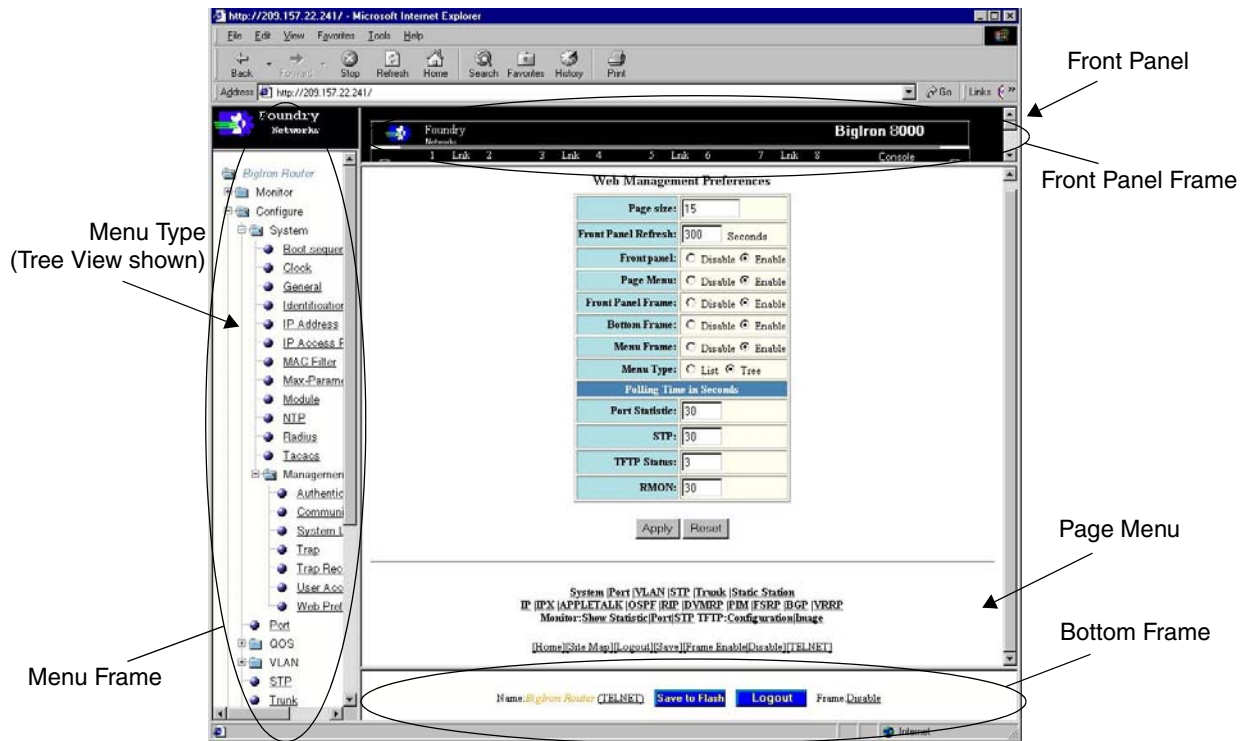
To disable the front panel frame:

```
BigIron(config)# no web-management front-panel
```

When you save the configuration with the **write memory** command, the changes will take place the next time you start the Web management interface, or if you are currently running the Web management interface, the changes will take place when you click the Refresh button on your browser.

USING THE WEB MANAGEMENT INTERFACE

1. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
2. Click on the plus sign next to System in the tree view to expand the list of system configuration links.
3. Click on the plus sign next to Management in the tree view to expand the list of system management links.
4. Click on the Web Preference link to display the Web Management Preferences panel.
5. Enable or disable elements on the Web management interface by clicking on the appropriate radio buttons on the panel. The following figure identifies the elements you can change.



NOTE: The tree view is available when you use the Web management interface with Netscape 4.0 or higher or Internet Explorer 4.0 or higher browsers. If you use the Web management interface with an older browser, the Web management interface displays the List view only, and the Web Management Preferences panel does not include an option to display the tree view.

6. When you have finished, click the Apply button on the panel, then click the Refresh button on your browser to activate the changes.
7. To save the configuration, click the plus sign next to the Command folder, then click the Save to Flash link.

NOTE: The only changes that become permanent are the settings to the Menu Type and the Front Panel Frame. Any other elements you enable or disable will go back to their default settings the next time you start the Web management interface.

Traps and Syslog Message for Web Management Interface

Beginning with software release 07.7.00, the Web management interface sends trap and Syslog messages when it is used to add, modify, or delete device configuration.

For example, if you click the Add button on the configuration panel below, the message “The change has been made.” appears at the top of the panel. Trap and Syslog messages are generated once a configuration change is made.

The change has been made.

Port VLAN

VLAN Id:	<input type="text" value="2"/>
Name:	<input type="text" value="admin"/>
QOS:	<input type="text" value="0"/>
Spanning Tree:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable

Port Members

<input type="checkbox"/> 1	<input type="checkbox"/> 2	<input checked="" type="checkbox"/> 3	<input type="checkbox"/> 4	<input type="checkbox"/> 5	<input type="checkbox"/> 6	<input type="checkbox"/> 7	<input type="checkbox"/> 8
<input type="checkbox"/> 9	<input type="checkbox"/> 10	<input type="checkbox"/> 11	<input type="checkbox"/> 12	<input checked="" type="checkbox"/> 13	<input type="checkbox"/> 14	<input type="checkbox"/> 15	<input type="checkbox"/> 16
<input type="checkbox"/> 17	<input type="checkbox"/> 18	<input type="checkbox"/> 19	<input type="checkbox"/> 20	<input type="checkbox"/> 21	<input checked="" type="checkbox"/> 22	<input type="checkbox"/> 23	<input type="checkbox"/> 24
<input type="checkbox"/> 25	<input type="checkbox"/> 26	<input type="checkbox"/> 27	<input type="checkbox"/> 28	<input type="checkbox"/> 29	<input type="checkbox"/> 30	<input checked="" type="checkbox"/> 31	<input type="checkbox"/> 32
<input type="checkbox"/> 33	<input type="checkbox"/> 34	<input type="checkbox"/> 35	<input type="checkbox"/> 36	<input type="checkbox"/> 37	<input type="checkbox"/> 38	<input type="checkbox"/> 39	<input type="checkbox"/> 40
<input type="checkbox"/> 41	<input type="checkbox"/> 42	<input type="checkbox"/> 43	<input type="checkbox"/> 44	<input type="checkbox"/> 45	<input type="checkbox"/> 46	<input type="checkbox"/> 47	<input type="checkbox"/> 48
<input type="checkbox"/> 49	<input type="checkbox"/> 50						

[Show][Protocol VLAN]

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

Logging on Through IronView Network Manager Network Manager

See the *Foundry IronView Network Management User's Guide* for information about using IronView Network Manager.

Creating an Alias for a CLI Command

In Enterprise software release 08.0.00 and later, you can create *aliases* for CLI commands. An alias serves as a shorthand version of a longer CLI command. For example, you can create an alias called **shoro** for the CLI command **show ip route**. Then when you enter **shoro** at the command prompt, the **show ip route** command is executed.

To create an alias called **shoro** for the CLI command **show ip route**, enter the following command:

```
BigIron(config)# alias shoro = show ip route
```

Syntax: [no] alias <alias-name> = <cli-command>

The <alias-name> must be a single word, without spaces. The <cli-command> can be up to <??> characters long, including spaces.

After the alias is configured, entering **shoro** at either the Privileged EXEC or CONFIG levels of the CLI, executes the **show ip route** command.

To create an alias called **wrsbc** for the CLI command **copy running-config tftp 10.10.10.10 test.cfg**, enter the following command:

```
BigIron(config)# alias wrsbc = copy running-config tftp 10.10.10.10 test.cfg
```

To remove the **wrsbc** alias from the Foundry device's configuration, enter one of the following commands:

```
BigIron(config)# no alias wrsbc
```

or

```
BigIron(config)# unalias wrsbc
```

Syntax: unalias <alias-name>

The specified <alias-name> must be the name of an alias already configured on the Foundry device.

To display the aliases currently configured on the Foundry device, enter the following command at either the Privileged EXEC or CONFIG levels of the CLI:

```
BigIron# alias
          wrsbc      copy running-config tftp 10.10.10.10 test.cfg
          shoro      show ip route
```

Syntax: alias

Notes

- You cannot include additional parameters with the alias at the command prompt. For example, after you create the **shoro** alias, **shoro bgp** would not be a valid command.
- If configured on the Foundry device, authentication, authorization, and accounting is performed on the actual command, not the alias entered for the command. If the actual command is available for execution at that level of the CLI and at the user's management privilege level, then the command is executed.
- To save an alias definition to the device's startup-config file, use the **write memory** command.

Swapping Modules (Chassis devices only)

Use the following procedures to swap out an old module and insert a new one.

Removing the Old Module

To remove a management module, pull the module out of the Chassis device.

To remove a forwarding module from a Chassis device, disable the module first before removing it from the Chassis device. Disabling the module before removing it prevents a brief service interruption on other forwarding modules. The brief interruption can be caused by the Chassis device reinitializing other modules in the chassis when you remove an enabled module.

NOTE: The **disable module** and **enable module** commands are not applicable to management modules. You do not need to disable a management module in software before removing it.

To disable a forwarding module, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# disable module 3
```

This command disables the module in slot 3.

Syntax: disable module <slot-num>

The <slot-num> parameter specifies the slot number.

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

NOTE: If you remove the module without first disabling it, the chassis re-initializes the other modules in the chassis, causing a brief interruption in service after which the chassis resumes normal operation.

If you decide after disabling a module that you do not want to remove the module, re-enable the module using the following command:

```
BigIron# enable module 3
```


Syntax: enable module <slot-num>

NOTE: You do not need to enable a module after inserting it in the chassis. The module is automatically enabled when you insert the module into a live chassis or when you power on the chassis.

NOTE: On all Chassis devices, if you plan to replace the removed module with a different type of module, you must configure the slot for the module. To configure a slot for a module, use the **module** command at the global CONFIG level of the CLI. See “Adding the New Module”.

Adding the New Module

After you physically insert a module into the Chassis device, you need to enter the location and type of module in the software if that slot was previously configured for a different module type.

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

NOTE: If the slot has never contained a module or you are swapping in exactly the same type of module, you do not need to use the **module** command. The slot requires configuration only if it has already been configured for another type of module.

USING THE CLI

To add a module to a Chassis device:

```
BigIron(config)# module 3 bi-8-port-gig-management-module
```

Syntax: module <slot-num> <module-type>

The <slot-num> parameter indicates the chassis slot number.

The <module-type> parameter specifies the platform, module type, and port configuration of the module.

NOTE: Module options that begin with “bi” and are for the Management 4 module also are applicable to the NetIron Internet Backbone router.

USING THE WEB MANAGEMENT INTERFACE

To configure a chassis slot for a module:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.

- Click on the [Module](#) link to display the Module panel, as shown in the following example.

Module

Slot	Module	Status	Ports	Starting MAC	
1	B8GM Fiber Management Module	OK	8	00e0.52f0.4f00	Delete
2	None				Delete
3	B24E Copper Switch Module	OK	24	00e0.52f0.4f40	Delete
4	B24E Copper Switch Module	OK	24	00e0.52f0.4f60	Delete
5	None				Delete
6	None				Delete
7	None				Delete
8	None				Delete

[\[Add Module\]](#)

[\[Home\]](#) [\[Site Map\]](#) [\[Logout\]](#) [\[Save\]](#) [\[Frame Enable\]](#) [\[Disable\]](#) [\[TELNET\]](#)

- Click the [Add Module](#) link to display the following panel.

Module

Slot:	1
Module Type:	bi-8-port-gig-management-module

[\[Show\]](#)

[\[Home\]](#) [\[Site Map\]](#) [\[Logout\]](#) [\[Save\]](#) [\[Frame Enable\]](#) [\[Disable\]](#) [\[TELNET\]](#)

- Select slot number from the Slot pulldown menu.
 - Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
 - Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
 - Slots in a 15-slot chassis are numbered 1 – 15, from left to right.
- Select the module type from the Module Type pulldown menu.
- Click the Add button to save the change to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Chapter 3

Using Redundant Management Modules

This chapter describes the redundant management modules and how to configure and manage them. Redundant management modules provide increased routing capacity and failover for BigIron, NetIron, and FastIron II .

See the following sections for information:

- “Configuring the Redundant Management Parameters” on page 3-3
- “File Synchronization Between the Active and Standby Redundant Management Modules” on page 3-11
- “Switching Over to the Standby Redundant Management Module” on page 3-16
- “PCMCIA Flash Card File Management Commands” on page 3-17 (applies only to the Management 4 module)
- “Using a 3Com Management Interface in the PCMCIA Slot” on page 3-33 (applies only to the Management 4 module)

NOTE: The NetIron Internet Backbone router requires a 512MB Management 4 module or higher, and is assembled at the factory with the appropriate model of Management 4 module. You cannot use a BigIron or FastIron II management module in a NetIron chassis.

The redundant management modules are fully-functional CPU management modules for Layer 3 Switches. You can use one or two redundant management modules in these devices.

You can use one or two redundant management modules in a Layer 3 Switch. Using two redundant management modules adds fault protection against system outage. The two modules work together as active and standby management modules. If the active module becomes unavailable, the standby module automatically takes over system operation.

NOTE: This chapter does not describe management features that are specific to the Velocity Management Module, such as logging on to individual CPUs. See “Using the Velocity Management Module” on page 4-1.

Configuration Considerations

- The Velocity Management Module and Management Modules 2, 3, and 4 support redundancy.
- You can use one or two redundant management modules in a Layer 3 Switch.
- You cannot use older management modules in the same Layer 3 Switch with redundant management modules.

Temperature Sensor

The redundant management modules contain a temperature sensor. You can use the CLI or Web management interface to display the active redundant management module's temperature and to change the warning and shutdown temperature levels. See "Using the Temperature Sensor" on page 9-52.

Switchover

When you power on or reload a Layer 3 Switch that contains two redundant management modules, the active redundant management module is selected based on the chassis slot previously specified by you or according to the lower slot number.

After the active module is selected, the active module loads its boot and flash code (boot and system software) and its system-config file and manages the system. The standby module also boots, using its own boot code but using the active module's flash code and system-config file. The standby module monitors the heartbeat of the active module. If the active module becomes unavailable, the standby module notices the absence of the heartbeat and assumes management control of the system.

NOTE: By default, the system does not use the boot code on the active module to boot the standby module. If you upgrade the boot code on the active module and the code contains a problem, you can still use the system by running the older boot code that is on the standby module. You can configure the standby to synchronize with the active module's boot code. See "File Synchronization Between the Active and Standby Redundant Management Modules" on page 3-11.

The standby module's system-config file is updated whenever the system-config file on the active module is updated. In addition, the running-config file on the standby module is updated at regular intervals to match the active module's running-config data. Thus, when a switchover occurs, the standby module also can reinstate the configuration data in the active module's running-config.

Following this switchover to the standby module, the standby module becomes the active module and continues to manage the system. When the other redundant management module (the one that used to be the active module) becomes available again or is replaced, that module becomes the standby module.

The active module also monitors the standby module. If the standby module becomes unavailable, the active module tries to reboot the standby module. You can display the status of each module using the CLI or the Web management interface, as described in "Determining Redundant Management Module Status" on page 3-8.

Management Sessions

You can establish management sessions only with the active redundant management module, not with the standby redundant management module. During switchover, all the CLI, Web management interface, and IronView Network Manager sessions open on the system are closed. To manage the system following a switchover, you must open a new management session. Although the system's MAC addresses change following switchover, the IP addresses do not. You can open new management sessions on the same IP addresses you were using before the switchover if desired.

To establish a serial connection to the CLI, you must move the serial cable to the serial port on the active redundant management module.

Syslog and SNMP Traps

When a switchover occurs, the software sends a Syslog message to the local Syslog buffer and also to the Syslog server, if you have configured the Foundry device to use one. In addition, if you have configured an SNMP trap receiver, the software sends an SNMP trap to the receiver.

When the system is powered on or otherwise reset normally, the software sends a cold start message and trap. However, if the system is reset as the result of switchover to the standby redundant management module, the software instead sends a switchover message and trap.

MAC Address Changes

The MAC addresses in the system are based on the MAC address of the active management module. During switchover, the system's MAC addresses change and the system sends out gratuitous ARP requests to flush the old MAC addresses from the ARP caches on attached IP devices, and update the caches with the Foundry device's new MAC addresses.

NOTE: The 15-slot chassis makes use of locally administered MAC addresses. If your site already uses locally administered MAC addresses of the Foundry OUI, which is 00e052, there could be a MAC address conflict with one of the ports on the Foundry device.

Configuring the Redundant Management Parameters

You can configure the following redundant management module parameters:

- Installation parameters:
 - Slot configuration. As with other module types, you must configure a chassis slot for the type of module you are installing in the slot.
 - Active redundant management module slot. By default, the redundant management module with the lower slot number is the active module.
- Operational parameters:
 - Boot code synchronization. By default, the standby redundant management module does not automatically synchronize to the boot code version installed on the active module. The standby module does automatically synchronize to the flash code (system software) on the active module.
 - Synchronization interval for running-config file
 - Warning and shutdown temperatures

Installing Redundant Management Modules

To install a redundant management module, perform the following tasks:

- Configure the chassis slot to receive the module.

NOTE: The system must be running a version of software that supports the module you want to install.

- Insert the module.
- Specify the default active module (if you do not want to use the system default, which is the redundant management module with the lower slot number).

In addition, if you use a TFTP or BootP server to boot the active module, you need to copy the flash code (system software) into the primary or secondary flash on the active redundant management module, then direct the active redundant management module to use the code to boot the standby module.

A standby redundant management module does not boot from a TFTP or BootP server.

NOTE: The slots in a 15-slot chassis are divided among 4 internal *regions*. Slots 1 – 4 belong to the same region; slots 5 – 8 belong to the same region; slots 9 – 12 belong to the same region, and slots 13 – 15 belong to the same region. If you are using redundant management modules, Foundry recommends that you place both management modules in slots belonging to the same region. For example, if you place one management module in slot 5, Foundry recommends that you place the other management module in slot 6, 7, or 8.

This note does not apply to 4-slot or 8-slot chassis.

Configuring the Chassis to Receive the Module

When you plan to insert a module into a chassis slot, you first must configure the slot to receive the module unless the slot already contains the same type of module.

USING THE CLI

To prepare slot 1 to receive an eight-port Gigabit redundant management module, enter the following commands at the global CONFIG level:

```
BigIron(config)# module 1 bi-8-port-gig-management-module
BigIron(config)# write memory
```

Syntax: module <slot-num> <module-type>

The <slot-num> parameter specifies the chassis slot to contain the module:

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

The <module-type> parameter specifies the platform and port configuration of the redundant management module.

NOTE: Module options that begin with “bi” and are for the Management 4 module also are applicable to the NetIron.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the [Module](#) link to display the Module panel, as shown in the following example.

Module

Slot	Module	Status	Ports	Starting MAC	
1	B8GMR Fiber Management Module	ACTIVE	8	00e0.5282.7aa0	Delete
2	None				Delete
3	B8GMR Fiber Management Module				Delete
4	None				Delete
5	None				Delete
6	B2P622 POS Module	OK	2	00e0.5282.7aa0	Delete
7	B24E Copper Switch Module	OK	24	00e0.5282.7aa0	Delete
8	None				Delete
Slot	Module	Status	Ports	Starting MAC	

[\[Add Module\]](#)

[\[Home\]](#)[\[Site Map\]](#)[\[Logout\]](#)[\[Save\]](#)[\[Frame Enable\]](#)[\[Disable\]](#)[\[TELNET\]](#)

- Click the [Add Module](#) link to display the following panel.

Module

Slot:	1
Module Type:	bi-8-port-gig-management-module

[Home](#) | [Site Map](#) | [Logout](#) | [Save](#) | [Frame Enable](#) | [Disable](#) | [TELNET](#)

- Select slot number from the Slot pulldown menu.
 - Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
 - Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
 - Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

NOTE: The slots in a 15-slot chassis are divided among 4 internal *regions*. Slots 1 – 4 belong to the same region; slots 5 – 8 belong to the same region; slots 9 – 12 belong to the same region, and slots 13 – 15 belong to the same region. If you are using redundant management modules, Foundry recommends that you place both management modules in slots belonging to the same region. For example, if you place one management module in slot 5, Foundry recommends that you place the other management module in slot 6, 7, or 8.

This note does not apply to 4-slot or 8-slot chassis.

- Select the module type from the Module Type pulldown menu.
- Click the Add button to save the change to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

The configuration change is saved to the active redundant management module's startup-config file. (The change is automatically sent to the standby module when the active module's system-config file is copied to the standby module.)

NOTE: You also can access the dialog for saving configuration changes by clicking on Command in the tree view, then clicking on [Save to Flash](#).

Specifying the Default Active Module

By default, the redundant management module in the lower slot number becomes the active redundant management module when you start the system. For example, if you install redundant management modules in slots 1 and 8 in a BigIron 8000 chassis, the default active module is the module in slot 1.

NOTE:

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
 - Slots in an 8-slot chassis are numbered 1 – 8, from top to bottom.
 - Slots in a 15-slot chassis are numbered 1 – 15, from left to right.
-

You can override the default and specify the active module.

NOTE: The change does not take effect until you reload the system. If you save the change to the active module's system-config file before reloading, the change persists across system reloads. Otherwise, the change affects only the next system reload.

USING THE CLI

To override the default and specify the active redundant management module, enter the following commands:

```
BigIron(config)# redundancy
BigIron(config-redundancy)# active-management 5
```

Syntax: active-management <slot-num>

The <slot-num> parameter specifies the chassis slot:

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

This command overrides the default and makes the redundant management module in slot 5 the active module following the next reload. The change affects only the next reload and does not remain in effect for future reloads.


To make the change permanent across future reloads, enter the **write memory** command to save the change to the startup-config file, as shown in the following example:

```
BigIron(config)# redundancy
BigIron(config-redundancy)# active-management 5
BigIron(config-redundancy)# write memory
```

NOTE: If you do not save the change to the startup-config file, the change affects only the next reload.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Select the Redundant link to display the following panel.



3. Select slot number for the active redundant management module from the Active Management Slot pulldown menu. If you use the default value, Auto Select, the Layer 3 Switch uses the redundant management module in the lower slot number.

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.

- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
 - Slots in a 15-slot chassis are numbered 1 – 15, from left to right.
4. Click the Apply button to send the configuration change to the active module's running-config file.
 5. If you want the change to remain in effect following the next system reload, select the [Save](#) link to save the configuration change to the active redundant management module's startup-config file. (The change is automatically sent to the standby module when the active module's system-config file is copied to the standby module.)

NOTE: If you do not save the change to the startup-config file, the change affects only the next reload.

NOTE: The other options on this panel are described in later sections.

Inserting the Module

You can remove and insert modules when the system is powered on. Make sure you adhere to the cautions noted in "Installation Precautions" on page 2-3.

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Remove the module or faceplate from the slot:
3. If you are replacing another module, loosen the two screws on the module you are removing.
 - Pull the card ejectors towards you, away from the module front panel. The card will unseat from the backplane.
 - Pull the module out of the chassis and place in an anti-static bag for storage.
4. If you are installing a redundant management module in an unoccupied module slot, remove the blank faceplate from the slot in which the module is to be installed. Place the blank faceplate in a safe place for future use.
5. Remove the redundant management module from its packaging.
6. Insert the module into the chassis slot and glide the card along the card guide until the card ejectors on the front of the module touch the chassis.
 - Modules for 4-slot chassis slide in horizontally with the module label on the left.
 - Modules for 8-slot chassis slide in vertically with the module label at the top.
 - Modules for 15-slot chassis slide in vertically with the module label at the top.
7. Push the ejectors toward the center of the module until they are flush with the front panel of the module. The module will be fully seated in the backplane.
8. Tighten the two screws at either end of the module.
9. If you do not use one or more of the slots, make sure that a slot faceplate is still attached over each unused slot for safe operation and proper system cooling.

Installing and Removing GBICs

The Management 4 models that have ports use GBIC modules for the ports.

WARNING: All fiber-optic interfaces use Class 1 Lasers.

To install a GBIC:

1. Put on an electrostatic discharge (ESD) wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

2. Remove the GBIC from its protective packaging.
3. Gently insert the GBIC into the slot on the front panel of the module until the GBIC clicks into place. The GBICs are keyed to prevent incorrect insertion.
4. Remove the protective covering from the port connectors and store the covering for future use.
5. Insert the interface cable.

To remove a GBIC:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Disconnect the interface cable from the GBIC.
3. Insert the protective covering into the port connectors.
4. Squeeze and hold the tabs on each side of the GBIC, then gently pull the GBIC out of the module.
5. Store the GBIC in a safe, static-free place.

Determining Redundant Management Module Status

You can determine the status of a redundant management module in the following ways:

- Status LED – The redundant management module has two green LEDs on the right side of the CLI serial port. The lower LED shows the management status.
- Module information in software – The module information displayed by the software indicates whether the module is the active module, the standby module, or has another status.

Status LED

If you are located near the device, you can determine which redundant management module is currently the active module and which one is the standby by observing the upper green LED to the right of the serial management port. If the upper green LED is lit, the module is currently the active redundant management module. If the LED is dark, the module is the standby. The lower green LED indicates the power status. If the lower LED is dark, the module is not receiving power. (A module without power will not function as the active or standby module.)

Software

You can display status information for the modules using either of the following methods.

NOTE:

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
 - Slots in an 8-slot chassis are numbered 1 – 8, from top to bottom.
 - Slots in a 15-slot chassis are numbered 1 – 15, from left to right.
-

USING THE CLI

To display the status of a redundant management module using the CLI, enter the following command at any CLI level:

```
BigIron> show module
```

Module	Status	Ports	Starting MAC
S1: B8GMR Fiber Management Module	ACTIVE	8	00e0.5202.a2d4
S2: B24E Copper Switch Module	OK	24	00e0.5202.a2d4
S3: B24E Copper Switch Module	OK	24	00e0.5202.a2d4
S4: B24E Copper Switch Module	OK	24	00e0.5202.a2d4
S5: B8GMR Fiber Management Module	STANDBY	8	00e0.5202.a334
S6: B24E Copper Switch Module	OK	24	00e0.5202.a2d4
S7: B24E Copper Switch Module	OK	24	00e0.5202.a2d4
S8: B24E Copper Switch Module	OK	24	00e0.5202.a2d4

Syntax: show module

NOTE: The module descriptions do not distinguish between SX and LX ports.

The Status column shows the module status. The redundant management modules can have one of the following statuses:

- **ACTIVE** – The module is currently the active management module.
- **STANDBY** – The module is the standby management module.
- **COMING UP** – The module is coming up as the standby module. This status can be observed during switchover.

The statuses above apply only to management modules. The following statuses apply only to host modules:

- **FAILED** – This status applies only to host modules, not to management modules. This status indicates that the host module failed to come up.
- **OK** – This status applies only to host modules, not to management modules. This status indicates that the module came up and is operating normally.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.

2. Click on the [Module](#) link to display the Module panel, as shown in the following example.

Module

Slot	Module	Status	Ports	Starting MAC	
1	B8GMR Fiber Management Module	ACTIVE	8	00e0.5282.7a00	Delete
2	None				Delete
3	B8GMR Fiber Management Module				Delete
4	None				Delete
5	None				Delete
6	B2P622 POS Module	OK	2	00e0.5282.7aa0	Delete
7	B24E Copper Switch Module	OK	24	00e0.5282.7a00	Delete
8	None				Delete
Slot	Module	Status	Ports	Starting MAC	

[\[Add Module\]](#)

[\[Home\]](#) [\[Site Map\]](#) [\[Logout\]](#) [\[Save\]](#) [\[Frame Enable\]](#) [\[Disable\]](#) [\[TELNET\]](#)

The Status column shows the module status. The redundant management modules can have one of the following statuses:

- ACTIVE – The module is currently the active management module.
- STANDBY – The module is the standby management module.

The statuses above apply only to management modules. The following statuses apply only to host modules:

- FAILED – This status applies only to host modules, not to management modules. This status indicates that the host module failed to come up.
- OK – This status applies only to host modules, not to management modules. This status indicates that the module came up and is operating normally.

Displaying Switchover Messages

You can determine whether a switchover has occurred by viewing the system log or the traps logged on an SNMP trap receiver.

USING THE CLI

To view the system log, enter the following command at any level of the CLI:

```
BigIron> show log

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
Buffer logging: level ACDMEINW, 8 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
I=informational N=notification W=warning

Static Log Buffer:

Dynamic Log Buffer (50 entries):

at 0 days 0 hours 0 minutes 0 seconds, level alert
Management module at slot 1 state changed,
changed state from standby to active
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Monitor in the tree view to display the Monitor options.
3. Select the System Log link to display the system log.

File Synchronization Between the Active and Standby Redundant Management Modules

Each redundant management module contains four files that can be synchronized between the two modules:

- **Boot code** – The code the module runs when it first starts up. By default, the boot code is not synchronized between redundant management modules. This ensures that the system can still operate if a new version of boot code contains a bug that prohibits normal operation. If the new code on the active module does not work properly, the system can still run using the older version of boot code on the standby module.

You can configure the standby redundant management module to synchronize with the active redundant management module's boot code whenever the boot code on the active module is updated or the system starts up.

- **Flash code (system software)** – The flash code is automatically synchronized between the redundant management modules. When the system starts up, the active redundant management module sends its flash code to the standby redundant management module to boot the module.

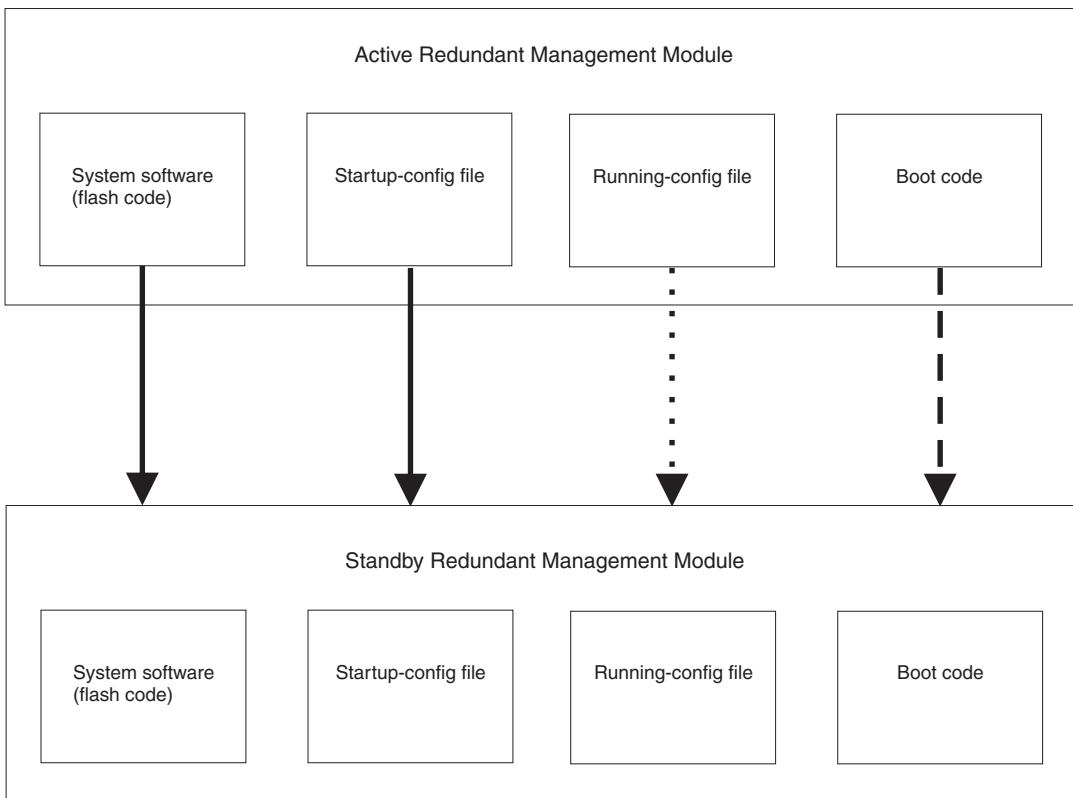

NOTE: The flash code on VM1 VSP CPUs (non-management CPUs) is not automatically synchronized. To synchronize the flash code on the VSP CPUs, use the **vm copy tftp flash** command, described in "Immediately Synchronizing Software" on page 3-14. The flash code on the VM CPU is automatically synchronized.

- **System-config file** – The system-config file is automatically copied from the active redundant management module to the standby redundant management module when the system starts up. The file is also copied to the standby module whenever you save changes to the file. If switchover occurs, the standby redundant management module loads system parameters from the running-config data that was last received from the active redundant management module. If the standby module did not receive running-config data from the active module, the standby module uses configuration information in the system-config file copied from the active module.

- **Running-config** – The running-config is automatically copied from the active redundant management module to the standby redundant management module at regular intervals. The default interval is 10 seconds. You can change the interval to 4 – 20 seconds. If you set the interval to 0, the configuration data is not copied to the standby redundant management module. As described above, if switchover occurs, the standby redundant management module loads system parameters from the running-config that was last received from the active redundant management module.

Figure 3.1 shows how the files are synchronized between the active redundant management module and the standby redundant management module.

Figure 3.1 Redundant management module file synchronization



Displaying the Synchronization Settings

You can independently synchronize the following types of software between the active and standby modules:

- boot code
- flash code (system software)
- startup-config file
- running-config

When you synchronize software between the modules, the active module copies its software to the standby module.

To display the current file synchronization settings, enter the following command:

```
BigIron# sync-standby

Sync code image: TRUE
Sync config data: TRUE
Sync boot image: FALSE
Running-config sync interval is 10 seconds
```

NOTE: The values shown in this example are the default values.

Syntax: sync-standby

NOTE: The **sync-standby** command has optional parameters. If you enter one of the parameters, the CLI synchronizes software between the modules. To display the synchronization settings instead of synchronizing software, enter the command without parameters.

This display shows the following information.

Table 3.1: CLI Display of Synchronization Settings

This Field...	Displays...
Sync code image	Indicates whether the active module is configured to automatically synchronize its flash code with the standby module. The value can be one of the following: <ul style="list-style-type: none"> FALSE – The code is not automatically synchronized. TRUE – The code is automatically synchronized.
Sync config data	Indicates whether the active module is configured to automatically synchronize its startup-config file with the standby module. The value can be one of the following: <ul style="list-style-type: none"> FALSE – The startup-config file is not automatically synchronized. TRUE – The startup-config file is automatically synchronized.
Sync boot image	Indicates whether the active module is configured to automatically synchronize its boot code with the standby module. The value can be one of the following: <ul style="list-style-type: none"> FALSE – The boot code is not automatically synchronized. TRUE – The boot code is automatically synchronized.
Running-config sync interval	Indicates whether the active module is configured to automatically synchronize its running-config with the standby module. The value can be one of the following: <ul style="list-style-type: none"> FALSE – The running-config is not automatically synchronized. TRUE – The running-config is automatically synchronized.

Immediately Synchronizing Software

You can immediately synchronize software between the active and standby management modules. When you synchronize software, the active module copies the software you specify to the standby module, replacing the software on the standby module.

To synchronize software, use either of the following methods.

USING THE CLI

To immediately synchronize the boot code on the standby module with the boot code on the active module, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# sync-standby boot
```

Syntax: sync-standby boot

To immediately synchronize the flash code (system software) on the standby module with the boot code on the active module, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# sync-standby code
```

Syntax: sync-standby code

NOTE: The **sync-standby code** command does not synchronize the VSP CPUs (non-management CPUs) on the VM1. To synchronize the VSP CPUs, use the following command:

vm copy tftp flash <tftp-server-ip-addr> <image-file-name> **primary | secondary**

This command upgrades the VSP CPU flash code on all VSP CPUs on both VM1 modules in the chassis.

To immediately synchronize the running-config on the standby module with the running-config on the active module, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# sync-standby running-config
```

Syntax: sync-standby running-config

To immediately synchronize the startup-config file on the standby module with the startup-config file on the active module, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# sync-standby startup-config
```

Syntax: sync-standby startup-config

USING THE WEB MANAGEMENT INTERFACE

NOTE: This procedure applies only to synchronizing the boot code and the running-config. To immediately synchronize the flash code or the startup-config file, use the CLI procedure above.

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.

- Select the [Redundant](#) link to display the following panel.

Redundant Management Modules

Active Management Slot:	Auto Select
Standby Management Module Synchronization	
Running Configuration Interval (sec):	10
	Synchronize Configuration Now
Boot Flash:	<input type="checkbox"/>
	Synchronize Boot Flash Now

[Switch-over Active Module]

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

- Click the button for the code or file you want to immediately synchronize:
 - To synchronize the running-config, select the Synchronize Configuration Now button.
 - To synchronize the boot flash code, select the Synchronize Boot Flash Now button.

As soon as you click the button, the Web management interface immediately performs the synchronization.

Automating Synchronization of Software

Automatic synchronization of the flash code, running-config, and system-config file is enabled by default. Automatic synchronization of the boot code is disabled by default.

To change the automatic synchronization setting, use one of the following methods.

USING THE CLI

The CLI commands for automating synchronization of software between the active and standby modules is the same as the syntax for immediately synchronizing the software. The only difference is the CLI level where you enter the commands.

- To immediately synchronize software, enter the command at the Privileged EXEC level.
- To automate synchronization starting with the next software reload or system reset and each reload or reset after that, enter the command at the Redundancy CONFIG level.

Automatic synchronization of the flash code, running-config, and system-config file is enabled by default. Automatic synchronization of the boot code is disabled by default. To change the automatic synchronization setting, use one of the following commands:

Syntax: [no] sync-standby boot

Syntax: [no] sync-standby code

Syntax: [no] sync-standby startup-config

Syntax: [no] sync-standby running-config [<num>]

To disable automatic synchronization of the boot code, flash code, or startup-config file, enter “no” in front of the command.

The <num> parameter with the **sync-standby running-config** command specifies the synchronization interval. You can specify from 4 – 20 seconds. The default is 10 seconds.

To disable automatic synchronization of the running-config, set the synchronization interval (the <num> parameter) to 0.

[USING THE WEB MANAGEMENT INTERFACE](#)

NOTE: This procedure applies only to synchronization of the boot code and running-config. To change automatic synchronization of other software, use the CLI procedure above.

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Select the [Redundant](#) link to display the following panel.

Redundant Management Modules

Active Management Slot:	Auto Select ▾
Standby Management Module Synchronization	
Running Configuration Interval (sec):	10 Synchronize Configuration Now
Boot Flash:	<input type="checkbox"/> Synchronize Boot Flash Now

[Switch-over Active Module]

[Home](#)
[Site Map](#)
[Logout](#)
[Save](#)
[Frame Enable](#)
[Disable](#)
[TELNET](#)

3. To enable automatic synchronization of the boot code, select the checkbox next to Boot Flash.
4. To change the synchronization interval for the running-config, enter the new value in the Running Configuration Interval field. To disable automatic synchronization of the running-config, enter 0 in the field.
5. Select the checkbox next to Boot Flash.

NOTE: Do not click the Synchronize Boot Flash Now button unless you want the active module to immediately copy its boot flash image to the standby module.

6. Click the Apply button to send the configuration change to the active module's running-config file.
7. If you want the change to remain in effect following the next system reload, select the [Save](#) link to save the configuration change to the active redundant management module's startup-config file. (The change is automatically sent to the standby module when the active module's system-config file is copied to the standby module.)

Switching Over to the Standby Redundant Management Module

If you reload the software using the **reload** command, the behavior of the management modules is the same as when you power the system on. The system selects the active module based on the slot you specified or based on the lower slot number if you did not specify a slot. Then both redundant management modules load their own boot code and load the active redundant management module's flash code (system software) and system-config file.

If you do not want to reload the system but you instead want to force the system to switch over to the standby module (and thus make it the active redundant management module), use one of the following methods.

[USING THE CLI](#)

To switch over to the other redundant management module, enter a command such as the following:

```
BigIron# reset 2
```

Syntax: reset <slot-num>

Specify the slot number containing the currently active management module. Do not specify the slot number containing the standby module to which you want to switch over.

The <slot-num> parameter specifies the chassis slot:

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Select the [Redundant](#) link to display the following panel.

Redundant Management Modules

Active Management Slot:	Auto Select ▾
Standby Management Module Synchronization	
Running Configuration Interval (sec):	10 Synchronize Configuration Now
Boot Flash:	<input type="checkbox"/> Synchronize Boot Flash Now

Apply Reset

[\[Switch-over Active Module\]](#)

[\[Home\]](#) [\[Site Map\]](#) [\[Logout\]](#) [\[Save\]](#) [\[Frame Enable\]](#) [\[Disable\]](#) [\[TELNET\]](#)

3. Select the [Switch-over Active Module](#) link. A message appears asking you to verify that you want to switch over from the active module to the standby.
4. Select Yes to switch over or No to cancel the switchover request.
5. Click the Add button to save the change to the device's running-config file.
6. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

PCMCIA Flash Card File Management Commands

This section describes the commands for managing files on the Management 4 module's PCMCIA flash card.

NOTE: These commands apply only to Chassis devices using Management 4 modules.

The Management 4 module's PCMCIA slot supports 8.3 file names (up to eight characters in the name and up to three additional characters in the extension). File name are not case sensitive. Thus, the software considers the name "test.cfg" and "TEST.CFG" to be the same.

You can use the file management commands to perform the following tasks:

- Format a flash card
- Determine the flash card that currently has the management focus.
- Switch the management focus from one flash card or subdirectory to another.
- Display a directory of the files on a flash card.
- Display the contents of a file.

- Display a hexadecimal dump of the data in a file.
- Create a subdirectory.
- Remove a subdirectory.
- Rename a file.
- Change the read-write attribute of a file.
- Delete a file from a flash card.
- Undelete a file.
- Append one file to another file (join two files).
- Perform the following copy operations:
 - Copy files from one flash card to the other.
 - Copy files between a flash card and the device's flash memory.
 - Copy files between a flash card and a TFTP server.
 - Copy a startup-config file between a flash card and the device's flash memory.
 - Copy the running-config to a flash card
 - Load a running-config from a flash card
 - Copy a POS image file from a flash card to a POS module's flash memory
- Load the startup-config file from a flash card during system startup

In the CLI, all the file management commands are at the Privileged EXEC level of the CLI.

CAUTION: Do not add or remove a flash card while a file operation involving the flash card's slot is in progress. Doing so can result in corruption of the flash card. If this occurs, you may need to reformat the flash card to make it usable again. Reformatting the card erases all data stored on the card.

PCMCIA Slots

The Management 4 module has two PCMCIA slots, numbered 1 and 2.

- In a 4-slot chassis, slot 1 is on top and slot 2 is on the bottom.
- In an 8-slot chassis, slot 1 is on the right and slot 2 is on the left.
- In a 15-slot chassis, slot 1 is on the right and slot 2 is on the left.

Some flash file management operations require that you specify the flash card slot(s) or subdirectory path(s) involved in the operation. For example, when you copy a file to or from a flash card, you must specify the flash card you are copying to or from and the subdirectory path for the file location. However, some other commands do not require a flash card slot or subdirectory path as a parameter. Instead, these commands assume that you want to perform the operation on the flash card slot and path that currently has the management focus.

The **management focus** determines the default flash card and subdirectory path (if applicable) for a file management operation. For example, when you list a directory of the files on a flash card, the PCMCIA slot and subdirectory path parameters are optional. If you do not specify the slot or subdirectory path, the software displays the contents of the subdirectory that currently has the management focus on the card in the slot that currently has the management focus. As another example, the command for deleting a file from a flash card does not require that you specify the PCMCIA slot or subdirectory path. If you do not specify the slot or subdirectory, the command deletes the file from the flash card and subdirectory that have the management focus.

When you power on or reload a device, if the management module contains only one flash card, the slot that contains the flash card receives the management focus by default. If both slots contain flash cards, slot 1 receives the management focus by default.

To determine the slot and subdirectory that currently have the management focus, enter the **pwd** command. (See “Determining the Flash Card Slot and Subdirectory Path That Currently Have the Management Focus” on page 3-21.) To change management focus to the other slot or subdirectory, enter the **cd...** or **chdir...** command. (See “Switching the Management Focus” on page 3-21.)

Commands that accept a slot number as a parameter can accept either of the following values:

- **slot1** – indicates slot 1
- **slot2** – indicates slot 2

The CLI provides commands to switch the focus from one PCMCIA slot to the other and to determine the slot that currently has the focus.

Subdirectories

The software supports subdirectories on the flash card.

Software release 07.1.00 enhances the PCMCIA flash card support by enabling you to create and navigate among DOS-like subdirectories on a flash card.

You can create up to 512 subdirectories from the root directory level. The actual number depends on how you name the subdirectories. If you use names that are not more than eight characters long for each directory, then you can create 512. If you use longer names, the name information occupies more space in the file system and you therefore can create fewer than 512 subdirectories.

The limit of 512 or fewer subdirectories applies only to ones you create at the root directory level. The number of subdirectories you can create at another subdirectory level is limited only on the amount of space on the flash card.

The full path name for a file's location can be a maximum of 260 characters. You can nest subdirectories as deep as you want so long as the full path name is 260 characters or less.

When you include a subdirectory path in a file management command, use a backslash between each level. For example, to create a subdirectory for flash code and copy a flash image file to the subdirectory, enter the following commands:

```
BigIron# mkdir slot1 \RouterCode\07100
BigIron# ncopy tftp 10.10.10.1 B2R07100.bin slot1 \RouterCode\07100\
```

These commands create two levels of subdirectories on the flash card in PCMCIA slot 1, then copy a flash image file named “B2R07100.bin” from a TFTP server into the new 07100 subdirectory. Since the file name for the copy destination is not specified, the software uses the same name for the copy (B2R07100.bin).

File and Subdirectory Naming Conventions

Files and subdirectory names can be up to 32 characters long. The following characters are valid in file and subdirectory names:

- All upper and lowercase letters
- All digits
- Spaces
- Any of the following special characters:
 - \$
 - %
 - ' (single quote)
 - - (hyphen)
 - _ (underscore)
 - @

- ~
- `
- !
- (
-)
- {
- }
- ^
- #
- &

You can use spaces in a file or subdirectory name if you enclose the name in double quotes. For example, to specify a subdirectory name that contains spaces, enter a string such as the following: "a long subdirectory name".

A subdirectory or file name can be a maximum of 256 characters long. A complete subdirectory path name cannot contain more than 260 characters.

There is no maximum file length. A file can be as large as the available flash card space.

Wildcards

Commands to display a directory of files, to change the read-write attribute of a file, or to delete files accept wildcards in the file name (<file-name>). When using these commands, you can use "*" (asterisk) as a wildcard for any part of the name. For example, all the following values are valid for <file-name>:

- startup.cfg
- start*.cfg
- B2P07000.bin
- *.bin
- B2P07000.*
- B2P*.bin
- B2P*.*

Formatting a Flash Card

The flash cards shipped with Management 4 modules are already formatted for the 16 FAT file system used by the modules. If you want to use a flash card that is not formatted for the 16 FAT file system, you need to reformat the flash card before you can store files on the card.

CAUTION: Make sure the flash card is empty or does not contain files you want to keep. Formatting a flash card completely erases all files on the card.

CAUTION: Once you start the formatting process, you cannot stop it. Even if you enter CTRL-C to stop the CLI output and a new prompt appears, the formatting continues. Make sure you want to format the card before you enter the command.

NOTE: When you reformat a flash card in the Management 4 module, the formatting takes about ten minutes. This is because the software checks every sector on the card, marks bad sectors if found, and skips over the bad sectors so that the bad sectors do not interfere with use of the card. If you do not want to use the Management 4 module to reformat the card, you can use a PC with a flash card drive instead.

USING THE CLI

To reformat a flash card, enter the following command:

```
BigIron# format slot2
Formatting Flash Card(256 clusters per dot) .....
.....
.....
Verifying Flash Card(256 clusters per dot) .....
.....
.....
80809984 bytes total card space.
80809984 bytes available on card.

    2048 bytes in each allocation unit.
    39458 allocation units available on card.
```

Flash card format done

As shown in this example, the software formats the sector on the flash card, then verifies the formatting. In this example, the software did not find any bad sectors, so all the bytes on the card are available.

Syntax: format slot1 | slot2 [<label>]

The **slot1 | slot2** parameter specifies the PCMCIA slot that contains the flash card you are formatting.

The <label> parameter specifies the label. You can specify up to 11 alphanumeric characters. You cannot use special characters or spaces.

Determining the Flash Card Slot and Subdirectory Path That Currently Have the Management Focus

If you are not sure which flash card slot and subdirectory path have the focus, use the following method to display this information.

USING THE CLI

To display which flash card slot and subdirectory path currently have the management focus, enter the following command:

```
BigIron# pwd
slot1 \
```

In this example, the management focus is at the root directory of the flash card in slot 1.

Syntax: pwd

In the following example, the management focus is at a subdirectory called "TEST" on the flash card in slot 1.

```
BigIron# pwd
slot1 \TEST
```

Switching the Management Focus

The effect of file management commands depends on the flash card and subdirectory that have the management focus. For example, if you enter a command to delete a file, the software deletes the specified file from the flash card and subdirectory that currently have the management focus.

To switch the focus of the CLI from one flash card to the other, enter a command such as the following:

```
BigIron# cd slot2
BigIron#
```

Syntax: cd | chdir slot1 | slot2

Syntax: cd | chdir <dir-name>

When you enter the **cd** command, the software changes the management focus to the slot or subdirectory path you specify, then displays a new command prompt.

If a slot you specify does not contain a flash card, the software displays the message shown in the following example.

```
BigIron# cd slot2
The system can not find the drive specified
```

To switch the management focus to a different subdirectory, enter a commands such as the following:

```
BigIron# cd PLOOK
Current directory of slot1 is: \PLOOK
```

This command changes the focus from the root directory level (\) to the subdirectory named “PLOOK”.

If you specify an invalid subdirectory path, the CLI displays a message such as the following:

```
BigIron# cd PLOOK
Path not found
```

If you are certain the path you specified exists, make sure you are at the correct level for reaching the path. For example, if you are already at the PLOOK level, the CLI cannot find the subdirectory “\PLOOK” because it is not a subdirectory from the level that currently has the management focus.

Displaying a Directory of the Files on a Flash Card

Use the following method to list the files on a flash card.

NOTE: By default, the software displays the contents of the flash card in the slot that has the management focus. However, you do not need to change the focus to list the files on another flash card. You can specify the other flash card when you display the files.

USING THE CLI

To display a directory of all the files on the flash card that has the management focus, enter the following command:

```
BigIron# dir
Volume in slot1 has no label
Volume Serial Number is 19ED-1725

Directory of slot1

01/01/2000  00:00a    685935      POS.BIN
01/01/2000  00:00a    2157693     M4R.BIN
01/01/2000  00:00a         184      A22.CFG
01/01/2000  00:00a         254      R CFG.CFG
01/01/2000  00:00a         256      STR.CFG
01/01/2000  00:00a    1027230     M5.BIN
01/01/2000  00:00a         184      A8.CFG
01/01/2000  00:00a    1029838     M4S.BIN
01/01/2000  00:00a     687026     P3R.BIN
01/01/2000  00:00a    1029838     MM.BIN
          10 File(s)          6618438 bytes
                               74180608 bytes free
```


Syntax: dir [slot1 | slot2] [<file-name>]

The following information is displayed for each file.

Table 3.2: CLI Display of Flash Card File Information

This Field...	Displays...
File date	The date on which the file was placed on the flash module, if the Foundry device's system clock is set.
Time of day	The time of day at which the file was placed on the flash module, if the Foundry device's system clock is set. If the clock is not set, the field shows 00:00a (12 AM), as shown in the example above.
File size	The number of bytes in the file.
Read-write attribute	If you have set the file's read-write attribute to read-only, "R" appears before the file name. If the file's read-write attribute is read-write (the default), no value appears in this column. For information, see "Changing the Read-Write Attribute of a File" on page 3-26.
File name	The file name.
Long file name	The longer file name if the file was created on a PC and the name is longer than the 8.3 format.

The directory also lists the total number of files that match the value for the name you specified, the total number of bytes used by all the files, and the number of bytes still free on the card.

To list only files that contain a specific pattern of characters in the name, enter a command such as the following:

```
BigIron# dir *.bin
Volume in slot1 has no label
Volume Serial Number is 19ED-1725

Directory of slot1

01/01/2000  00:00a    685935      POS.BIN
01/01/2000  00:00a    2157693     M4R.BIN
01/01/2000  00:00a    1027230     M5.BIN
01/01/2000  00:00a    1029838     M4S.BIN
01/01/2000  00:00a    687026      P3R.BIN
01/01/2000  00:00a    1029838     MM.BIN
           6 File(s)          6617560 bytes
                          74180608 bytes free
```

The command in this example lists all the image files on the flash card in the slot that has the management focus. (More specifically, the command lists all the files that end with ".bin".)

Displaying the Contents of a File

Use the following method to display the data in a file on a flash card.

USING THE CLI

To display the contents of a file, enter a command such as the following:

```
BigIron# more cfg.cfg
ver 06.5.00T51
module 1 bi-4-port-gig-m4-management-module
module 2 bi-24-port-copper-module
module 3 bi-4-port-gig-m4-management-module
```

```

module 4 bi-4-port-gig-m4-management-module
!
!
!
!
m2 active-management 3
ip address 192.168.2.58 255.255.255.0
end

```

Syntax: more [slot1 | slot2] <file-name>

This example shows the contents of a simple configuration file.

NOTE: The syntax for the m2 active-management command is changed to active-management. This example is from a software release before the change.

Display a Hexadecimal Dump of the Data in a File

Use the following method to display the data in a file in hexadecimal format.

USING THE CLI

To display the data in a file in hexadecimal format, enter a command such as the following:

```

BigIron# hd cfg.cfg
00000000: 76657220 30362e35 2e303054 35310a6d      ver 06.5.00T51 m
00000010: 6f64756c 65203120 62692d34 2d706f72      odule 1 bi-4-por
00000020: 742d6769 672d6d34 2d6d616e 6167656d      t-gig-m4-managem
00000030: 656e742d 6d6f6475 6c650a6d 6f64756c      ent-module modul
00000040: 65203220 62692d32 342d706f 72742d63      e 2 bi-24-port-c
00000050: 6f707065 722d6d6f 64756c65 0a6d6f64      opper-module mod
00000060: 756c6520 33206269 2d342d70 6f72742d      ule 3 bi-4-port-
00000070: 6769672d 6d342d6d 616e6167 656d656e      gig-m4-managemen
00000080: 742d6d6f 64756c65 0a6d6f64 756c6520      t-module module
00000090: 34206269 2d342d70 6f72742d 6769672d      4 bi-4-port-gig-
000000a0: 6d342d6d 616e6167 656d656e 742d6d6f      m4-management-mo
000000b0: 64756c65 0a210a21 0a210a21 0a6d3220      dule ! ! ! m2
000000c0: 61637469 76652d6d 616e6167 656d656e      active-managemen
000000d0: 7420330a 69702061 64647265 73732031      t 3 ip address 1
000000e0: 39322e31 36382e32 2e353820 3235352e      92.168.2.58 255.
000000f0: 3235352e 3235352e 300a656e                255.255.0 end

```

Syntax: hd [slot1 | slot2] <file-name>

Each row of hexadecimal output contains the following parts:

- The byte offset of the data that is displayed to the right of the offset
- A row of hexadecimal data
- The ASCII equivalent of the hexadecimal data shown in the row

Creating a Subdirectory

To create a subdirectory on a flash card, enter a command such as the following:

```
BigIron# mkdir slot1 \TEST
```

To verify successful creation of the subdirectory, enter a command to change to the new subdirectory level:

```
BigIron# chdir \TEST
Current directory of slot1 is: \TEST
```

Syntax: md | mkdir [slot1 | slot2] <dir-name>

You can enter either **md** or **mkdir** for the command name.

The **slot1 | slot2** parameter specifies a PCMCIA slot. If you do not specify a slot, the command applies to the slot that currently has the management focus.

The <dir-name> parameter specifies the subdirectory name. You can enter a name that contains any combination of the following characters. Do not enter a backslash “/” in front of the name.

- All upper and lowercase letters
- All digits
- Spaces
- Any of the following special characters:
 - \$
 - %
 - '
 - -
 - _
 - @
 - ~
 - `
 - !
 - (
 -)
 - {
 - }
 - ^
 - #
 - &

You can use spaces in a file or subdirectory name if you enclose the name in double quotes. For example, to specify a subdirectory name that contains spaces, enter a string such as the following: “a long subdirectory name”.

A subdirectory or file name can be a maximum of 256 characters long. A complete subdirectory path name cannot contain more than 260 characters.

The name is not case sensitive. You can enter upper- or lowercase letters. The CLI displays the name using uppercase letters.

Removing a Subdirectory

To remove a subdirectory, enter a command such as the following:

```
BigIron# rmdir \TEST
```

Syntax: rd | rmdir [slot1 | slot2] <dir-name>

You can enter either **rd** or **rmdir** for the command name.

The **slot1 | slot2** parameter specifies a PCMCIA slot.

The <dir-name> parameter specifies the subdirectory you want to delete. You can enter a path name if the subdirectory is not in the current directory.

NOTE: You can remove a subdirectory only if the subdirectory does not contain files or other subdirectories.

If you receive a message such as the following, enter the **pwd** command to verify that the management focus is at the appropriate level of the directory tree.

```
BigIron# rmdir \TEST
File not found
```

Renaming a File

Use the following method to rename a file on a flash card.

USING THE CLI

To rename a file, enter a command such as the following:

```
BigIron# rename oldname newname
```

Syntax: rename [slot1 | slot2] <old-name> <new-name>

If the command is successful, the CLI displays a new command prompt.

Changing the Read-Write Attribute of a File

The read-write attribute specifies whether a file on a flash card can be changed or deleted.

- Read-only – You can display or copy the file but you cannot replace (copy over) or delete the file.
- Read-write – You can replace (copy over) or delete the file. This is the default.

Use the following method to change the read-write attribute of a file.

USING THE CLI

To protect a file from accidental changes by changing the read-write attribute from read-write to read-only, enter a command such as the following:

```
BigIron# attrib ro goodcfg.cfg
```

Syntax: attrib [slot1 | slot2] ro | rw <file-name>

To determine the read-write attribute of a file, use the **dir** command to list the directory information for the file. Files set to read-only are listed with “R” in front of the file name. See “Displaying a Directory of the Files on a Flash Card” on page 3-22.

To change all files on a flash card to read-only, enter a command such as the following:

```
BigIron# attrib ro *.*
```

This command changes the read-write attribute for all files on the flash card that currently has the management focus to read-only.

Deleting a File from a Flash Card

To delete a file from a flash card, use the following method.

CAUTION: By default, the delete option deletes all files on the flash card. Make sure you specify the files you want to delete.

CAUTION: The software does not have an undelete option. Make sure you really want to delete the file.

USING THE CLI

To delete a file on the flash card that has the management focus, enter a command such as the following:

```
BigIron# delete cfg.cfg
```

If the command is successful, the CLI displays a new command prompt.

Syntax: delete [slot1 | slot2] [<file-name>]

The command in this example deletes the specified file. To delete all files that contain a specific string of characters, enter a command such as the following:

```
BigIron# delete test*.*
```

This command deletes all files whose names start with “test”. To delete all the files on a flash card, enter a command such as the following:

```
BigIron# delete slot2
```

The command in this example deletes all files on the flash card in slot 2. In this example, slot 1 has the management focus, but the files to be deleted are on the flash card in slot 2.

Recovering (“Undeleting”) a File

You can undelete a command you have deleted from a flash card. To do so, enter a command such as the following:

```
BigIron# undelete
Undelete file "?LD.CFG" ? (enter 'y' or 'n'): y
Input one character: O
File recovered successfully and named to OLD.CFG
```

The command in this example starts the undelete process for the flash card and subdirectory that currently have the management focus. For each file that can be undeleted, the CLI displays the remaining name entry in the file directory and prompts you for the first character of the file name. You can enter any valid file name character. You do not need to enter the character that was used before in the deleted file name.

Once you enter a character and the CLI undeletes the file, the CLI continues with the next file that can be undeleted. For each file, specify “y” or “n”, and specify a first character for the files that you select to undelete.

To end the undelete process, enter the CTRL + C key combination.

Syntax: undelete [slot1 | slot2] [<to-dir-path>]

NOTE: When you delete a file from a flash card, the CLI leaves the file intact but removes the first letter in the file name from the file directory. However, if you save file changes or new files that use part of the space occupied by the deleted file, you cannot undelete the file. The **undelete** command lists only the files that can be undeleted.

Appending a File to Another File

You can append a file on a PCMCIA flash card to the end of another file. To append one file to another one, enter a command such as the following:

```
BigIron# append newacfs.cfg startup-config.cfg
```

This command appends a file called “newacfs.cfg” to the end of a file called “startup-config.cfg” file. This example assumes that both files are present on the PCMCIA slot and in the subdirectory level that currently have the management focus.

The following command appends a file in the current subdirectory to the end of a file in another subdirectory:

```
BigIron# append newacfs.cfg \TEST\startup-config.cfg
```

Syntax: append [<from-card> <to-card>] [<from-dir-path>]<from-name> [<to-dir-path>]<to-name>

The <from-card> and <to-card> parameters specify the source and destination flash cards when you are appending a file on one flash card to a file located on another flash card.

The [\<from-dir-path>]\<from-name> parameter specifies the file you are adding to the end of another file. If the file is not located in the current subdirectory (the subdirectory that currently has the management focus), specify the subdirectory path in front of the file name.

The [\<to-dir-path>]\<to-name> parameter specifies the file to which you are appending the other file. If the file is not located in the current subdirectory, specify the subdirectory path in front of the file name.

Copying Files

You can perform the following copy operations involving flash cards:

- Copy files from one flash card to the other.
- Copy files between a flash card and the device's flash memory.
- Copy files between a flash card and a TFTP server.
- Copy a startup-config file between a flash card and the device's flash memory.
- Copy the running-config file to a flash card
- Load a running-config file from a flash card into the device's running configuration (for loading ACLs only)
- Copy a POS image file from a flash card to a POS module's flash memory

NOTE: The copy options require you to explicitly specify the flash card. Therefore, you can perform a copy regardless of the flash card that currently has the management focus.

For convenience, the CLI provides two forms of each copy command. One form begins with **copy** and the other form begins with **ncopy**. The commands also differ in the order you specify the parameters.

- In the **copy** commands, you specify the device you are copying from, the device you are copying to, and then additional parameters if applicable. The additional parameters can include information such as the source file name and the new file name. For example, to copy a file from a flash card to flash memory, enter the following **copy** command: **copy slot1 | slot2 flash <from-name> primary | secondary**
- In the **ncopy** commands, you specify all the source information, then all the destination information. For example, to copy a file from a flash card to flash memory, enter the following **ncopy** command: **ncopy slot 1 | slot2 <from-name> flash primary | secondary**

The two forms of a copy command provide the same function. Use the form you are more comfortable with. The following sections provide examples using the **copy** forms and list the complete command syntax for both forms.

Copying Files from One Flash Card to the Other

Use the following methods to copy files from one flash card to the other.

USING THE CLI

To copy a file from one flash card to the other, enter the following command:

```
BigIron# copy slot1 slot2 sales.cfg
```

Syntax: copy <from-card> <to-card> [\<from-dir-path>]\<from-name> [[\<to-dir-path>]\<to-name>]

Syntax: ncopy <from-card> [\<from-dir-path>]\<from-name> <to-card> [[\<to-dir-path>]\<to-name>]

The command shown in the example above copies a file from the flash card in slot 1 to the flash card in slot 2. In this case, the software uses the same name for the original file and for the copy. Optionally, you can specify a different file name for the copy.

Copying Files Between a Flash Card and Flash Memory

Use the following methods to copy files between a flash card and the management module's flash memory.

USING THE CLI

To copy a file from a flash card to the primary area in flash memory, enter a command such as the following:

```
BigIron# copy slot1 flash B2P07000.bin primary
BigIron# Flash Erase -----
Flash Memory Write (8192 bytes per dot) .....
.....
.....code flash copy done
```

Syntax: copy slot1 | slot2 flash [*<from-dir-path>*]*<from-name>* primary | secondary

Syntax: ncopy slot1 | slot2 [*<from-dir-path>*]*<from-name>* flash primary | secondary

To copy a file from flash memory to a flash card, enter a command such as the following:

```
BigIron# copy flash slot2 BIS07000.bin primary
Flash Card Write (128 KBytes per dot) .....
Write to slot2 BIS07000.bin succeeded
```

The command in this example copies a software image file from the primary area in flash memory onto the flash card in slot 2.

If the copy does not succeed, the software lists messages to indicate the reason the copy did not work. For example, the following messages indicate that the copy did not work because the slot specified for the copy does not contain a flash card.

```
BigIron# copy flash slot2 m4s.car secondary
The system can not find the drive specified
Write to slot2 m4s.car failed
```

Syntax: copy flash slot1 | slot2 [*<to-dir-path>*]*<to-name>* primary | secondary

Syntax: ncopy flash primary | secondary slot1 | slot2 [*<to-dir-path>*]*<to-name>*

Copying Files Between a Flash Card and a TFTP Server

Use the following methods to copy files between a flash card and a TFTP server.

NOTE: The Foundry device must have network access to the TFTP server.

USING THE CLI

To copy a file from a flash card to a TFTP server, enter a command such as the following:

```
BigIron# copy slot1 tftp 192.168.1.17 notes.txt
Uploading 254 bytes to tftp server ...
Upload to TFTP server done.
```

Syntax: copy slot1 | slot2 tftp *<ip-addr>* [*<from-dir-path>*]*<from-name>* [*<to-name>*]

Syntax: ncopy slot1 | slot2 [*<from-dir-path>*]*<from-name>* tftp *<ip-addr>* [*<to-name>*]

To copy a file from a TFTP server to a flash card, enter a command such as the following:

```
BigIron# copy tftp slot1 192.168.1.17 notes.txt
Downloading from tftp server ...
Tftp 254 bytes done, copy to slot1 ...
Write to slot1 cfg.cfg succeeded
```

Syntax: copy tftp slot1 | slot2 *<ip-addr>* *<from-name>* [[*<to-dir-path>*]*<to-name>*]

Syntax: ncopy tftp *<ip-addr>* *<from-name>* slot1 | slot2 [[*<to-dir-path>*]*<to-name>*]

If the file name you specify is not on the TFTP server, the CLI displays messages such as those shown in the following example:

```
BigIron# copy tftp slot1 192.168.1.17 nots.txt
Downloading from tftp server ...
```

```
TFTP: received error request -- code 1 message File not found: C:/TFTP/notes.txt.  
Error - can't download data from TFTP server, error code 17. Abort!
```

To simplify troubleshooting, especially when the file is present on your server but the command doesn't find it, the messages list the complete TFTP path name on your TFTP server.

Copying the Startup-Config File Between a Flash Card and Flash Memory

Use the following methods to copy a startup-config file between flash memory and a flash card. By default, the device uses the startup-config in the primary area of flash memory to configure itself when you boot or reload the device.

NOTE: The device cannot use a startup-config file on a flash card to configure itself. You cannot boot or reload from a flash card.

USING THE CLI

To copy a startup-config file from a flash card to flash memory, enter a command such as the following:

```
BigIron# copy slot1 start test2.cfg  
..Write startup-config done.
```

Syntax: copy slot1 | slot2 start [*<from-dir-path>*]*<from-name>*

Syntax: ncopy slot1 | slot2 [*<from-dir-path>*]*<from-name>* start

This command copies a configuration file named test2.cfg from the flash card in slot 2 into the device's flash memory. The next time you reboot or reload the device, it uses the configuration information in test2.cfg.

To copy the device's startup-config file from flash memory onto a flash card, enter a command such as the following:

```
BigIron# copy start slot1 mfgtest.cfg  
Write to slot1 cfgtest.cfg succeeded
```

Syntax: copy start slot1 | slot2 [*<to-dir-path>*]*<to-name>*

Syntax: ncopy start slot1 | slot2 [*<to-dir-path>*]*<to-name>*

Copying the Running-Config to a Flash Card

Use the following method to copy the device's running-config to a flash card. The running-config contains the device's currently active configuration information. When you copy the running-config to a flash card, you are making a copy of the device's current configuration, including any configuration changes you have not saved to the startup-config file.

USING THE CLI

To copy the device's running configuration into a file on a flash card, enter a command such as the following:

```
BigIron# copy running slot1 runip.1  
Write to slot1 run.sw succeeded
```

Syntax: copy running slot1 | slot2 [*<to-dir-path>*]*<to-name>*

Syntax: ncopy running slot1 | slot2 [*<to-dir-path>*]*<to-name>*

Loading a Running-Config from a Flash Card

Use the following method to load configuration commands into the device's active configuration.

NOTE: A configuration file that you create must follow the same syntax rules as the startup-config file the device creates. See the release notes for a product for upgrade procedure for a software release.

USING THE CLI

To copy a running-config from a flash card, enter a command such as the following:

```
BigIron# copy slot2 running runip.2
```


Syntax: copy slot1 | slot2 running [\<<from-dir-path>\]<from-name>

Syntax: ncopy slot1 | slot2 [\<<from-dir-path>\]<from-name> running

The command in this example changes the device's active configuration based on the information in the file.

Copying a POS Image File from a Flash Card to a POS Module's Flash Memory

To copy a POS image file from a flash card to a POS module's flash memory, use the following method.

USING THE CLI

To copy a POS image file from a flash card onto all the POS modules in the chassis, enter a command such as the following:

```
BigIron# pos copy slot1 flash P2R07000.bin primary
```

Syntax: pos copy slot1 | slot2 flash <pos-image-file-name> primary | secondary [slot]

The command in this example copies a POS image file named P2R07000.bin from the flash card in slot 1 to all the POS modules in the chassis.

To copy a POS image file from a flash card onto a specific POS module, enter a command such as the following:

```
BigIron# pos copy slot1 P2R07000.bin flash primary 4
```

The command in this example copies the specified image file onto the POS module in chassis slot 4 only, but does not copy the file to other POS modules in the chassis.

The following command copies a POS image file from a TFTP server to flash memory. This command also is present in earlier software releases.

Syntax: pos copy tftp flash <ip-addr> <pos-image-file-name> primary | secondary [slot]

Loading the Software from a PCMCIA Flash Card

You can boot the flash software (system software) from a PCMCIA flash card.

Rebooting from the Flash Card

To reboot the device using a software image file on the flash card, enter a command such as the following at the Privileged Exec level of the CLI:

```
BigIron# boot system slot1 B2R07100.bin
```

The command in this example reboots the device using the image file B2R07100.bin located on the PCMCIA flash card in slot 1. This example assumes the image file is in the root directory on the flash card. If the image file is in a subdirectory, specify the subdirectory path. For example, to boot using an image in a subdirectory called "B2R", enter command such as the following:

```
BigIron# boot system slot1 \B2R\B2R07100.bin
```

Syntax: boot system slot1 | slot2 [\<<dir-path>\]<file-name>

The **slot1** | **slot2** parameter indicates the flash card slot.

The <file-name> parameter specifies the file name. If the file is in a subdirectory, specify the subdirectory path in front of the file name. If the file name you specify is not a full path name, the CLI assumes that the name (and path, if applicable) you enter are relative to the subdirectory that currently has the management focus.

NOTE: This command also is supported at the boot PROM.

Configuring the Flash Card as the Boot Source for Future Reboots

To set a flash card as the primary boot source, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# boot system slot1 B2R07100.bin
```

The command in this example sets PCMCIA slot 1 as the primary boot source for the device. When you reload the software or power cycle the device, the device looks for the flash image file you specify on the flash card in the slot you specify.

Syntax: boot system slot1 | slot2 <file-name>

NOTE: The command syntax is the same for immediately reloading and for changing the primary boot source, except the <file-name> must be the full path name. You cannot specify a relative path name. If the first character in the path name is not a backslash (\), the CLI treats the name you specify as relative to the root directory.

The device's response to the command depends on whether you enter the command at the Privileged EXEC level or the global CONFIG level.

If you enter multiple **boot system** commands at the global CONFIG level, the software places them in the running-config in the order you enter them, and saves them to the startup-config file in the same order when you save the configuration. When you reload or power cycle the device, the device tries the boot sources in the order they appear in the startup-config file and running-config.

Saving Configuration Changes to a PCMCIA Flash Card

You can configure the device to save configuration changes to a configuration file on a PCMCIA flash card.

Displaying the Current Location for Saving Configuration Changes

Enter the following command at the Privileged EXEC level of the CLI to display the current save location for the startup-config file:

```
BigIron# locate startup-config
```

Syntax: locate startup-config

Specifying the Location for Saving Configuration Changes

By default, when you save configuration changes, the changes are saved to the startup-config file on the device's flash memory module. If you want to change the save location to a PCMCIA slot, enter a command such as the following:

```
BigIron# locate startup-config slot1 router1.cfg
BigIron# write memory
```

The first command in this example sets the device to save configuration changes to the file named "router1.cfg" in the flash card in PCMCIA slot 1. The second command saves the running-config to the router1.cfg file on the flash card in slot 1.

NOTE: In this example, after you save the configuration changes using the **write memory** command, the router1.cfg file will include the command that designates PCMCIA slot 1 as the save location for configuration changes.

Syntax: locate startup-config [slot1 | slot2] <file-name>

You can specify a relative path name or full path name as part of the file name.

File Management Messages

The following table lists the messages the CLI can display in response to file management commands.

Table 3.3: Flash Card File Management Messages

This Message...	Means...
File not found	You specified a file name that the software couldn't find. Verify the command you entered to make sure the command matches the source and destination you intended for the file operation.
Current directory of <slot-num> is: <dir-path>	You have successfully changed the management focus to the slot and subdirectory indicated by the message.
Path not found	You specified an invalid path.
There is not enough space on the card	The flash card does not have enough space to hold the file you are trying to copy to it.
Warning: Can not copy file bigger than 3145728 bytes. File will be truncated	The file you are trying to copy exceeds the maximum file size allowed for copy operations.
Access is denied	You tried to copy or delete a file that has the read-only attribute.
A duplicate file name exists	You tried to rename a file using a name that is already in use by another file.
Fatal error, can not read or write media	A hardware error has occurred. One possible cause of this message is if you removed the flash card while a file operation involving the card was in progress.
The system can not find the drive specified	The slot you specified does not contain a flash card.
Fatal error, File Allocation Table corrupted	The flash card has been corrupted. Try the file operation you were attempting again. If you receive the same message, you may need to reformat the flash card.
Media is unformatted or contains unknown file system	The card has not been formatted or has been formatted for a file system the Management 4 module does not recognize. To use the card in the Management 4 module, you need to format it for the 16 FAT file system.
There is sharing conflict between format command and other read/write operations	The flash card is currently undergoing formatting. This message also can show up if you enter a command to format the card while the card is being accessed for another file operation.
Invalid DOS file name	A filename you entered contains an invalid character (for example, "." or "\").
File recovered successfully and named <file-name>	A file you tried to recover was successfully recovered under the name indicated in the message

Using a 3Com Management Interface in the PCMCIA Slot

You can insert a 3Com Megahertz 10/100 LAN PC Card (model 3CCFE574BT) into a PCMCIA flash card slot on a Management 4 module and use the PC Card as a management interface to the device.

NOTE: This feature supports 3Com Megahertz 10/100 LAN PC Card model 3CCFE574BT only.

NOTE: This feature applies only to the Management 4 module.

For PC Card information, see www.pc-card.com.

Usage Notes

- You can insert the PC Card in PCMCIA slot 1 or 2.

NOTE: You can remove and insert a PC Card in the same slot without reloading the software. However, if you remove the card from one slot and insert it into the other slot, you must reload the software (using the **reload** command) to place the change into effect.

- Only one management port is supported.
- The port supports Telnet, TFTP, and ICMP pings.
- The management port does not perform traffic forwarding (switching).
- The management port is a member of the default VLAN.
- The management port cannot be a member of another VLAN and cannot be configured as part of a trunk group. The port does not support rate limiting, multicast, or jumbo packets. CLI commands for configuring and managing other ports (switching ports) do not display information about the management port and cannot configure the port.
- The power saving option is not supported.

Installing the PC Card

1. When the management module is either powered down or is powered on and has completely booted, insert the card into a PCMCIA slot. Push gently yet firmly until the card is fully seated.
2. Insert the PC Card cable that came with the card into the port on the card. Make sure the icon on the cable connector is facing upward (toward the top of the PCMCIA slot).
3. Insert one end of a crossover or straight-through cable into the other end of the PC Card cable:
 - If you are connecting directly to a management device (such as a PC), use a Cat-5 crossover cable.
 - If you are connecting to a hub or switch, use a Cat-5 straight-through cable.
4. Insert the other end of the cable into the management device, hub, or switch. Once a link is established, the LED on the PCMCIA slot will light up.

Removing the PC Card

- If you need to remove the card, squeeze the release clips on the sides of the PC Card cable connector, then pull the connector straight out of the card.
- When the management module is either powered down or is powered on and has completely booted, pull the card out of the PCMCIA slot.

Changing Parameter Settings

You can insert the PC Card into PCMCIA slot 1 or slot 2. No configuration is required, although you can change the following parameter settings.

- Port speed (10 Mbps, 100 Mbps, or auto-sensing) – The default is 10 Mbps.
- Mode (half-duplex or full-duplex) – The default is half-duplex.
- 802.3x Flow control (applies only to full-duplex mode) – Disabled by default.

- IP address of the management port (applies only to Layer 3 code) – None configured by default.

Changing the Port Speed

The default port speed is 10 Mbps. To change the port speed, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# set pcmcia slot1 auto
```

Syntax: [no] set pcmcia slot1 | slot2 10 | 100 | auto

The **slot1** | **slot2** parameter specifies the PCMCIA slot. There is no default.

The **10** | **100** | **auto** parameter specifies the port speed. The default is 10 Mbps.

Changing the Mode

To change the mode, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# set pcmcia slot1 full
```

Syntax: [no] set pcmcia slot1 | slot2 full | half

The **slot1** | **slot2** parameter specifies the PCMCIA slot. There is no default.

The **full** | **half** parameter specifies the mode. The default is **half** (half-duplex).

Enabling 802.3x Flow Control

To enable 802.3x flow control, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# set pcmcia slot2 flow-control
```

Syntax: [no] set pcmcia slot1 | slot2 flow-control

Configuring an IP Address

To configure an IP address on the management port, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# set pcmcia slot1 ip 10.10.10.2/24
```

Syntax: [no] set pcmcia slot1 | slot2 ip <ip-addr> <ip-mask>

or

Syntax: [no] set pcmcia slot1 | slot2 ip <ip-addr>/<mask-bits>

Displaying Management Port Information

To display information for a management port, enter the following command:

```
BigIron(config)# show pcmcia all
3Com Megahertz 574B lan card present at slot 1
  Hardware is 100mbitEthernet,   Mac address is 0001.03aa.2902
  Speed is 10mb, Duplex is hdx, Flow_control is off
  Member of L2 Vlan ID 1, port is untagged, not a member of any trunk.
```

```
PCMCIA card is not present at slot 2
```

Syntax: show pcmcia all | slot1 | slot2

Chapter 4

Using the Velocity Management Module

The Velocity Management Module version 1 (VM1) is a redundant management module for BigIron Layer 3 Switch. The VM1 supports all of the features supported by Management 2, 3, and 4 modules, but enhances feature performance using new hardware architecture.

NOTE: This chapter does not describe how to configure redundancy parameters. See “Using Redundant Management Modules” on page 3-1.


Overview

The VM1 provides enhanced performance using distributed processing among multiple CPUs. The multiple CPUs enable the VM1 to perform the following in hardware:

- Process Access Control Lists (ACLs)
- Perform Policy-Based Routing (PBRs)
- Perform Network Address Translation (NAT)
- Collect statistics and export them for NetFlow-based accounting and billing

Figure 4.1 shows the VM1.

Figure 4.1 Velocity Management Module version 1 (VM1)



The VM1 does not have network interfaces but does have a serial management interface. In addition, the module has status LEDs for its Management Processor (MP) and Velocity Switching Processor (VSPs), described in “Management and Co-Processing CPUs” on page 4-1 and “Status LEDs” on page 4-13.

Management and Co-Processing CPUs

The VM1 contains four CPUs:


- One MP (Management Processor) – The MP performs management functions for the entire chassis.

- Three Velocity Switching Processor (VSPs) – The VSPs perform Layer 2 and Layer 3 switching for the forwarding modules.

The MP and the VSP have their own flash memory with primary and secondary areas.

Figure 4.2 illustrates the architecture of the VM1.

Figure 4.2 Architecture of VM1



Feature Coexistence

The VM1 architecture allows all the following features to be configured and active on a given port at the same time.

- Input ACLs
- Input rate limiting
- NetFlow Export
- sFlow Export
- Network Address Translation (NAT)
- Policy-Based Routing (PBR)
- Output ACLs
- Output rate limiting

When two or more of these features are applicable for a packet, the VM1 processes the features in the order listed above.

Temperature Sensor

The VM1 also contains a temperature sensor. The sensor generates a Syslog message and SNMP trap if the module's temperature exceeds a specified warning level or shutdown level. You can use the CLI or Web management interface to display the management module's temperature and to change the warning and shutdown temperature levels. See "Using the Temperature Sensor" on page 9-52.

Management Redundancy

The VM1 supports management redundancy. You can install a second VM1 to act as a backup and take over management of the Layer 3 Switch if the active VM1 becomes unavailable.

Management redundancy is described in “Using Redundant Management Modules” on page 3-1. Management redundancy using a pair of VM1 modules works as described in the chapter, with the following important differences:

- The VSP CPUs on both modules actively process traffic. Only the MP CPU on the standby module is in backup mode. The VSP CPUs on the standby module actively process traffic.
- The VSP CPU flash code is not automatically synchronized. To synchronize the flash code on the VSP CPUs, use the **vm copy tftp flash** command, described in “File Synchronization Between the Active and Standby Redundant Management Modules” on page 3-11. The flash code on the CPU is automatically synchronized.
- If you use a pair of VM1 management modules in a chassis for redundancy, the device does not reassign the forwarding modules assigned to the VSP CPUs on the active module to the other module following a hot swap. See the next section.

Management Redundancy and Hot Swap

If you use a pair of VM1 management modules in a chassis for redundancy, the device does not reassign the forwarding modules assigned to the VSP CPUs on the active module to the other module following a hot swap. This is true in the following cases:

- If you insert a standby VM1 into an active device, the device does not replicate the assignments of the forwarding modules to the VSP CPUs on the standby module. To work around this issue, use the **vm-map** command to assign the forwarding modules to the VSP CPUs on the standby module after you insert the module.
- If you remove a standby VM1 module that has taken over forwarding on an active device, the forwarding modules assigned to the VSP CPUs on the standby module are not reassigned to the VSP CPUs on the default active module. To avoid traffic interruption, use the **vm-map** command to assign the forwarding modules to the VSP CPUs on the default active VM1 module **before removing the standby module**.

To list the VSP CPU assignments, enter the following command: **show vm-map**

To assign forwarding modules to VSP CPUs, enter the following command:

```
vm vm-map <from-slotnum> vm-slot <to-slotnum> vm-cpu <cpunum>
```

The <from-slotnum> parameter specifies the slot that contains the forwarding module.

The <to-slotnum> parameter specifies the slot that contains the VM1.

The <cpunum> parameter specifies the VSM CPU on <to-slotnum> that will perform the processing. The VSM CPUs are numbered from 1 – 3.

VSP Load Sharing

The VM1 optimizes performance by distributing responsibility for the forwarding modules across the VSPs, so that each VSP has sole responsibility for a given forwarding module and the modules are as evenly distributed across the VSPs in terms of bandwidth.

When you power on or reset the VM1, the module assigns each of the forwarding modules to a VSP according to each module’s weight. A forwarding module’s weight is a number that represents its total forwarding capacity. The weight is measured in units of 1 for each 100 Mbps. For example, Table 4.1 shows the weights for some common forwarding module types. Notice that the weight for 10/100 modules is based on the higher bandwidth (100 Mbps instead of 10 Mbps) for all ports.

Table 4.1: Forwarding Module Weights

Module type	Total Mbps capacity	Weight
24-port 10/100 Mbps	2400	24
4-port 1000 Mbps	4000	40
8-port 1000 Mbps	8000	80

The device assigns the forwarding modules to VSPs in numerical order (always starting with VSP 1) and beginning with the module with the highest weight and working down to the module with the lowest weight.

The device assigns a forwarding module's ports to only one VSP. A single module's ports are never distributed across multiple VSPs.

The allocations determine the VSP that will process traffic received on a forwarding module's ports. For example, if an 8-port Gigabit module in slot 3 is allocated to VSP 1, then that CPU processes all the traffic received on the module's ports.

NOTE: If you hot-swap a module into or out of the chassis after the allocations have taken place at startup, the device does not re-allocate modules to even out the load sharing. Instead, the device allocates the module you insert to the VSP that currently has the least weight allocated to it. If you remove a module, the device subtracts the module's weight from the VSP to which the module was allocated.

Here are some examples of load sharing allocations for various configurations. Notice that for a four-slot chassis, each forwarding module is allocated to its own VSP. The module's weights determine the VSPs to which they are allocated. For a chassis with more than four slots, some VSPs are allocated more than one module. Nonetheless, the allocations are based on the forwarding modules' weights and provide the most even distribution possible.

Example Configuration 1

Table 4.2 shows a module configuration and the resulting VSP allocations for a four-slot chassis. Notice that since the VM1 does not have any forwarding ports, the module does not need to be allocated to a VSP.

Table 4.2: Example Configuration 1

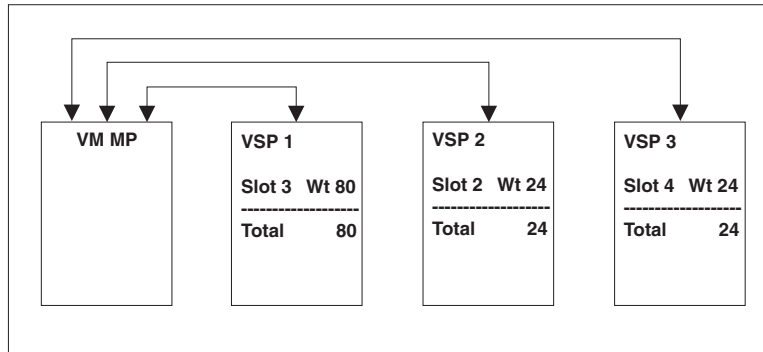
Slot	Module type	Weight	Order allocated	VSP
1	VM1	n/a	n/a	n/a
2	24-port 10/100	24	2	VSP 2
3	8-port Gigabit	80	1	VSP 1
4	24-port 10/100	24	3	VSP 3

Figure 4.3 shows the VSP allocations for this configuration.

Figure 4.3 VSP allocations for example configuration 1

VM MP = Management Processor

VSP = Velocity Switching Processor



The device begins with the highest-weight module, in this case the 8-port Gigabit module in slot 3, and allocates that module's ports to VSP 1. The device then allocates the module with the second-highest weight, in this case the 24-port 10/100 module in slot 2, to the next VSP with the lowest allocated weight, which is VSP 2. Finally, the device allocates the last forwarding module, the 24-port 10/100 module in slot 4, to the next VSP with the lowest allocated weight, VSP 3.

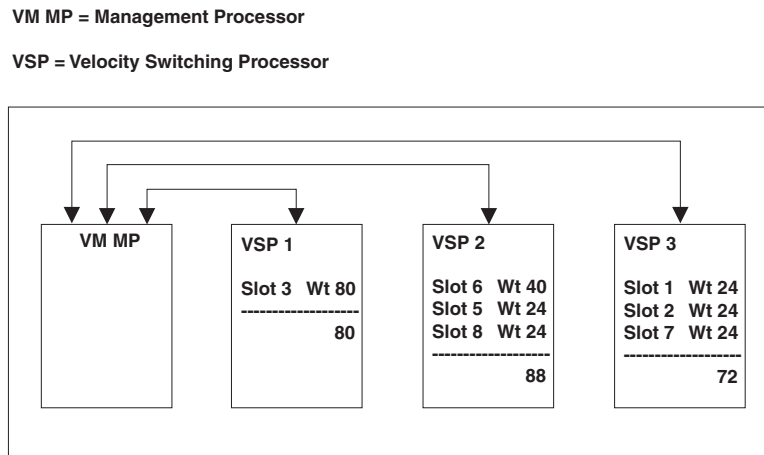
Example Configuration 2

Table 4.3: Example Configuration 2

Slot	Module type	Weight	Order allocated	VSP
1	24-port 10/100	24	3	VSP 3
2	24-port 10/100	24	4	VSP 3
3	8-port Gigabit	80	1	VSP 1
4	VM1	n/a	n/a	n/a
5	24-port 10/100	24	5	VSP 2
6	4-port Gigabit	40	2	VSP 2
7	24-port 10/100	24	6	VSP 3
8	24-port 10/100	24	7	VSP 2

Figure 4.4 shows the VSP allocations for this configuration.

Figure 4.4 VSP allocations for example configuration 2



As in the previous example, the device starts with the highest-weight module, in this case the 8-port Gigabit module in slot 3, and allocates that module to VSP 1. The device then allocates the second-highest weighted module to VSP 2, and the third-highest weighted module to VSP 3. For the next module, the device selects the VSP with the lowest allocated weight; in this case, that is VSP 3. And so on. As shown in this example, the resulting distribution is fairly even among the three CPUs.

Displaying the Slot Allocations for the VSPs

To display the allocations, enter the **show vm-map** command. See “Determining the Slot Allocations for the VSPs” on page 4-14.

Changing Slot Allocations

The default allocations are applicable to almost all configurations. However, you can remap a module to another VSP CPU. To do so, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# vm vm-map slot 3 vm-slot 2 vm-cpu 1
```

This command remaps processing for the modules in slot 3 to VSP CPU 1 on the VM1 in slot 2.

Syntax: vm vm-map <from-slotnum> vm-slot <to-slotnum> vm-cpu <cpunum>

The <from-slotnum> parameter specifies the slot that contains the forwarding module.

The <to-slotnum> parameter specifies the slot that contains the VM1.

The <cpunum> parameter specifies the VSM CPU on <to-slotnum> that will perform the processing. The VSM CPUs are numbered from 1 – 3.


NOTE: Starting with release 07.6.01, you can assign individual MPLS-enabled Gigabit ports to specific VSPs.

VSP Load Sharing on a Per-DMA Basis

Starting in release 07.6.03, the VM1 supports VSP load sharing on a per-DMA basis. You can configure the VM1 to use either per-module or per-DMA VSP load sharing, and you can statically assign ports or slots to individual VSPs.

In releases prior to 07.6.03, the VM1 distributes the load to the VSPs on a per-module basis. When the Foundry device is powered on or reset, the VM1 assigns each of the forwarding modules to a VSP.

In release 07.6.03, the VM1 can distribute the load to the VSPs on a per-DMA basis. DMAs are packet processors that control ports on Ethernet modules. Ethernet modules have multiple DMAs, each controlling a set of ports on the module. For example, an IronCore 8-port Gigabit Ethernet module has four DMAs, each controlling two ports. The following diagram illustrates the relationship between ports and DMAs on an IronCore 8-port Gigabit Ethernet module.

Figure 4.5 DMAs and ports on an IronCore 8-port Gigabit Ethernet module

On an IronCore 8-port Gigabit Ethernet module, separate DMAs control ports 1 – 2, 3 – 4, 5 – 6, and 7 – 8. When per-DMA VSP load sharing is enabled, the VM1 assigns forwarding responsibility for each DMA's ports to a VSP so that the forwarding load is balanced among the VSPs. This means that a single module's ports can be distributed across multiple VSPs. In previous releases, the VM1 assigned all of a module's ports to only one VSP.

NOTE: In release 07.6.03, per-DMA VSP load sharing is supported only for IronCore 8-port Gigabit Ethernet modules. Modules that do not support per-DMA VSP load sharing will have all of their ports assigned to a single VSP, even if per-DMA VSP load sharing is enabled on the device.

Configuring Per-DMA VSP Load Sharing

To configure the VM1 to use per-DMA VSP load sharing, you can do the following:

- Assign ports to individual VSPs statically. Forwarding for all of the ports controlled by the specified port's DMA are handled by the specified VSP.
- Allow the VM1 to assign DMAs to VSPs dynamically. When the device is started or reset, the VM1 load balances processing by assigning DMAs to the VSPs according to the total bandwidth of the DMAs.

The following command assigns the DMA that controls port 2/1 to VSP 1 on the VM1 in slot 1:

```
BigIron(config)# vm vm-map port-dma 2/1 vm-slot 1 vm-cpu 1
```

Syntax: [no] vm vm-map port-dma <port> vm-slot <slot> vm-cpu <VSP-cpu>

To assign a DMA to a VSP, you specify any of the ports controlled by the DMA as the <port> parameter. Forwarding for all of the ports controlled by the DMA is then handled by the specified VSP. This command is similar to the **vm vm-map slot** command, which allows specific modules to be assigned to specific VSPs.

To configure the VM1 to assign DMAs to VSPs dynamically at startup, enter the following command:

```
BigIron(config)# vm vm-map per-port-dma
```

Syntax: [no] vm vm-map per-port-dma

If the **vm vm-map per-port-dma** command is in the Foundry device's configuration when the device is started or reset, the VM1 uses per-DMA VSP load sharing. Otherwise, the VM1 uses per-module VSP load sharing to balance forwarding among the VSPs.

If any ports or modules are statically assigned to VSPs, then those assignments are made prior to any dynamic assignments. You can have both per-module static assignments and per-DMA static assignments in a configuration.

Static per-DMA assignments take precedence over static per-module assignments. For example, if the ports controlled by DMA 1 in slot 2 are statically assigned to VSP 1 in slot 1, and the module in slot 2 is statically assigned to VSP 2 in slot 1, then all the ports controlled by all the DMAs except DMA 1 are assigned to VSP 2, and the ports controlled by DMA 1 are assigned to VSP 1.

Changing the Management Session from the MP to a VSP

By default, management sessions you open with the VM1 are established with the MP. However, you can establish a session directly with a VSP. Each VSP supports some commands at the Privileged EXEC level.

NOTE: You can enter configuration commands only to the MP, not directly to a VSP.

The CLI provides a remote login facility for changing the management session to a VSP. When you log in to a VSP, the CLI management session changes from the MP to the VSP. At this point, commands apply only to the VSP. To enter commands to the MP, you must log out of the VSP. The CLI prompt changes to indicate the chassis slot number and VSP you are logged on to.

Logging In to a VSP

To log in to a VSP, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# rconsole 2 1
BigIron2/1 #
```

This command changes the management session from the MP to VSP 1 on the VM1 in slot 2. Notice that the end of the command prompt changes to indicate the slot number and VSP number.

Syntax: rconsole <slotnum> <cpunum>

The <slotnum> parameter specifies the chassis slot that contains the module.

- Slots in a four-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an eight-slot chassis are numbered 1 – 8, from left to right.
- Slots in a fifteen-slot chassis are numbered 1 – 15, from left to right.

The <cpunum> parameter specifies the VSP. The VSPs are numbered from 1 – 3.

Logging Out from the VSP

To log out from a management session with a VSP, enter the following command at the User EXEC or Privileged EXEC level:

```
BigIron2/1 # rconsole-exit
BigIron#
```

Syntax: rconsole-exit

NOTE: You must enter the entire command name (**rconsole-exit**). The CLI will not accept abbreviated forms of the command.

VSP Commands

The following commands are supported at the VSP command prompt:

- **rconsole-exit** – Logs out of the VSP.
- **show ?** – Displays the available show commands. The following show commands are available:
 - **show arp** – Displays the ARP table.
 - **show filter** – Displays configured filters.
 - **show ip access-lists** – Shows the configured ACLs.
 - **show ip cache** – Shows the IP cache.
 - **show ip nat** – Shows NAT information.
 - **show ip route** – Shows the IP route table.

- **show mac-address** – Shows the MAC table.
- **show running-config** – Shows the running-config.
- **show usage** – Shows Layer 4 session table information.
- **show trunk** – Shows trunk group information.
- **show vlans** – Shows VLAN information.
- **write terminal** – Displays the running-config on the management console.

With a few exceptions, the command syntax and displays are the same as described in the *Foundry Switch and Router Command Line Interface Reference*. Here are the exceptions:

- The **show ip route** command displays only 20 entries at a time. The command has an optional parameter, <num>, that indicates the entry at which you want the display to begin.
- The output of the **show trunk** and **show vlans** commands is different from the output format for these commands when entered on the MP.

Displaying VM1 Module Information

You can display the following VM1 information:

- Software versions – see “Displaying the Software Version Running on the Module” on page 4-9
- General module information – “Displaying General Module Information” on page 4-11
- Module status – see “Determining Module Status” on page 4-12
- Slot allocations for the VSPs – see “Determining the Slot Allocations for the VSPs” on page 4-14

The commands in this section are supported on the MP, not on the VSPs.

Displaying the Software Version Running on the Module

To display the software version running on the VM1, use either of the following methods.

USING THE CLI

To display the software version running on the module, enter the following command at any CLI level:

```
MON-BigIron# show version
  SW: Version 07.5.00T53 Copyright (c) 1996-2001 Foundry Networks, Inc.
      Compiled on Oct 28 2001 at 15:54:49 labeled as VM1R07500
      (2852369 bytes) from Primary VM1r07500.bin
  HW: BigIron 8000 Router, SYSIF version 21
=====
SL 2: VM1 Management Module, SYSIF II, VM, ACTIVE
  0 MB SHM, 3 Application Processors
8192 KB BRAM, SMC version 1, ICBM version 21
SW: (1)07.5.00b2SPT72 (2)07.5.00b2SPT72 (3)07.5.00b2SPT72
=====
SL 3: B24E Copper Switch Module
  2048 KB BRAM, SMC version 2, ICBM version 21
  256 KB PRAM(256K+0K) and 2048*8 CAM entries for DMA 8, version 0808
  256 KB PRAM(256K+0K) and shared CAM entries for DMA 9, version 0808
  256 KB PRAM(256K+0K) and shared CAM entries for DMA 10, version 0808
=====
SL 4: B8G Fiber Switch Module
  2048 KB BRAM, SMC version 1, ICBM version 21
  256 KB PRAM(256K+0K) and 2048*8 CAM entries for DMA 12, version 0209
  256 KB PRAM(256K+0K) and shared CAM entries for DMA 13, version 0209
  256 KB PRAM(256K+0K) and 2048*8 CAM entries for DMA 14, version 0209
  256 KB PRAM(256K+0K) and shared CAM entries for DMA 15, version 0209
=====
Active management module:
  500 MHz Power PC processor 750 (version 8/8302) 66 MHz bus
  512 KB boot flash memory
16384 KB code flash memory
  256 KB SRAM
  512 MB DRAM
Monitor Option is on
The system uptime is 42 minutes 6 seconds
The system : started=warm start   reloaded=by "reload"
```

Syntax: show version

The command shows information about the VM1 and also lists all the software versions running on the device. The VM1 information is shown in this example in bold text.

USING THE WEB MANAGEMENT INTERFACE

You cannot display the module software versions using the Web management interface.

Displaying the Software Versions Installed on the Module

To display the software versions installed in the flash areas of the MP and the VSPs, use the following method.

USING THE CLI

To display the software in the device's flash areas, enter the following command at any CLI level:

```
MON-BigIron(config)# show flash
Active management module:
Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304, Unit: 4
Boot Flash Type: AMD 29F040, Size: 8 * 65536 = 524288
Compressed Pri Code size = 2852369, Version 07.5.00b2SPT23 (VM1r07500.bin)
Compressed Sec Code size = 2848200, Version 07.5.00bT22
Maximum Code Image Size Supported: 7011840 (0x006afe00)
Boot Image size = 162664, Version 07.03.00 (m2b.bin)
Monitor Image Version 4, for DRAM size 268435456
VM 2/1: Pri (1231492, 07.5.00b2SPT72), Sec (1004047, 07.2.11T71) Boot
(07.01.00)

VM 2/2: Pri (1231492, 07.5.00b2SPT72), Sec (1004047, 07.2.11T71) Boot
(07.01.00)

VM 2/3: Pri (1231492, 07.5.00b2SPT72), Sec (1004047, 07.2.11T71) Boot
(07.01.00)
```

Syntax: show flash

The lines highlighted in bold in this example list the software installed on the module:

- The Compressed Pri Code and Compressed Sec Code lines list the flash code installed in the flash areas on the module.
- The Boot Image line lists the boot code.
- The VM lines list the flash images and boot code installed on the VSPs. The numbers following "VM" indicate the chassis slot number that contains the VM1 and the VSP number on the VM1.

Displaying General Module Information

To display general module information, use the following method.

USING THE CLI

To display general information for a VM1, enter the following command at any CLI level:

```
MON-BigIron(config)# show vm-state
=====
VM MODULE (2) App CPU    0 MB SHM, 3 Application Processors
      CPU 0 in state of VM_STATE_RUNNING
      CPU 1 in state of VM_STATE_RUNNING
      CPU 2 in state of VM_STATE_RUNNING
-----
Module 2 App CPU 1, SW: Version 07.3.00b2SPT72
Compiled on Jun 04 2001 at 17:14:08 labeled as VSP07300b2SP
DRAM 268M, BRAM 262K, FPGA Version 0050
Code Flash 4M: Primary (1231492 bytes, 07.3.00b2SPT72),
                Secondary (1004047 bytes, 07.2.11T71)
Boot Flash 131K, Boot Version 07.01.00
The system uptime is 0 day 0 hour 44 minute 15 second
General Status: 0 ipc msg rec, 2 ipc msg sent
-----
Module 2 App CPU 2, SW: Version 07.3.00b2SPT72

Compiled on Jun 04 2001 at 17:14:08 labeled as VSP07300b2SP
DRAM 268M, BRAM 262K, FPGA Version 0050
Code Flash 4M: Primary (1231492 bytes, 07.3.00b2SPT72),
                Secondary (1004047 bytes, 07.2.11T71)
Boot Flash 131K, Boot Version 07.01.00
The system uptime is 0 day 0 hour 44 minute 15 second
General Status: 0 ipc msg rec, 2 ipc msg sent
-----
Module 2 App CPU 3, SW: Version 07.3.00b2SPT72
Compiled on Jun 04 2001 at 17:14:08 labeled as VSP07300b2SP
DRAM 268M, BRAM 262K, FPGA Version 0050
Code Flash 4M: Primary (1231492 bytes, 07.3.00b2SPT72),
                Secondary (1004047 bytes, 07.2.11T71)
Boot Flash 131K, Boot Version 07.01.00
The system uptime is 0 day 0 hour 44 minute 15 second
General Status: 0 ipc msg rec, 2 ipc msg sent
```

Syntax: show vm-state

This command displays the state of the VM1, the software version running on the module, and detailed information for each VSP on the module.

USING THE WEB MANAGEMENT INTERFACE

You cannot display general VM1 information using the Web management interface.

Determining Module Status

You can determine the status of a VM1 in the following ways:

- Status LEDs – Each VSP has LEDs that show send and receive activity for the processor. The MP has LEDs for data activity (both send and receive) and power.
- Module information in software – The module information displayed by the software indicates whether the module came up properly.

Status LEDs

You can determine the status of a VM1 processor by observing its LEDs. The processors have the following LEDs. Each VSP has its own column of TxAct and RxAct LEDs. The left column shows activity for VSP 1, the middle column shows activity for VSP 2, and the right column shows activity for VSP 3.

Table 4.4: VM1 LEDs

LED	Position	State	Meaning
Active	Upper LED to the left of the serial interface	On	The MP is active.
		Off	The MP is not active.
Power	Lower LED to the left of the serial interface	On	The power status is good.
		Off	The power status is not good.
TxAct	Upper LED near the middle of the module	Blinking	The VSP is transmitting data.
RxAct	Lower LED near the middle of the module	Blinking	The VSP is receiving data.

Software

You can display status information for a VM1 using either of the following methods.

NOTE:

- Slots in a four-slot chassis are numbered 1 – 4, from top to bottom.
 - Slots in an eight-slot chassis are numbered 1 – 8, from left to right.
 - Slots in a fifteen-slot chassis are numbered 1 – 15, from left to right.
-

USING THE CLI

To display the status of a VM1 using the CLI, enter the following command at any CLI level:

```
BigIron(config)# show module
Module                Status    Ports Starting MAC
S1:
S2: Configured as B0GMR VM Management Module
S3: B24E Copper Switch Module    OK        24    00e0.52c2.9f40
S4: B24E Copper Switch Module    OK        24    00e0.52c2.9f60
S5:
S6: B0GMR VM Management Module    ACTIV    0
S7:
S8:
```

Syntax: show module

The Status column shows the module status. A VM1 can have one of the following statuses:

- ACTIVE – The module is currently the active management module.

- STANDBY – The module is the standby management module. (This applies only to management modules that support redundancy.)
- COMING UP – The module is coming up as the standby module. This status can be observed during switchover.
- FAILED – This status indicates that the host module failed to come up.
- OK – This status indicates that the module came up and is operating normally.

NOTE: The ACTIVE, STANDBY, and COMING UP status values apply only to management modules.

USING THE WEB MANAGEMENT INTERFACE

1. Select the [Home](#) link to display the System configuration sheet, if not already displayed.
2. Select the [Module](#) link to display the Module panel. The Status column shows the module status. A Web Switching module can have one of the following statuses:
 - ACTIVE – The module is currently the active management module.
 - STANDBY – The module is the standby management module. (This applies only to management modules that support redundancy.)
 - COMING UP – The module is coming up as the standby module. This status can be observed during switchover.
 - FAILED – This status indicates that the host module failed to come up.
 - OK – This status indicates that the module came up and is operating normally.

NOTE: The ACTIVE, STANDBY, and COMING UP status values apply only to management modules.

Determining the Slot Allocations for the VSPs

The VM1 automatically load balances processing by allocating chassis slots to the VSPs according to the total bandwidth of the modules in the slots. To list the slot allocations, use the following CLI method.

USING THE CLI

To display the slot allocations for the VSPs, enter the following command at any CLI level:

```
BigIron(config)# show vm-map
slot 2 (weight 24 x 100M) is processed by VSP 1/2 (weight 24)
slot 3 (weight 8 x 1000M) is processed by VSP 1/1 (weight 80)
slot 4 (weight 24 x 100M) is processed by VSP 1/3 (weight 24)
```

Syntax: show vm-map

This example shows the slot allocations for a four-slot chassis. The output displays rows only for the slots that contain forwarding modules. No information is displayed for empty slots.

Each row shows the following information:

- The chassis slot (“slot 2” in the first row of the example above)
- The weight of the module in the slot (“weight 24 x 100M” in the first row of the example above)
- The chassis slot that contains the VM1 and the VSP to which the forwarding module described by this row is allocated (“is processed by VSP 1/2”). The “1” in this example indicates the VM1 is in chassis slot 1. The “2” in this example indicates that VSP 2 is handling the forwarding module in slot 2.

- The total weight assigned to the VSP ("weight 24" in the first row of this example).

NOTE: If the ports on a module are not up, the output says "will be processed" instead of "is processed" and the weight is listed as "0". In this case, the VM1 reserves a VSP for the module but does not add weight for the module's ports to the reserved VSP.

NOTE: For reference, this example matches "Example Configuration 1" on page 4-4.

When per-DMA VSP load sharing is enabled on the device, the show **vm-map** command displays static VSP assignments. For example:

```
BigIron# show vm-map
slot 1 (weight 80 x 100M):
  e  1/5-1/6  is processed by VSP processor 4/2
  e  1/7-1/8  is processed by VSP processor 4/3
slot 2 (weight 24 x 100M) is processed by VSP processor 4/1

Static configuration:
slot 1 (weight 80 x 100M):
  e  1/1-1/2  is processed by VSP processor 4/3
  e  1/3-1/4  is processed by VSP processor 4/1
```

In the example above, per-DMA VSP load sharing has been enabled on the device. The module in slot 1 supports per-DMA VSP load sharing, but the module in slot 2 does not. The VM1 is located in slot 4.

On the module in slot 1, the DMAs controlling ports 1 – 2 and 3 – 4 have been statically assigned to VSPs. The DMAs controlling the other ports on the module have been dynamically assigned to VSPs based on the weight of the DMAs.

All of the ports on the module in slot 2 are assigned to VSP 4/1. Since the module does not support per-DMA VSP load sharing, all of its ports are assigned to a single VSP.

Chapter 5

Using 10 Gigabit Ethernet Modules

This chapter describes the Foundry 10 Gigabit Ethernet modules. It contains the following topics:

- “1-Port 10 Gigabit Ethernet Module” below
- “10 Gigabit Ethernet Modules with XENPAK Optics” on page 5-2
- “Cleaning the Fiber Optic Connectors” on page 5-4
- “Cabling 10 Gigabit Ethernet Modules” on page 5-4
- “Port LEDs” on page 5-5
- “Troubleshooting Network Connections” on page 5-5
- “Upgrading an FPGA on a 10 Gigabit Ethernet Module” on page 5-9

1-Port 10 Gigabit Ethernet Module

The 1-port 10 Gigabit Ethernet module provides the following types of 10 Gigabit Ethernet interfaces:

- 1310nm serial for single-mode fiber – part number B10Gx-LR
- 1510nm serial for single-mode fiber – part number B10Gx-ER

Figure 5.1 shows the front panel of a 1-port 10 Gigabit Ethernet module.

Figure 5.1 Front panel of 1-port 10 Gigabit Ethernet module



This module provides one 10 Gigabit interface. The interfaces operate at full duplex. For the serial port types listed above, use the matching fiber type with an SC connector. For example, if you are using the 1310nm serial module for single-mode fiber, attach a 1310nm single-mode fiber cable that has an SC connector.

Foundry 10 Gigabit Ethernet modules are compliant with the IEEE 802.3ae 10-Gigabit Ethernet standard.

System Requirements

The 1-port 10 Gigabit Ethernet modules are supported in the following products:

- BigIron 4000, BigIron 8000, and BigIron 15000
- NetIron 400, NetIron 800, and NetIron 1500 Metro routers

- FastIron 400, FastIron 800, and FastIron 1500

You need an M2, M3, or M4 management module that supports redundancy or a VM1 management module.

NOTE: To use a 1-port 10 Gigabit Ethernet module in a chassis managed by a VM1, you must be running software release 07.6.01 or later.

Hardware on the 1-Port 10 Gigabit Ethernet Module

Each 1-port 10 Gigabit Ethernet module uses a single 10 Gigabit Ethernet MAC controller, separate transmit and receive controllers, and five Field-Programmable Gate Arrays (FPGAs). The FPGAs enable you to easily implement architecture upgrades without changing the hardware. The software includes a CLI command you can use to upgrade the FPGAs if needed. See “Upgrading an FPGA on a 10 Gigabit Ethernet Module” on page 5-9.

The 10 Gigabit Ethernet standard does not include link auto-negotiation. A Foundry 10 Gigabit Ethernet port is unable to detect a link failure at the other end of the link if the failure is on the receive side of the remote link. However, the Foundry 10 Gigabit Ethernet port can detect a link failure if the failure occurs on the transmit side of the remote link.

NOTE: The non-XENPAK 10 Gigabit Ethernet module can function in the same chassis with Foundry’s XENPAK-based 10 Gigabit Ethernet modules.

Features Not Supported on the 1-Port 10 Gigabit Ethernet Module

The following features are not supported on the non-XENPAK 10 Gigabit Ethernet module in the current release:

- Rate limiting
- IronClad QoS
- IP multicast on tagged ports
- Jumbo packets, if the module is used in a chassis that contains IronCore (non-JetCore) modules. When you use the module in a chassis containing JetCore modules, jumbo packets are supported.

Replacing the Optics on the 1-Port 10 Gigabit Ethernet Module

If you need to replace the optics on the non-XENPAK 10 Gigabit Ethernet module, contact Foundry Networks.

10 Gigabit Ethernet Modules with XENPAK Optics

Software release 07.6.03 introduced support for 1-port and 2-port 10 Gigabit Ethernet modules with XENPAK optics.


Figure 5.2 shows the front panel of a 1-port 10 Gigabit Ethernet Module.

Figure 5.2 Front panel of 1-port 10 Gigabit Ethernet module



Figure 5.3 shows the front panel of a 2-port 10 Gigabit Ethernet Module.

Figure 5.3 Front panel of 2-port 10 Gigabit Ethernet module



The 10 Gigabit Ethernet interfaces operate at full duplex. The module uses GBIC-like XENPAK Multisource Agreement (MSA) optics. The XENPAK optics are hot-swappable, allowing you to change the optics without removing the module from the chassis.

The following kinds of XENPAK optics are supported:

- 1310nm serial for single-mode fiber
- 1510nm serial for single-mode fiber

For the XENPAK optic types listed above, use the matching fiber type with an SC connector. For example, if you are using the 1310nm serial module for single-mode fiber, attach a 1310nm single-mode fiber cable that has an SC connector.

System Requirements

The XENPAK-based 10 Gigabit Ethernet modules are supported in the following products:

- BigIron 4000, BigIron 8000, and BigIron 15000.
- FastIron 400, FastIron 800, and FastIron 1500.
- NetIron 400, NetIron 800, and NetIron 1500.

You need an M4 or VM1 management module.

Hardware on the XENPAK-Based 10 Gigabit Ethernet Module

Each port on the XENPAK-based 10 Gigabit Ethernet modules has a 10 Gigabit Ethernet MAC controller and separate transmit and receive controllers. The modules have two kinds of FPGAs. The FPGAs enable you to easily implement architecture upgrades without changing the hardware. The software includes a CLI command you can use to upgrade the FPGAs if needed. See “Upgrading an FPGA on a 10 Gigabit Ethernet Module” on page 5-9.

NOTE: The XENPAK-based 10 Gigabit Ethernet modules can function in the same chassis with the non-XENPAK-based 1-port 10 Gigabit Ethernet modules.

Features Not Supported on XENPAK-based 10 Gigabit Ethernet Modules

The XENPAK-based 10 Gigabit Ethernet modules support all of the applicable Layer 2 and Layer 3 features in software release 07.6.03 and earlier. The following features are not supported in the current release:

- Rate limiting
- Jumbo packets, if the module is used in a chassis that contains IronCore (non-JetCore) modules. When you use the module in a chassis containing JetCore modules, jumbo packets are supported.

Removing and Installing XENPAK Optics

You can remove a XENPAK optic from a 10 Gigabit Ethernet module and replace it with a new one while the Foundry device is powered on and running.

Before performing either of these tasks, have the following on hand:

- An electrostatic discharge (ESD) wrist strap

WARNING: For safety reasons, the ESD wrist strap should contain a series 1 meg ohm resistor.

- The protective covering that you removed from the port connectors when you initially installed the XENPAK optic
- The new XENPAK optic (if you are installing one)
- A small flathead screwdriver

Removing a XENPAK Optic

To remove a XENPAK optic from a 10 Gigabit Ethernet module, do the following:

1. Put on the ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Disconnect the two fiber cable connectors from the port connectors.
3. Insert the protective covering into the port connectors.
4. Using the flathead screwdriver if necessary, loosen the two thumbscrews on the ends of the XENPAK optic.
5. Pull the XENPAK optic out of the port, and place it in an anti-static bag for storage if desired.
6. Install a new XENPAK optic in the module, if necessary. For information about performing this task, see "Installing a XENPAK Optic" below.

Installing a XENPAK Optic

To install a XENPAK optic in a 10 Gigabit Ethernet module, do the following:

1. Put on the ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.
2. Remove the new XENPAK optic from its protective packaging.
3. Gently insert the XENPAK optic into the module until it clicks into place. The XENPAK optics are keyed to prevent incorrect insertion.
4. Secure the XENPAK optic by tightening the two thumb-screws. If desired, you can further tighten the thumb-screws using the flathead screwdriver.

Cleaning the Fiber Optic Connectors

To avoid problems with the connection between the fiber-optic module connectors and the fiber cable connectors, Foundry strongly recommends cleaning both connectors each time you disconnect and reconnect them. In particular, dust can accumulate in the connectors and cause problems such as reducing the optic launch power.

To clean the fiber cable connectors, Foundry recommends using a fiber-optic reel-type cleaner. You can purchase this type of cleaner from the following Web site:

http://www.fisfiber.com/fisfiber.com/Home_Page.asp

To clean the fiber-optic module connectors, Foundry recommends using a product that dispenses dust-free air, such as Micro-Blast. You can purchase such a product from the following Web site:

<http://www.microcare.com/product/solvents/PS-50.html>

When cleaning a fiber-optic module connector, do not use unfiltered air from an air compressor, cotton swabs, or other types of swab applicators. These types of products may leave lint or dust in the connector.

Also, when not using a fiber-optic module connector, make sure to keep the protective covering on.

Cabling 10 Gigabit Ethernet Modules

To cable a 10 Gigabit Ethernet module, do the following:

1. Remove the protective covering from the fiber-optic port connectors and store the covering for future use.
2. Before attaching cables to the module, Foundry strongly recommends cleaning the cable connectors and the port connectors. For more information, see "Cleaning the Fiber Optic Connectors".
3. Gently insert the two cable connectors (a tab on each connector should face upward) into the port connectors until the tabs lock into place.

4. Observe the link and active LEDs to determine if the network connections are functioning properly. For more information about the LED indicators, see Table 5.1.

Port LEDs

The LEDs listed in Table 5.1 provide status information for 10 Gigabit Ethernet ports. All types of Foundry 10 Gigabit Ethernet modules use the same port LEDs.

Table 5.1: LEDs for 10 Gigabit Ethernet Ports

LED	Position	State	Meaning
Link	Top	On	Port is connected.
		Off	No port connection exists.
Activity	Bottom	On	Traffic is being transmitted and received on that port.
		Off	No traffic is being transmitted.
		Blinking	Traffic is being transmitted and received on that port.

Troubleshooting Network Connections

After you attach cables to the 10 Gigabit Ethernet modules, you can observe the LEDs to determine if the network connections are functioning properly. Table 5.2 outlines possible abnormal states of each LED, and what to do if an LED indicates an abnormal state.

Table 5.2: Network Connection-Related LED States

LED	Abnormal State	Meaning/Action
Link	Off	<p>A link is not established with the remote port. You can do the following:</p> <ul style="list-style-type: none"> Verify that the connection to the other network device has been properly made. Also, make certain that the other network device is powered on and operating correctly. Verify that the transmit port on the Foundry device is connected to the receive port on the other network device, and that the receive port on the Foundry device is connected to the transmit port on the other network device. If you are not certain, remove the two cable connectors from the port connector and reinsert them in the port connector, reversing their order. Dust may have accumulated in the cable connector or port connector. For information about cleaning the connectors, see "Cleaning the Fiber Optic Connectors" on page 5-4. If the other actions don't resolve the problem, try using a different port or a different cable.

Table 5.2: Network Connection-Related LED States

LED	Abnormal State	Meaning/Action
Activity	Off for an extended period.	<p>The port is not transmitting or receiving user packets. You can do the following:</p> <ul style="list-style-type: none"> • Check the Link LED to make sure the link is still established with the remote port. If not, take the actions described in the Meaning/Action column for the Link LED. • Verify that the port has not been disabled through a configuration change. You can use the CLI to do this. If you have configured an IP address on the device, you also can use the Web management interface or IronView Network Manager.

If a problem persists after taking these actions, contact Foundry Technical Support.

Link Fault Signaling (LFS)

NOTE: This feature is supported in software releases 07.6.02 and later.

Link Fault Signalling (LFS) is a physical layer protocol that enables communication on a link between two 10 Gigabit Ethernet devices. When configured on a Foundry 10 Gigabit Ethernet port, the port can detect and report fault conditions on transmit and receive ports.

Foundry's 10 Gigabit Ethernet devices include the following:

- First generation device:
 - 1-port 10 Gigabit Ethernet Module
- Second generation devices:
 - 2-port 10 Gigabit Ethernet Module with XENPAK optics

Foundry introduced LFS in software release 07.6.02, thereby enabling Foundry's 10 Gigabit Ethernet devices to communicate critical link status information.

In release 07.6.03, Foundry's implementation of LFS became compliant with the IEEE standard, however, you could enable it only between two First generation 10 Gigabit Ethernet devices, or between two Second generation 10 Gigabit Ethernet devices. LFS support was not available between First generation and Second generation 10 Gigabit Ethernet devices.

In software release 07.6.04 and later, Foundry's software supports LFS among all 10 Gigabit Ethernet devices, including LFS support between First and Second generation devices.

Determining the 10 Gigabit Ethernet Module Installed in Your System

To determine which 10 Gigabit Ethernet Module your system is running, enter the **show flash** command at the Privileged EXEC level of the CLI. The output of the **show flash** command will help you determine which module is installed, as follows:

- First generation devices (1-port 10 Gigabit Ethernet Modules) will show five FPGA images.
- Second generation devices (1-port and 2-port 10 Gigabit Ethernet Modules with XENPAK optics) will display two FPGA images.

Configuring Link Fault Signalling

Configuration procedures for LFS differ depending on your 10 Gigabit Ethernet hardware configuration. See "Determining the 10 Gigabit Ethernet Module Installed in Your System" on page 5-6.

To configure LFS, follow the appropriate procedures, below.

Enabling LFS Between Two First Generation Devices or Between Two Second Generation Devices

To enable LFS between two First generation 10 Gigabit Ethernet devices, or between two Second generation 10 Gigabit Ethernet devices, enter commands such as the following on both ends of the link:

```
BigIron(config)# interface e 1/1
BigIron(config-if-e1000-1/1)# link-fault-signal
```

Syntax: [no] link-fault-signal

Use the no form of the command to disable LFS.

LFS is OFF by default.

Enabling LFS Between a First Generation Device and a Second Generation Device

To enable LFS between a First generation 10 Gigabit Ethernet device and a Second generation 10 Gigabit Ethernet device, enter commands such as the following:

1. On the First generation 10 Gigabit Ethernet device, enter commands such as the following:

```
BigIron(config)# interface e 1/1
BigIron(config-if-e1000-1/1)# link-fault-signal
```

2. On the Second generation 10 Gigabit Ethernet device, enter commands such as the following:

```
BigIron(config)# interface e 1/1
BigIron(config-if-e1000-1/1)# link-fault-signal legacy
```

Syntax: [no] link-fault-signal legacy

Use the no form of the command to disable LFS.

LFS is OFF by default.

Remote Fault Notification (RFN) on Fiber Connections

For fiber-optic connections, you can optionally configure a transmit port to notify the receive port on the remote device whenever the transmit port becomes disabled.

When you enable this feature, the transmit port notifies the remote port whenever the fiber cable is either physically disconnected or has failed. When this occurs and the feature is enabled, the device disables the link and turns OFF both LEDs associated with the ports.

By default, Remote Fault Notification (RFN) is disabled. In this case, if the transmit port becomes physically disabled or fails, the link still appears as though it is enabled and the LEDs for both ports remain ON.

Configuration Notes

- RFN is supported in software releases 07.6.05 and later, and in Service Provider releases 09.1.00 and later.
- This feature is only available for Gigabit Ethernet Fiber ports. It is not available for 10/100 ports and Gigabit Ethernet Copper ports.

RFN Enhancements in 07.8.00

Software releases prior to 07.8.00 provide support for Remote Fault Notification (RFN) on individual interfaces only.

Software release 07.8.00 provides enhanced support for RFN by making it available:

- Globally, on the entire device
- On a trunk group

- On an individual interface

Configuration Considerations

- Configuring RFN on a global basis enables RFN on every port and takes precedence over RFN configurations on trunk groups and on individual interfaces.
- When RFN is enabled both globally and on a trunk group, and you remove a secondary port in a trunk group, the secondary port becomes disabled and the primary port in the trunk group remains enabled.
- When RFN is enabled both globally and on a trunk group, the trunk group behaves as follows:
 - If a primary port in a trunk group fails, only the primary port reports an RFN error. None of the other (secondary) ports in the trunk group will report an RFN error.
 - If a secondary port in a trunk group fails, the secondary port reports an RFN error even though the primary port has not failed and is still enabled.
- When RFN is not enabled globally, but is enabled on a trunk group, the trunk group behaves as follows:
 - If the primary port in a trunk group fails, the primary port will report an error. All other trunk ports will not report errors and will remain enabled.
 - If a secondary port in a trunk group fails and the primary port is still enabled and has not failed, the secondary trunk port will report an error and all other trunk ports will remain enabled.

Enabling Remote Fault Notification

To enable RFN globally, on the entire device, enter the following command:

```
BigIron(config)# gig-default auto-gig rfn
```

To enable RFN on a trunk group, enter the command on the primary port of the trunk group, as shown in the following example:

```
BigIron(config)# int e 1/1  
BigIron(config-if-e1000-1/1)# gig-default auto-gig rfn
```

To enable RFN on an individual interface, enter commands such as the following:

```
BigIron(config)# int e 1/2  
BigIron(config-if-e1000-1/2)# gig-default auto-gig rfn
```

Syntax: [no] gig-default auto-gig rfn

To disable RFN after enabling it, use the **no** parameter with the command.

Viewing Which Fiber Ports Have RFN Enabled

Use the **show run** command to view which fiber ports have RFN enabled. The following shows an example output.

```
BigIron Router# sh run

Building configuration...
Current configuration : 349 bytes
ver 07.7.00b1T51

module 1 bi-jc-8-port-gig-m4-management-module
module 2 bi-jc-16-port-gig-fiber-module
module 3 bi-jc-8-port-gig-module
module 8 bi-jc-16-port-gig-copper-module

no global-stp

no spanning-tree

ip address 2.2.2.2 255.255.255.255

auto-cam-repaint

pram-write-retry

interface ethernet 1/2

    gig-default auto-gig rfn
```

Upgrading an FPGA on a 10 Gigabit Ethernet Module

NOTE: If an upgrade is required for any of the FPGA files, you must upgrade all the FPGA files.

1. Complete the upgrades of the boot code and flash code, if required.
2. Enter commands such as the following at the Privileged EXEC level of the CLI:

```
BigIron# 10gig copy tftp 10.10.10.10 rxbmgr.bin
BigIron# 10gig copy tftp 10.10.10.10 rxpp.bin
BigIron# 10gig copy tftp 10.10.10.10 txaccum.bin
BigIron# 10gig copy tftp 10.10.10.10 txpp.bin
BigIron# 10gig copy tftp 10.10.10.10 ageram.bin
```

Syntax: 10gig copy tftp | slot1 | slot2 flash <ip-addr> <filename> [module <slotnum>]

where:

- **tftp | slot1 | slot2** – specifies the location of the FPGA file. The **tftp** parameter indicates that the file is on a TFTP server. The **slot1** and **slot2** parameters indicate that the file is on a PCMCIA flash card. Specify **slot1** if the file is on the flash card in PCMCIA slot 1. Specify **slot2** if the file is on the flash card in PCMCIA slot 2.
- <ip-addr> – specifies the IP address of the TFTP server, if you specify **tftp**.
- <filename> – specifies the FPGA file name.

NOTE: You can store and copy the FPGA files using any valid filename; however, Foundry recommends that you use the file names listed in the “Software Image Files” section of the release notes. The device uses information within the files to install them in the correct FPGAs. The **show flash** command lists the FPGAs. For an example of the **show flash** output, see “Displaying the Installed FPGA Revisions” on page 5-10.

- **module** <slotnum> – optionally, specifies the modules on which you want to install the upgrade. If you do not specify a slot number, the command upgrades the FPGA on all 10 Gigabit Ethernet modules in the chassis.
3. Reload the software by entering one of the following commands:
- **reload** (this command boots from the default boot source, which is the primary flash area by default)
 - **boot system flash primary | secondary**

NOTE: The **show flash** command will list the new FPGA code versions but the new versions do not take effect until you reload the software.

Displaying the Installed FPGA Revisions

To display the software versions installed in flash memory on the management module and the FPGA versions installed on the 10 Gigabit Ethernet modules, enter the following command:

```
BigIron# show flash
Active management module:
Code Flash Type: AMD 29LV033C, Size: 64 * 65536 = 4194304, Unit: 4
Boot Flash Type: AMD 29LV040B, Size: 8 * 65536 = 524288
Compressed Pri Code size = 2813111, Version 07.6.03b130T53 (b2r07603b130.bin)
Compressed Sec Code size = 2799367, Version 07.6.03b79T53 (b2r07603b79.bin)
Maximum Code Image Size Supported: 6815232 (0x0067fe00)
Boot Image size = 275128, Version 07.06.03 (m2b07603.bin)
Monitor Image Version 4, for DRAM size 268435456
Used Configuration Flash Size=2092, Max Configuration Flash Size=524288.

10 GIG module slot 2
X10G RXBMGR FGPA      version: 80  revision: 6          2001/11/15  15:38:43
X10G RXPP FGPA       version: 81  revision: 13         2002/06/17  17:28:53
X10G TXACCUM FGPA    version: 82  revision: 6          2001/12/12  18:51:43
X10G TXPP FGPA       version: 83  revision: 11         2002/08/16  19:15:36
X10G AGERAM FGPA     version: 84  revision: 4          2001/10/26  19:53:24

2x10 GIG module slot 3
2X10G XTM FGPA       version: 89  revision: 34         2003/02/07  01:35:52
2X10G XPPE FGPA     version: 88  revision: 34         2003/02/25  05:08:30
```

Syntax: show flash

The boot code and flash code versions are listed in the "Compressed Pri Code size", "Compressed Sec Code size", and "Boot Image size" lines of the display. The FPGA versions are listed separately for each 10 Gigabit Ethernet module. In this example, the chassis contains a non-XENPAK-based 1-port 10 Gigabit Ethernet module in slot 2, and a XENPAK-based 2-port 10 Gigabit Ethernet module in slot 3. Notice that the FPGA names match the file names listed in the release notes.

Chapter 6

Using Power Over Ethernet Modules

This chapter describes the JetCore 24-Port Power Over Ethernet (POE) modules and how to configure them.

Power over Ethernet on JetCore Chassis devices is compliant with the IEEE 802.3af specification. Support for Power over Ethernet consists of the JetCore J-B24E-POE and J-F24E-POE modules, which can supply power and data to POE-enabled devices, and the RPS-POE Shelf, which is a separately installed power shelf used for supplying power to the J-B24E-POE and J-F24E-POE modules.

The 802.3af specification provides the standard for delivering power over existing network cabling infrastructure, enabling multicast-enabled full streaming audio and video applications for converged services, such as Voice over IP (VoIP), WLAN access points, IP surveillance cameras, and other IP powered devices. POE technology eliminates the need for an electrical outlet and dedicated UPS near IP-powered devices. When the POE modules are installed in a Chassis device, power can be consolidated and centralized in the wiring closets, improving the reliability and resiliency of the network.

This chapter consists of the following sections:

- “Power Over Ethernet Hardware Description”, below describes the J-B24E-POE and J-F24E-POE modules and the RPS-POE power shelf.
- “Installing the POE Hardware” on page 6-4 provides instructions for installing the J-B24E-POE or J-F24E-POE modules in a Chassis device and installing RPS-POE power shelf in an equipment rack.
- “Configuring the POE Software” on page 6-11 describes how to configure Power over Ethernet on J-B24E-POE or J-F24E-POE ports using CLI commands.
- “Displaying Power over Ethernet Information” on page 6-13 shows how to display information about POE-enabled ports on a Foundry device.


Power Over Ethernet Hardware Description

This section describes the features of the J-B24E-POE and J-F24E-POE modules and the RPS-POE power shelf.

J-B24E-POE and J-F24E-POE Modules

Figure 6.1 shows the front panel of the J-B24E-POE module.

Figure 6.1 Front panel of the J-B24E-POE Power Over Ethernet Module



The J-B24E-POE and J-F24E-POE modules provide 24 RJ-45 connectors for Cat5 cabling. You can connect each port to a 10 Mbps or 100 Mbps segment. The ports automatically detect the speed of the network and configure themselves accordingly. The pin assignments and the status LEDs are the same as the ones for the 10/100 Mbps ports on other Foundry modules. The 10/100 ports on the J-B24E-POE and J-F24E-POE modules can detect 802.3af compatible IP devices and provide power accordingly.

The J-B24E-POE and J-F24E-POE modules contain a Foundry-proprietary 4-pin connector, labelled POE-IN. The connector is used for connecting a cable between the module and the RPS-POE power shelf. In order for the module to provide power to POE-enabled devices, you must connect the POE-IN connector to one of the POE output connectors on the RPS-POE power shelf. Each J-B24E-POE or J-F24E-POE module installed in the Chassis device must be connected to the RPS-POE power shelf in this way.

The 10/100 Mbps ports provide status information using the LEDs listed in the following table:

Table 6.1: LEDs on J-B24E-POE and J-F24E-POE modules

LED	Position	State	Meaning
Link/Activity	Left LED above each 10/100 port	On	Link is up.
		Off	Link is down.
		Blinking	Port is transmitting or receiving.
FDX	Right LED above each 10/100 port	On	Full-duplex connection found or configured. Note: This LED also is lit if you configure the port to 10 Mbps full-duplex or 100 Mbps full-duplex. This is true even when no link is present.
		Off	Half-duplex connection or no port connection exists.
		Blinking	Collisions are being detected.
Port Power	Below each row of 10/100 ports	On	The port is enabled, a power-consuming device has been detected, and the module is supplying power to the device.
		Off	The port is not providing in-line power.


Table 6.1: LEDs on J-B24E-POE and J-F24E-POE modules (Continued)

LED	Position	State	Meaning
POE Power	Far left side of module	Green	The RPS-POE shelf is supplying –48V power to the module in an acceptable range (between –44V and –57V).
		Yellow	The RPS-POE shelf is supplying –48V power to the module outside of the acceptable range (between –44V and –57V).
		Off	No –48V power supply is detected. The module may be disconnected from the RPS-POE shelf, or the RPS-POE shelf is not supplying power to the module.

RPS-POE Power Shelf

Figure 6.2 shows the front panel of the RPS-POE power shelf.

Figure 6.2 Front panel of the RPS-POE power shelf



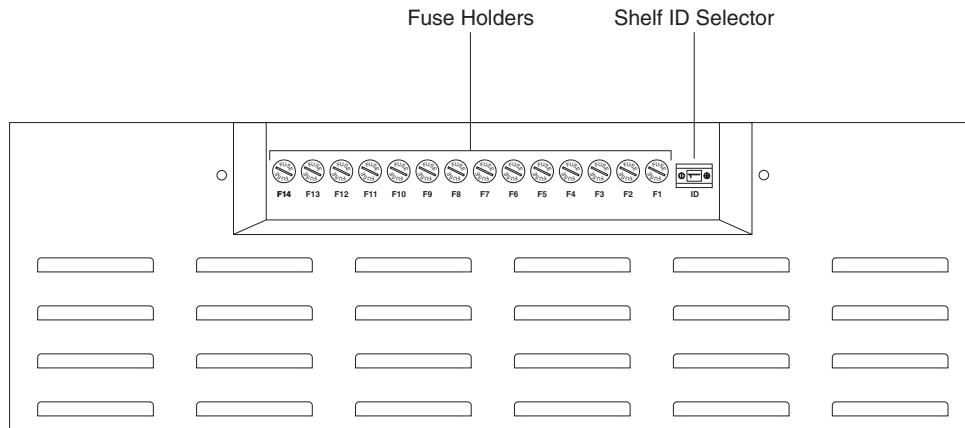
The RPS-POE power shelf provides power to the J-B24E-POE or J-F24E-POE modules installed in a JetCore Chassis device. Each RPS-POE power shelf can supply power to as many as 14 J-B24E-POE or J-F24E-POE modules. To provide power to a J-B24E-POE or J-F24E-POE module, you connect a power cable (Foundry-supplied) from one of the POE output connectors on the RPS-POE power shelf to the POE-IN connector on the J-B24E-POE or J-F24E-POE module. Each POE output connector has a Power LED. When power is being supplied to a POE module using the connector, the connector's Power LED is green.

The RPS-POE shelf contains six slots for AC power supplies. Each power supply can provide power to up to three J-B24E-POE or J-F24E-POE modules. Foundry recommends that you install additional power supplies to provide redundancy. Each AC power supply has a maximum power rating of 1140 Watts. If six power supplies are installed in the RPS-POE power shelf, then the RPS-POE power shelf can provide up to 6,840 Watts. This would provide sufficient power for J-B24E-POE or J-F24E-POE modules connected to all 14 of the RPS-POE power shelf's POE output connectors, with each of the J-B24E-POE or J-F24E-POE modules providing power to 24 power-consuming devices, each drawing maximum power (15.4 Watts) simultaneously.

Each AC power supply has an LED on its faceplate. If the LED is on (green), the power supply is providing power to the chassis components. If the LED is amber or off, the power supply is not providing power to the POE modules in the Chassis device.

Figure 6.3 shows the rear panel of the RPS-POE power shelf.

Figure 6.3 Rear panel of the RPS-POE power shelf



On the rear panel of the RPS-POE shelf are fuses for each of the POE output connectors. If the power LED for one of the POE output connectors goes out, you may need to replace the corresponding fuse. See “Replacing a Fuse in the RPS-POE Shelf” on page 6-11 for information on how to do this.

If you are using more than one RPS-POE shelf to supply power to POE modules in a single Chassis device, then you must assign an ID to each RPS-POE shelf. To do this, use the shelf ID selector to assign each RPS-POE shelf a unique ID. See “Setting the ID of the RPS-POE Shelf (If Necessary)” on page 6-10.

Installing the POE Hardware

This section describes how to install a J-B24E-POE or J-F24E-POE module in a Foundry Chassis device and how to install the RPS-POE shelf. In addition, this section contains a procedure for setting the ID of the RPS-POE shelf if more than one RPS-POE shelf will provide power to a single Foundry Chassis device, as well as a procedure for replacing a fuse on the RPS-POE shelf.

Installing a J-B24E-POE or J-F24E-POE Module in a Foundry Chassis Device

Before installing a J-B24E-POE or J-F24E-POE module in a Foundry Chassis device, have the following on hand:

- An electrostatic discharge (ESD) wrist strap

WARNING: For safety reasons, the ESD wrist strap should contain a series 1 meg ohm resistor.

- A large flat-head screwdriver

To install a J-B24E-POE or J-F24E-POE module in a Foundry Chassis device, do the following:

1. Put on an ESD wrist strap and attach the clip end to a metal surface (such as an equipment rack) to act as ground.

WARNING: To avoid risk of shock, do not attach the clip end to the air flow panel of a power supply.

2. Remove the blank face plate from the slot in which the module will be installed. Place the blank face plate in a safe place for future use.
3. Remove the module from its packaging.
4. Insert the module into the chassis slot and slide the card along the card guide until the card ejectors on the front of the module touch the chassis.

NOTE: Modules for the 8-slot and 15-slot Chassis devices slide in vertically with port number 1 at the top. Modules for the 4-slot Chassis devices slide in horizontally with port number 1 on the left.

5. Push the ejectors toward the center of the module until they are flush with the front panel of the module. The module will be fully seated in the backplane.
6. Tighten the two screws at either end of the module.

CAUTION: To provide additional safety and proper airflow to the device, make sure that slot cover plates are installed on all chassis slots that do not have either a module or power supply installed.

NOTE: If installing a module into a slot *previously occupied by a different type of module*, you must use the CLI to configure the new module (use the CLI command **module** <slot-num> <module-type>) and then use the **write memory** command to save the configuration and the **reload** command to reset the device. See “Swapping Modules (Chassis devices only)” in the *Foundry Switch and Router Installation and Basic Configuration Guide*. If the slot has never contained a module or you are swapping in exactly the same type of module, you do not need to enter these commands.

Installing the RPS-POE Shelf

Installing the RPS-POE Shelf consists of the following tasks:

- Attaching mounting brackets to the RPS-POE shelf
- Mounting the RPS-POE Shelf in a rack
- Installing power supplies in the RPS-POE shelf
- Cabling the module and the RPS-POE shelf
- Applying power to the RPS-POE shelf

NOTE: If you are using more than one RPS-POE shelf to supply power to POE modules in a single Foundry Chassis device, then you must assign an ID to each RPS-POE shelf. To do this, use the shelf ID selector to assign each RPS-POE shelf a unique ID. See “Setting the ID of the RPS-POE Shelf (If Necessary)” on page 6-10 for information on how to do this.


Attaching Mounting Brackets to the RPS-POE Shelf

The RPS-POE shelf ships with a rack mount kit. The kit includes two L-shaped mounting brackets and mounting screws.

The sides of the RPS-POE shelf chassis have three sets of screw holes: one set for attaching the mounting brackets close to the chassis front, another set for attaching the brackets toward the chassis center, and another set for attaching the brackets close to the chassis rear.

Attach the mounting brackets to the sides of the chassis as illustrated in Figure 6.4. Note that the narrow portion of the keyhole slot is up and the wide portion of the slot is down.

Figure 6.4 Attaching mounting brackets to the RPS-POE Shelf



NOTE: Figure 6.4 shows the mounting brackets being attached to the RPS-POE shelf in the front position. The procedure for attaching the brackets in the center or rear position is exactly the same.

Mounting the RPS-POE Shelf in a Rack

Keep the following in mind when mounting an RPS-POE shelf in a rack:

WARNING: Do not use the handles on the power supply units to lift or carry an RPS-POE shelf.


WARNING: Make sure the rack or cabinet housing the device is adequately secured to prevent it from becoming unstable or falling over.

WARNING: Mount the devices you install in a rack or cabinet as low as possible. Place the heaviest device at the bottom and progressively place lighter devices above.


For each RPS-POE shelf that you install in a rack, you must provide four screws on which to mount and secure the chassis.

To mount an RPS-POE shelf in a rack, do the following:

1. Determine the position of the RPS-POE shelf in the rack; for example, above the Chassis device where the J-B24E-POE or J-F24E-POE module is installed.
2. Position the two screws for the RPS-POE shelf according to the spacings of the keyhole slots on the mounting brackets as shown in Figure 6.5. Do not secure the screws completely; leave approximately 1/4 in of clearance between the back of the screw head and the rack.

Figure 6.5 Positioning the screws in a rack

3. Mount the RPS-POE shelf in the rack as shown in Figure 6.6. Slip the wide portion of each keyhole slot over the corresponding screw in the rack.

Figure 6.6 Mounting the RPS-POE shelf in a rack

4. Slide the RPS-POE shelf down so that the screw heads are in the narrow portion of the keyhole slots.
5. Tighten the screws to secure the RPS-POE shelf in place.

Installing a Power Supply in the RPS-POE Shelf

You install the power supplies starting in the leftmost power supply slots of the RPS-POE shelf. You need a small Phillips or flat-head screwdriver to perform this task.

WARNING: The front panel of a power supply includes a handle that locks the power supply in the RPS-POE shelf. This handle is a locking mechanism only and should not be used to lift and carry the power supply. You may sustain physical injury or harm if you attempt to lift and carry a power supply using the locking handle.

To install a power supply in the RPS-POE shelf, do the following:

1. Remove the blank power supply faceplate, and expose the empty power supply slot.
2. Remove the power supply from its packaging.
3. Insert the power supply into the empty power supply slot, using the guides provided on either side of the slot.


CAUTION: Carefully follow the mechanical guides on each side of the power supply slot and make sure the power supply is properly inserted in the guides. Never insert the power supply upside down.

4. After the power supply is fully inserted, push the power supply front panel toward the back of the RPS-POE shelf until the front panel is flush with the rest of the chassis. This action causes the power supply connector to lock into the backplane connector.
5. Gently pull the handle on the power supply front panel upward and toward the top of the power supply front panel. This action locks the power supply in place.
6. Use the screwdriver to tighten the two screws on either side of the power supply front panel.
7. Repeat Step 1 through Step 6 for each power supply to be installed in the RPS-POE shelf.

Cabling the J-B24E-POE or J-F24E-POE Module and the RPS-POE Shelf

After you have installed the J-B24E-POE or J-F24E-POE module in the Chassis device and the RPS-POE shelf in the rack, connect the Foundry-supplied 4-pin POE cable between the module and the RPS-POE shelf. The male end of the cable attaches to the one of the POE connectors on the RPS-POE shelf, and the female end of the cable connects to the POE-IN connector on the J-B24E-POE or J-F24E-POE module. The connectors are keyed to prevent improper insertion of the cable.

Figure 6.7 illustrates a cable connection between the RPS-POE shelf and a J-B24E-POE module.

Figure 6.7 POE cable attached between the RPS-POE shelf and a J-B24E-POE module in a BigIron 8000

NOTE: When you connect the cable to the connector, it is locked into place. To remove the cable from the connector, grasp the cable by its metal sheath and pull it from the connector.

Applying Power to the RPS-POE Shelf

To apply power to an RPS-POE shelf, attach one end of a Foundry-supplied AC power cord into the front of an installed power supply and insert the other end into a 115V/120V wall outlet. Do this for each installed power supply.

WARNING: If the installation requires a different power cord than the one supplied with the device, make sure you use a power cord displaying the mark of the safety agency that defines the regulations for power cords in your country. The mark is your assurance that the power cord can be used safely with the device.

CAUTION: Ensure that the device does not overload the power circuits, wiring, and over-current protection. To determine the possibility of overloading the supply circuits, add the ampere (amp) ratings of all devices installed on the same circuit as the device. Compare this total with the rating limit for the circuit. The maximum ampere ratings are usually printed on the devices near the input power connectors.

CAUTION: Foundry recommends using a separate branch circuit for each AC power cord, which provides redundancy in case one of the circuits fails.

NOTE: If the wall outlet is not rated 115/120V and 20A, stop and get the appropriate cable for the outlet. Make sure you obtain a power cord displaying the mark of the safety agency that defines the regulations for power cords in your country. The mark is your assurance that the power cord can be used safely with the device.


NOTE: The wall outlet should be installed near the equipment and should be easily accessible.

Setting the ID of the RPS-POE Shelf (If Necessary)

If you are using more than one RPS-POE shelf to supply power to J-B24E-POE or J-F24E-POE modules in a single Chassis device, each RPS-POE shelf must have a unique identifier. The identifier for the RPS-POE shelf is set with the shelf ID selector, located on the rear panel of the device.


Figure 6.8 shows the location of the shelf ID selector on the rear panel of the RPS-POE shelf.

Figure 6.8 Location of the shelf ID selector on the RPS-POE power shelf rear panel



By default each RPS-POE shelf has an ID of 1. To change the ID of the RPS-POE power shelf, use a small stylus (such as a ballpoint pen) to press one of the buttons on either side of the shelf ID selector until the desired ID appears in the window. The button on the right advances the ID, and the button on the left reverses the ID.

Figure 6.9 Shelf ID selector on the RPS-POE shelf



Replacing a Fuse in the RPS-POE Shelf

On the rear panel of the RPS-POE shelf are fuses for each of the POE output connectors (see Figure 6.8 on page 6-10). If the power LED for one of the POE output connectors goes out while the RPS-POE shelf is powered on, you may need to replace the corresponding fuse. The RPS-POE shelf uses an 8 amp fuse.

To remove a fuse from the RPS-POE shelf, use a small flathead screwdriver to twist the fuse holder so that it comes out of its housing in the chassis.

Replace the old fuse with a new 8 amp fuse, then place the fuse holder back into its housing in the RPS-POE shelf chassis. Using the flathead screwdriver, press the fuse holder into its housing in the chassis and twist the fuse holder to lock it in place.

Configuring the POE Software

This section describes how to configure Power over Ethernet on J-B24E-POE or J-F24E-POE ports using CLI commands. It contains the following topics:

- “Enabling Power Over Ethernet”, below
- “Specifying the Power Limit for a Port” on page 6-12
- “Assigning Priority to POE Ports” on page 6-12

Enabling Power Over Ethernet

By default, Power over Ethernet is enabled on the J-B24E-POE or J-F24E-POE ports for 802.3af-compliant devices and disabled for 802.3af non-compliant devices. You can manually disable or enable Power over Ethernet for 802.3af-compliant or 802.3af non-compliant devices connected to the port, or you can configure the Foundry device to automatically detect the type and class of the power-consuming devices connected to the port.

Disabling and Re-enabling Power Over Ethernet For 802.3af-Compliant Devices

On the J-B24E-POE or J-F24E-POE modules, Power over Ethernet is enabled by default for 802.3af-compliant devices connected to the port. To manually disable Power over Ethernet on an interface, enter commands such as the following:

```
BigIron# interface e 3/11
BigIron(config-if-e100-3/11)# no inline power
```

After disabling Power over Ethernet on a port, you can re-enable it by entering commands such as the following:

```
BigIron# interface e 3/11
BigIron(config-if-e100-3/11)# inline power
```

Syntax: [no] inline power

NOTE: When you re-enable Power over Ethernet on a port, the power-consuming device connected to the port is power-cycled. Consequently, you may want to enter this command only when no power-consuming device is connected to the port.

NOTE: If you move a power-consuming device from one port to another, do not connect the device to a new port until the detection status of the old port reflects the current status of the port. The detection status of a port can be displayed with the **show inline power** command. See page 6-14 for more information.

Enabling Power Over Ethernet for 802.3af Non-Compliant Devices

On the J-B24E-POE or J-F24E-POE modules, Power over Ethernet is disabled by default for 802.3af non-compliant devices connected to the port. You can enable Power over Ethernet for 802.3af non-compliant devices on a per-slot basis.

For example, to configure the ports on a J-B24E-POE or J-F24E-POE module in slot 6 to provide power for devices that are not 802.3af-compliant (for example, legacy devices such as Cisco VoIP phones), enter a command such as the following:

```
BigIron(config)# inline power legacy 6
```

Syntax: [no] inline power legacy <slot>

The following legacy devices are currently supported on the J-B24E-POE or J-F24E-POE modules. Other legacy devices may have been tested with the J-B24E-POE or J-F24E-POE modules after the release of this document. Contact your Foundry account representative about installing legacy devices that are not included in this list.

Table 6.2: Legacy Devices Supported on the J-B24E-POE or J-F24E-POE modules

Legacy Device	Firmware
Cisco IP Phone 7910, 7940, and 7960 Series	Cisco Call Manager version 3.1
Cisco Aironet 350 and 1200 Series Access Point	EnterpriseAP version 12.0
Intel PRO/Wireless 5000 LAN Access Point and PRO/Wireless 5000 Dual Access Point	Version 1.2
Sony SNC-VL10N Video Network Color Camera	Version 1.4.6

NOTE: Although Foundry has attempted to provide accurate information in these materials, Foundry assumes no legal responsibility for the accuracy or completeness of the information. More specific information is available on request from Foundry. Please note that Foundry's product information does not constitute or contain any guarantee, warranty or legally binding representation, unless expressly identified as such in a duly signed writing.

Specifying the Power Limit for a Port

By default, each port on the J-B24E-POE or J-F24E-POE modules can provide up to 16.8 watts of power to each POE power consuming device connected to the switch. You can configure the maximum amount of power that a port can provide to a power consuming device. You can specify from 1 to 15.4 watts of power for each device connected to the switch.

To configure the maximum power level for a port on a J-B24E-POE or J-F24E-POE module, enter commands such as the following:

```
BigIron(config)# interface e 3/11
BigIron(config-if-e100-3/11)# inline power maxpower 5000
```

These commands enable in-line power on Ethernet interface 3/11 and set the POE power level to 5,000 milliwatts (5 watts).

Syntax: [no] inline power maxpower <power level>

where <power level> is the number of milliwatts, between 3000 and 16800. The default is 16800 milliwatts.

Assigning Priority to POE Ports

When the Foundry device is not receiving sufficient power to provide power to all of the POE-enabled ports, it stops providing power to some of the ports. You can specify a priority to one or more of the POE-enabled ports on the module. If the Foundry device is not receiving enough power for all of the POE-enabled ports, power is removed from the higher-priority ports only after power is removed from the lower-priority ports.

To assign a priority to a POE port, enter commands such as the following:

```
BigIron(config)# interface e 3/11
BigIron(config-if-e100-3/11)# inline power priority 1
```

Syntax: [no] inline power priority <number>

The priority <number> can be 1 (critical priority), 2 (high priority), or 3 (low priority). The default for all POE-enabled ports is 3 (low priority).

When the device is not receiving sufficient power to provide power to all of the ports, it first disables power to low-priority ports, starting with the highest-numbered low-priority port on the J-B24E-POE or J-F24E-POE module in the highest-numbered slot. If there is still insufficient power after all of the low-priority ports have been disabled, then the device starts disabling power to high-priority ports, starting with the highest-numbered high-priority port on the J-B24E-POE or J-F24E-POE module in the highest-numbered slot. If there is still insufficient power after all of the high-priority ports have been disabled, then the device starts disabling power to critical-priority ports, starting with the highest-numbered critical-priority port on the J-B24E-POE or J-F24E-POE module in the highest-numbered slot.

If additional power is subsequently applied to the device, then power is enabled on POE ports in reverse order of how it was disabled: critical-priority ports are enabled starting with the lowest-numbered critical-priority port in the lowest-numbered slot, and so on.

Displaying Power over Ethernet Information

To display information about POE-enabled ports on the Foundry device, enter the following command:

```
BigIron# show inline power
```

Port	Detection	Class:mwatts	Power	Enable	Power	Priority
6/1	802.3AF	Class3: 7000	ON		Good	High
6/2	OPEN	Unknown: 0	ON		No	Low
6/3	OPEN	Unknown: 0	ON		No	Low
6/4	OPEN	Unknown: 0	ON		No	Low
6/5	OPEN	Unknown: 0	ON		No	Low
6/6	802.3AF	Class1: 4000	ON		Good	Low
6/7	OPEN	Unknown: 0	ON		No	Low
6/8	OPEN	Unknown: 0	ON		No	Low
6/9	OPEN	Unknown: 0	ON		No	Low
6/10	LEGACY	Unknown: 2100		ON	Good	Critical

Syntax: show inline power [<slot> | <portnum>]

You can display Power over Ethernet information about the ports for a specified slot or for a specified port number. If you do not specify a <slot> or <portnum>, then information is displayed for all POE-enabled ports on the Foundry device.

Table 6.3 describes the output of the **show inline power** command.

Table 6.3: Output of the show inline power command

This Field...	Displays...
Port	The number of each POE port that has been successfully initialized by the software.

Table 6.3: Output of the show inline power command (Continued)

This Field...	Displays...
Detection	<p>Information about the power-consuming device detected on the port. This can be one of the following:</p> <ul style="list-style-type: none"> • PENDING – Detection is in progress. • OPEN – No connection was detected. • SHORT – A non-power-consuming device is connected to the port. • LEGACY – A legacy (non-802.3af compliant) power-consuming device was detected. • 802.3AF – An 802.3af power-consuming device was detected. • OFF – Power over Ethernet capability was manually disabled. • DENIED – A power-consuming device was detected on the port; however, there was insufficient power available to power the device. Power was disabled to the port. <p>Note: If you move a power-consuming device from one port to another, do not connect the device to a new port until the detection status of the old port reflects the current status of the port.</p> <p>For example, a port connected to an 802.3af power-consuming device has a detection status of 802.3AF. If you then disconnect the device from the port, you must wait until the detection status of the port changes to OPEN before connecting the power-consuming device to a new port.</p>
Class:mwatts	<p>The maximum amount of power a powered device receives, as well as the number of milliwatts currently applied. This value can be one of the following:</p> <ul style="list-style-type: none"> • Class0 – This is the default. Requires 15.4 watts maximum. • Class1 – Requires 4 watts maximum • Class2 – Requires 7 watts maximum • Class3 – Requires 15.4 watts maximum • Class4 – Not supported at this time • Unknown – Indicates that the device attached to the port cannot advertise its class. This can happen when the classification process has not been completed, no power-consuming device was detected, or a legacy (non-802.3af compliant) power-consuming device was detected.
Power Enable	<p>Whether Power over Ethernet has been enabled on the port. This value can be one of the following</p> <ul style="list-style-type: none"> • ON – This port has been configured to provide inline power. • OFF – This port has not been configured to provide inline power.
Power	<p>The status of the power provided to the powered device. This value can be one of the following:</p> <ul style="list-style-type: none"> • No – The port is not providing inline power. • Good – Indicates power is being transmitted through the POE port to the power-consuming device connected to it. The incoming –48V power is present and within operating range.

Table 6.3: Output of the show inline power command (Continued)

This Field...	Displays...
Priority	The configured priority for the port (Critical, High, or Low) for power management purposes. If the Foundry device is not receiving enough power for all of the POE-enabled ports, power is removed from the higher-priority ports only after power is removed from lower-priority ports.

To display detailed information about POE configuration on the Foundry device, enter the following command:

```
BigIron# show inline power detail
```

```
External Power Supply Status
+++++
```

```
Fan: ON
```

```
Power
```

```
Supply      Status      Power
```

```
-----
PS1         None         0 W
PS2         None         0 W
PS3         Good        1140 W
PS4         Good        1140 W
PS5         Good        1140 W
PS6         Good        1140 W
Total                          4560 W
```

```
System Inline Power Consumption
```

```
+++++
```

Module Number	POE Capable	Available Power	Used Power
1	No	N/A	N/A
2	No	N/A	N/A
3	Yes	384.0 W	2.5 W
4	No	N/A	N/A
5	No	N/A	N/A
6	Yes	384.0 W	17.1 W
7	No	N/A	N/A
8	No	N/A	N/A
Total:		768.0 W	19.6 W

```
System Inline Power Status
```

```
+++++
```

```
Slot Inline Power Capability
```

```
-----
1  None
2  None
3  None
4  Off
5  None
6  None
7  802.3AF and Legacy
8  None
```

```

Port      Detection      Class:mwatts      Power Enable  Power      Priority
-----
6/1      802.3AF          Class3: 7000      ON           Good      High
6/2      OPEN            Unknown: 0        ON           No        Low
6/3      OPEN            Unknown: 0        ON           No        Low
6/4      OPEN            Unknown: 0        ON           No        Low
6/5      OPEN            Unknown: 0        ON           No        Low
6/6      802.3AF          Class1: 4000      ON           Good      Low
6/7      OPEN            Unknown: 0        ON           No        Low
6/8      OPEN            Unknown: 0        ON           No        Low
6/9      OPEN            Unknown: 0        ON           No        Low
6/10     LEGACY          Unknown: 2100     ON           Good      Critical
    
```

System Total Port Detections

+++++

```

802.3AF Legacy Open Short Off Pending Denied Total
-----
3      1      32      2      2      0      0      48
    
```

Syntax: show inline power detail

The output of the **show inline power detail** command is divided into four sections. The first section displays information about the RPS-POE shelf, and the other three sections display information about the POE ports. The information displayed in the System Inline Power Status section is the same as that described in Table 6.3 on page 6-13. The information in the other sections is described below.

Table 6.4 lists the information displayed in the External Power Supply Status section of the **show inline power detail** output.

Table 6.4: External power supply status information

This Field...	Displays...
Fan	The status of the fan on the RPS-POE shelf, either ON or OFF.
Power Supply	Information about the power supplies in each of the six slots on the RPS-POE shelf.
Status	The status of each power supply installed in the RPS-POE shelf. This can be one of the following: GOOD – The power supply is present in the slot and working. NONE – A power supply is not present in the slot. FAILED – The power supply is present in the slot but not working.
Power	The power output of each power supply
Total	The total power output of all of the power supplies that have a status of GOOD.

Table 6.5 lists the information displayed in the System Inline Power Consumption section of the **show inline power detail** output.

Table 6.5: System Inline Power Consumption information

This Field...	Displays...
Module Number	Information about the module in each slot on the Foundry device.
POE Capable	Whether the module is capable of delivering Power over Ethernet.
Available Power	The amount of power available to the module.
Used Power	The amount of power used by the module.
Total	The total amount of power available to the modules.

Table 6.7 lists the information displayed in the System Inline Power Status section of the **show inline power detail** output.

Table 6.6: System Inline Power Status information

This Field...	Displays...
Slot	Information about the module in each slot on the Foundry device.
Inline Power Capability	Information about the POE capability of the modules in each slot. Possible values are: None – The module in the slot is not capable of delivering Power over Ethernet. Off – The module is capable of delivering Power over Ethernet, but no power is being supplied to the module. 802.3AF and Legacy – The module is capable of delivering power to 802.3af-compliant and 802.3af non-compliant devices.

The rest of the information displayed in the System Inline Power Status section is the same as that described in Table 6.3 on page 6-13.

Table 6.7 lists the information displayed in the System Total Port Detections section of the **show inline power detail** output.

Table 6.7: System Total Port Detections information

This Field...	Displays...
802.3AF	The number of detections where the powered device connected to the port was 802.3af-compliant.
Legacy	The number of detections where the powered device connected to the port was a legacy product (not 802.3af-compliant).
Open	The number of ports where Power over Ethernet is configured, but no power-consuming device is connected.
Short	The number of times where a short was detected during the detection of the power-consuming device, and power was disabled to the port.
Off	The number of ports where Inline power is not enabled.
Pending	The number of ports for which detection is pending.
Denied	The number of times a power-consuming device was detected on the port; however, there was insufficient power available to power the device, so power was disabled to the port.
Total	The total number of port detections of all types.

Chapter 7

Using Packet Over SONET Modules

This chapter describes the Foundry POS (Packet Over SONET) modules and how to configure and manage them.

Overview

SONET (Synchronous Optical Network) is based on a worldwide standard for fiber optic transmission, modified for North American asynchronous rates. In the rest of the world, this technology is known as Synchronous Digital Hierarchy (SDH). Foundry's Packet over SONET (POS) is fully compatible with SONET and Synchronous Digital Hierarchy (SDH) network facilities and is compliant with RFC 2615, "PPP over SONET/SDH" (and also RFC 1619, which 2615 obsoletes), and RFC 1662, "PPP in HDLC-like Framing".

POS (Packet over SONET) is the serial transmission of data over SONET frames through the use of Point-to-Point Protocol (PPP). The Foundry POS modules allow direct connection to interfaces within the SONET. POS is a transport technology that encapsulates packet data such as an IP datagram directly into SONET.

The POS modules are available on NetIron Internet Backbone routers and BigIron Layer 3 Switches with redundant management modules. You can use multiple POS modules in a chassis.

You can install the following types of POS modules in a BigIron chassis running Layer 3 Switch software:

- N2P2488-SR – Contains two OC-48c (2488 Mbps) POS/SDH ports for single-mode fiber, short reach (15 Km)
- N2P2488-IR – Contains two OC-48c POS/SDH ports for single-mode fiber, intermediate reach (40 Km)
- N2P2488-A-SR – Contains the same interfaces as the N2P2488-SR and has a larger FPGA than model N2P2488-SR, to support performance and feature enhancements
- N2P2488-A-IR – Contains the same interfaces as the N2P2488-IR and has a larger FPGA than model N2P2488-IR, to support performance and feature enhancements
- B2P2488-SR – Contains two OC-48c ports providing 2488 Mbps each.
- B2P622 – Contains two OC-12c ports providing 622 Mbps each.
- B2P155 – Contains two OC-3c ports providing 155 Mbps each.


NOTE: The N2P2488 and N2P2488-A models use the Foundry Network Processor Architecture (NPA) and are supported only on the NetIron Internet Backbone router.

NOTE: The N2P2488, N2P2488-A, and B2P models use different image files. See "Upgrading the Flash Code" on page 7-5.

NOTE: Software release 07.5.00 and later supports model N2P2488-A of the Network Processor Architecture (NPA) OC-48 POS module, as well as all models of the OC-3 and OC-12 modules. These releases do not support model N2P2488 of the OC-48 NPA POS module.

Figure 7.1 shows the front panel of a POS module. This is the B2P622. The front panels for all models look similar.

Figure 7.1 POS Module



The LEDs are described in “Status LEDs” on page 7-29.

Cable Specifications

Table 7.1 lists the cable specifications for Foundry POS modules.

Table 7.1: POS Interface Specifications

Transceiver	Power Budget	Launch Window	Transmit Power	Receive Power	Maximum Distance
OC-3c POS interfaces					
Single-mode short-reach	13 dB	1270 to 1380 nm	-28 to -8 dBm	-31 to -8 dBm	9.75 miles (15 Km)
Single-mode intermediate-reach	29 dB	1280 to 1335 nm	-5 to 0 dBm	-34 to -8 dBm	26 miles (40 Km)
Multimode	11.5 dB	1270 to 1380 nm	-18 to -14 dBm	-30 to -14 dBm	1.3 miles (2 Km)
OC-12c POS interfaces					
Single-mode long-reach	25 dBm	1280 to 1335 nm	-3 to 2 dBm	-28 to -8 dBm	65 miles (100 Km)
Single-mode Intermediate-reach	13 dBm	1274 to 1356 nm	-15 to -8 dBm	-28 to -7 dBm	9.32 miles (40 Km) ^a
Multimode short-reach	6 dBm	1270 to 1380 nm	-20 to -14 dBm	-26 to -14 dBm	1640ft. (500 m)
OC-48c POS interfaces					
Single-mode short-reach	8 dB	1260 to 1580 nm	-10 to -3 dBm	-18 to -3 dBm	9.75 miles (15 Km)
Single-mode intermediate-reach	13 dB	1260 to 1360 nm	-5 to 0 dBm	-18 to 0 dBm	26 miles (40 Km)

a.If the transceiver part number for the OC-12c POS module is HFCT-5208B (before May 2000), the maximum distance is 15 Km, not 40 Km.

NOTE: The OC-12c specification is the same, regardless of whether it is configured for 622 Mbps or 155 Mbps.

Network Processor Architecture POS Modules

The N2P2488 and N2P2488-A models use the Foundry Network Processor Architecture (NPA). NPA provides the following features:

- A dedicated packet processor (CPU) for each POS port.
- Ternary Content Addressable Memory (CAM), which the module can use for Layer 2, Layer 3, and Layer 4 fast lookups. The device can use the entries for wire-speed forwarding and wire-speed ACLs.
- Dedicated memory for certain applications including Adaptive Rate Limiting.

The ternary CAM and the application-specific memory allow the CPU to use more memory for other tasks.

Each module uses the same front panel features as other 2-port POS modules. Here is an example.



System Requirements

The 2-port OC-48c NPA modules are supported in the NetIron Internet Backbone router. The management module on the NetIron Internet Backbone router must be running the B2P or N2P IP-only image.

NOTE: Foundry recommends that you do not use NPA POS modules and non-NPA POS modules in the same circuit. For example, if you have an NPA module on one end of a POS link, do not use a non-NPA module on the other end of the link.

Installing a POS Module

To install a POS module, perform the following tasks:

- Configure the chassis slot to receive the module.
- Insert the module.

Configuring the Chassis to Receive the Module

When you plan to insert a module into a chassis slot, you first must configure the slot to receive the module unless the slot already contains the same type of module.

NOTE: If you are swapping out another module, you must disable the module before removing it from the Chassis device. See “Removing the Old Module” on page 2-40.

USING THE CLI

To prepare slot 1 to receive a 622 Mbps POS module, enter the following commands at the global CONFIG level:

```
BigIron(config)# module 1 bi-pos-2-port-622m-module
BigIron(config)# write memory
```

Syntax: module <slot-num> <module-type>

The <slot-num> parameter specifies the chassis slot:

- Slots on a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots on an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots on a 15-slot chassis are numbered 1 – 15, from left to right.

In the current software release, the <module-type> for a POS module can be one of the following:

- bi-pos-2-port-2488m-module
- bi-pos-2-port-622m-module
- bi-pos-2-port-155m-module

USING THE WEB MANAGEMENT INTERFACE

1. Enter the BigIron's IP address in your Web browser's Location or Address field, then press Enter.
2. Log on to the BigIron using a valid user name and password for read-write access.
3. Select the [Home](#) link to display the System configuration sheet (if not already displayed).
4. Select the [Module](#) link to display the Module panel.
5. Select the [Add Module](#) link.
6. Select the chassis slot that will receive the module from the Slot field's pulldown menu.
7. Select the module type from the Module Type field's pulldown menu. In the current software release, the following module types apply to POS modules:
 - bi-pos-2-port-2488m-module
 - bi-pos-2-port-622m-module
 - bi-pos-4-port-155m-module
 - bi-pos-2-port-155m-module
8. Click the Add button to send the configuration information to the chassis.
9. Select the [Save](#) link to save the configuration change.

Upgrading POS Software from a TFTP Server

The POS modules contain their own flash memory from which they can boot. To upgrade the boot code or flash code (system software) on a POS module, copy the upgrade onto a TFTP server to which the Layer 3 Switch has access, then download the code from the TFTP server to the POS modules in the chassis.

By default, the code on all the POS modules in the chassis is upgraded. If you want to upgrade code only a particular module, you can specify the module's slot number.

To upgrade the POS software, use the following CLI methods.

Upgrading the Boot Code

To upgrade the POS boot code from a TFTP server, enter a command such as the following:

```
BigIron# pos copy tftp flash 109.157.22.26 P2B06000.bin boot
```

This command upgrades the boot code on all POS modules in the chassis.

Syntax: pos copy tftp flash <ip-addr> <image-file-name> boot

Upgrading the Flash Code

The software required by each model of POS module differs. Make sure you install the proper software.

Table 7.2: POS Image Files

Model	Boot Image	Flash Image
B2P2488 B2P622 B2P155	• P2Bxxxxx.bin	• P2Rxxxxx.bin
N2P2488	• P2Bxxxxx.bin	• O2Rxxxxx.bin
N2P2488-A	• P2Bxxxxx.bin	• L3Pxxxxx.bin or • L2Pxxxxx.bin

NOTE: For the model N2P2488-A, the L3P image provides Layer 3 features including MPLS. The L2P image provides Layer 2 features such as Ethernet over POS. For the model N2P2488, the O2R image provides Layer 3 features only. Both ports on a module must run the same software, either L2P or L3P.

NOTE: Software release 07.5.00 supports model N2P2488-A, as well as all models of the non-NPA OC-3 and OC-12 modules. This release does not support model N2P2488 of the OC-48 NPA POS module.

Upgrading the Flash Code on a POS Module

NOTE: To upgrade flash code on POS NPA OC-48 model N2P2488-A, use the procedure in “Upgrading the Flash Code on POS NPA OC-48 Model N2P2488-A” .

To upgrade flash code on a POS module:

1. Place the new flash code on a TFTP server to which the Foundry device has access.
2. Enter the following command at the Privileged EXEC level of the CLI (example: `NetIron#`) to copy the flash code from the TFTP server into the flash memory of the each POS module:
 - **pos copy tftp flash** <tftp-server-ip-addr> <pos-image-file-name> **primary** | **secondary** [<slot>]

NOTE: If you specify a slot number (<slot> parameter), the software copies the new flash code only to the POS module in the specified slot.

3. Reload the software by entering one of the following commands:
 - **reload** (this command boots from the default boot source, which is the primary flash area by default)
 - **boot system flash primary** | **secondary**

Upgrading the Flash Code on POS NPA OC-48 Model N2P2488-A

1. Place the new flash code on a TFTP server to which the Foundry device has access.
2. Enter the following command at the Privileged EXEC level of the CLI (example: `NetIron#`) to copy the flash code from the TFTP server into the flash memory of the each POS module:
 - **pos copy tftp flash** <tftp-server-ip-addr> <pos-image-file-name> **primary** | **secondary** [<slotnum>]

NOTE: If you specify a slot number (<slotnum> parameter), the software copies the new flash code only to the module in the specified slot.

NOTE: If you do not specify a slot number (<slotnum> parameter), only the modules that are running the same image type (L2P or L3P) are upgraded. For example, if you specify an L3P image, and the chassis contains three N2P2488-A modules and two have L3P images, only the modules with the L3P images are upgraded. The L2P image on the other module is not affected.

3. Reload the software by entering one of the following commands:
 - **reload** (this command boots from the default boot source, which is the primary flash area by default)
 - **boot system flash primary | secondary**

Interactively Upgrading the POS Flash Code

If the flash code versions of the management module and a POS module do not match, the software disables the POS module when you reboot or reload the software. After disabling the POS module, the software prompts you to specify a boot source for the POS module. At this point, you can boot the module, then copy the flash code upgrade onto the module.

The software also generates a Syslog message to indicate that the POS module has been disabled due to a software mismatch.

NOTE: The software release version on the management module and POS modules must be exactly the same. For example, if you are upgrading from a Beta release, flash code versions 07.2.05 and 07.2.05B1 are not the same.

Here is an example of the messages that are displayed if you reload a device that has different flash code versions on the management module and a POS module:

```
BigIron#
!!! MGMT and POS(slot 1) modules are running incompatible SW
      MGMT SW version 07.2.05B1, POS SW version 07.2.05
!!! POS module will be brought down and then wait for boot instruction from MGMT
module

Taking down module 1 ...
Module 1 is now deleted
Detected module 1 being inserted
Bringing up module 1 ...
Please use "pos boot pri/sec/tftp" command to boot POS module with matching SW
```

In this example, the POS module in slot 1 has a flash code version that is different from the version on the management module. The software prompts you to interactively boot the POS module. Interactively booting the module allows you to specify a boot source, such as a TFTP server, that contains the correct version.

After the POS module boots, you can copy the correct flash code version onto the module's flash memory, using the procedure in "Upgrading the Flash Code on a POS Module" on page 7-5 or "Upgrading the Flash Code on POS NPA OC-48 Model N2P2488-A" on page 7-5.

To interactively boot a POS module, enter the following command at the Privileged EXEC level of the CLI (example: BigIron#):

- **pos boot tftp** <tftp-server-ip-addr> <pos-image-file-name>

After the module boots, copy the correct flash code version onto the module's flash memory.

Configuring POS Boot Parameters

The POS module has its own system software and boots after the management module boots. By default, the POS boots from the software image in its own primary flash. You can configure a POS module to boot from one of the following sources:

- POS module's primary flash
- POS module's secondary flash

To boot the POS module from a TFTP server, you must use the interactive boot mode, then enter the **pos boot tftp...** command after the module comes up.

Changing the Boot Source

To change the boot source for the POS module, use either of the following methods.

USING THE CLI

To change the boot source from the POS module's primary flash to its secondary flash, enter the following commands:

```
BigIron(config)# pos boot secondary
BigIron(config)# write memory
```

Syntax: pos boot interactive | primary | secondary

The **primary** and **secondary** parameters identify either the primary or secondary flash on the POS module.

The **interactive** parameter enables you to enter a separate command after the module comes up to boot the module from a TFTP server. If you use this method, you also need to use the **pos boot tftp...** command to boot the module after the module comes up. See "Reloading the Software from TFTP" on page 7-8.

Copying a POS Image File from a Flash Card to a POS Module's Flash Memory

To copy a POS image file from a flash card to a POS module's flash memory, use the following method.

USING THE CLI

To copy a POS image file from a flash card onto all the POS modules in the chassis, enter a command such as the following:

```
BigIron# pos copy slot1 flash P2R07000.bin primary
```

Syntax: pos copy slot1 | slot2 flash <pos-image-file-name> primary | secondary [slot]

The command in this example copies a POS image file named P2R07000.bin from the flash card in slot 1 to all the POS modules in the chassis.

To copy a POS image file from a flash card onto a specific POS module, enter a command such as the following:

```
BigIron# pos copy slot1 P2R07000.bin flash primary 4
```

The command in this example copies the specified image file onto the POS module in chassis slot 4 only, but does not copy the file to other POS modules in the chassis.

The following command copies a POS image file from a TFTP server to flash memory. This command also is present in earlier software releases.

Syntax: pos copy tftp flash <ip-addr> <pos-image-file-name> primary | secondary [<slot>]

Rebooting

You can reboot (reload the software) on an individual module or on all modules in the chassis, including the management module. You also can reload using a flash image on a TFTP server.

Reloading the Software on All Modules

To reload the software on all POS modules and all other modules in the chassis, including the management module, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# reload
```

Syntax: reload

Reloading the Software on an Individual Module

You can reload the software on an individual POS module, without also reloading the management module. To reload a POS module, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# reload pos 2
POS MODULE (2) App CPU in running mode:
      CPU 1 in state of POS_STATE_RUNNING
      CPU 2 in state of POS_STATE_RUNNING
Taking down module 2 ...
Module 2 is now deleted
BigIron Router#Detected module 2 being inserted
Bringing up module 2 ...
POS module at slot 2 is up and running
All POS Modules Up (1)
```

This command reloads the POS module in slot 2. Messages are displayed in the CLI to show the status of the reload. The management module is not also reloaded and thus continues to operate while the POS module is being reloaded.

Syntax: reload pos <slotnum>

Reloading the Software from TFTP

To boot the POS module from a TFTP server, you must use the interactive boot method, then use the following method to load the software after the module comes up.

USING THE CLI

To boot the POS module from a TFTP server, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# pos boot tftp 209.157.22.26 B2R06000
```

Syntax: pos boot tftp <tftp-server-ip-addr> <pos-image-file-name>

The <tftp-server-ip-addr> parameter specifies the IP address of the TFTP server.

The <pos-image-file-name> parameter lists the name of the image file you want the module to boot from the TFTP server.

Configuring POS Interfaces

To configure a Layer 3 POS interface, you need to add an IP address to the interface. Each POS interface also has the following additional parameters. The parameters have factory defaults but you can modify the values if needed for your network.

NOTE: To configure a Layer POS interface for remote bridging, see “Configuring POS for Layer 2 Switching” on page 7-16.

- Encapsulation type – You can configure a POS interface to use PPP (Point-to-Point Protocol), HDLC (High-Level Data Link Control), or Frame Relay encapsulation. The default is PPP.

NOTE: HDLC is not supported on Layer 2 POS.

- Clock source – You can configure a POS interface to use the POS module's internal clock or use the network as the clock source. By default, Foundry POS interfaces use the internal clock as the clock source.
- Loopback path – During startup, the POS module tests each POS interface by performing loopback tests. You can configure the path used by the loopback test. The path can consist of the POS module's interface circuitry alone, or can include both the local and remote POS interfaces. The loopback path and tests are disabled by default.
- Bandwidth – You can change a 622 Mbps port to run at 155 Mbps if needed.
- MTU (Maximum Transmission Unit) – You can specify the maximum IP packet size or SONET/SDH frame size. The size can be from 60 – 4470 bytes. The default is 4470 bytes.
- Frame type – You can configure a Foundry POS interface to transmit and receive SDH (Synchronous Digital Hierarchy) frames or SONET (Synchronous Optical Network) frames. The default is SONET.
- Keepalive messages – You can disable or reenable a POS interface to send PPP or HDLC keepalive messages to the POS interface at the other end of the link. Keepalive messages are enabled by default.
- ATM scramble mode – You can enable or disable scrambling of the Synchronous Payload Envelope (SPE). When you enable scrambling, the data in the SONET packet is scrambled for security. ATM SPE scrambling is disabled by default.
- CRC (Cyclic Redundancy Check) – You can specify the length of the CRC field in each packet sent by the POS interface. The default length is 32 bits. You can change the length to 16 bits.
- SONET overhead Flags – You can change the following values in the SONET frame header:
 - c2 – Path signal identifier, which identifies the payload content type. The default is 0xCF, which specifies PPP or HDLC. This is part of the path overhead.
 - j1 – Bits 5 and 6 of the payload pointer byte, which indicates the frame type. The default is 0x00 (SONET). You can change this value to 0x02 (SDH). This is part of the path overhead.
 - j0 – Section trace byte, which can be configured to allow interoperability with various other types of SDH devices. The default value is 0xCC. This is part of the section overhead.

To configure POS interface parameters, use the procedures in the following sections.

Adding an IP Address

You can add up to 24 IP subnet interfaces to each POS interface.

USING THE CLI

To add an IP address to POS interface 2/1, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# ip address 209.157.22.26/24
BigIron(config-posif-2/1)# write memory
```

Syntax: [no] ip address <ip-addr> <ip-mask> [secondary]

or

Syntax: [no] ip address <ip-addr>/<mask-bits> [secondary]

Use the **secondary** parameter if you have already configured an IP address within the same subnet on the interface.

Changing the Interface State

The POS interfaces are enabled by default. To disable or reenable an interface, use the following method.

USING THE CLI

To disable POS interface 2/1, enter the following commands:

```
BigIron(config)# interface pos 2/1
```

```
BigIron(config-posif-2/1)# disable
BigIron(config-posif-2/1)# write memory
```

Syntax: disable

To reenable POS interface 2/1, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# enable
BigIron(config-posif-2/1)# write memory
```

Syntax: enable

Changing the Encapsulation Type

Foundry POS interfaces use the PPP encapsulation type by default. You can change the encapsulation type of an interface to HDLC or Frame Relay, or back to PPP using the following method.

NOTE: Both ends of the POS link must use the same encapsulation type.

NOTE: HDLC encapsulation is not supported on Layer 2 POS. You must use PPP encapsulation.

USING THE CLI

To configure POS interface 2/1 to use HDLC, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# encapsulation hdlc
BigIron(config-posif-2/1)# write memory
```

Syntax: [no] encapsulation hdlc | ppp | frame-relay [ietf]

NOTE: If you are configuring a Frame Relay interface, see “Configuring POS for Frame Relay” on page 7-15.

Changing the Clock Source

By default, Foundry POS interfaces use the internal clock, which means that clocking information comes from the POS module itself. You can change the clock source for an interface to network.

If you are connecting two Foundry POS modules back-to-back or the POS interfaces are connected by a fiber link that has no clocking information on it, use the internal clock source. Otherwise, use the network (line) as the source.

To change the clock source, use the following method.

USING THE CLI

To change the clock source for POS interface 2/1 to internal (the POS module itself), enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# clock internal
BigIron(config-posif-2/1)# write memory
```

Syntax: clock internal | line

The **internal** and **line** parameters specify whether the clock source is on the POS module (internal) or on the network (line).

Changing the Loopback Path

Foundry POS interfaces can use the following loopback configurations for self tests:

- Internal – Packets that the router transmits on the interface are looped back to the interface’s POS framer.

The internal loopback configuration is useful for checking the POS circuitry.

- **Line** – The interface's transmit and receive fibers are logically linked so that packets received on the receive fiber are sent back out on the transmit fiber. Use this mode on the POS interfaces on both ends of a link to test the interfaces along with the link.

By default, loopback is disabled. Do not enable loopback unless you are testing the interface.

USING THE CLI

To configure POS interface 2/1 for internal loopback, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# loop internal
BigIron(config-posif-2/1)# write memory
```

Syntax: loop internal | line

The **internal** and **line** parameters specify the path for the loopback. The **internal** parameter loops packets transmitted on the interface back to the framer. The **line** parameter loops packets that are received on the receive fiber of the port back out on the transmit fiber.

Changing the MTU

The MTU (Maximum Transmission Unit) specifies the maximum number of bytes a frame transmitted on the interface can contain. You can configure the MTU to a value from 60 – 4470 bytes. The default is 4470 bytes.

USING THE CLI

To change the MTU for POS interface 2/1 to 1200 bytes, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# mtu 1200
BigIron(config-posif-2/1)# write memory
```

Syntax: mtu <length>

The <length> can be from 60 – 4470 bytes.

Changing the CRC Length

The CRC (Cyclic Redundancy Check) length specifies whether the CRC portion of each frame transmitted on the interface is 16 bits or 32 bits long. The default is 32 bits. You can change the CRC length using the following method.

USING THE CLI

To change the CRC length for POS interface 2/1 to 16 bits, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# crc 16
BigIron(config-posif-2/1)# write memory
```

Syntax: crc 16 | 32

The **16** and **32** parameters specify how many bits in each frame transmitted on the interface contain the CRC data.

Disabling or Reenabling Keepalive Messages

You can disable or reenabling a POS interface to send keepalive messages to the POS interface at the other end of the link. Keepalive messages are enabled by default.

The message interval is 10 seconds and is not configurable. Every ten seconds, the POS interface sends a keepalive message addressed to the POS interface at the other end of the connection. The timeout is 30 seconds. If the other interface (remote end of the connection) does not respond as expected to a keepalive message after 30 seconds, the Layer 3 Switch takes the local interface down. The link remains up but the POS

interface is down. If the Layer 3 Switch receives a keepalive response from the other interface, the Layer 3 Switch brings the local interface back up.

For PPP connections, one or both interfaces can send the keepalive messages. For HDLC connections, both interfaces must send keepalive messages. If the keepalive messages are enabled on only one interface in an HDLC connection, the interface will conclude based on the absence of keepalive messages from the other interface that the connection is down and will take the local interface down.

To disable or re-enable the POS keepalive messages, use the following method:

USING THE CLI

To disable keepalive messages on POS interface 2/1, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# no keepalive
BigIron(config-posif-2/1)# write memory
```

Syntax: [no] keepalive

To reenable the messages, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# keepalive
BigIron(config-posif-2/1)# write memory
```

Changing the Bandwidth

Depending on the POS module you have installed, the interfaces operate at 155 Mbps or 622 Mbps by default.

- The ports on model number B2P622 run at 622 Mbps by default. If needed, you can reduce the port bandwidth to 155 Mbps on an individual port basis.
- The port bandwidth on other models is not configurable.

To change the bandwidth of a POS port, use the following method.

USING THE CLI

To change the bandwidth of POS interface 2/1 from 622 Mbps to 155 Mbps, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# bandwidth 155
BigIron(config-posif-2/1)# write memory
```

Syntax: bandwidth 155 | 622

The **155** and **622** parameters specify how many megabits per second the port can transmit.

To change the port bandwidth of POS interface 2/1 back to 622 Mbps, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# bandwidth 622
BigIron(config-posif-2/1)# write memory
```

Changing the POS Flags

The POS flags set POS “overhead” values, which are values in the packet headers that identify the packet payload content type, the packet’s compatibility with some other SDH devices, and the frame type.

- c2 – Identifies the payload type. This flag can have one of the following values:
 - cf (0xCF) – PPP or HDLC. This is the default.
 - 13 (0x13) – ATM
- j0 – This is the section trace byte, formerly the C1 byte. You can set this flag to 1 (0x01) for interoperability with certain SDH devices. The default value is cc (0xCC).
- h1 – Identifies the frame type. This flag is part of the payload pointer byte (bits 5 and 6 of the H1 number 1

payload pointer byte). The flag can have one of the following values:

- 0 (0x00) – SONET. This is the default.
- 2 (0x02) – SDH

To change a flag, use the following method.

USING THE CLI

To change the j0 flag from the default value to 1, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# pos flag j0 1
BigIron(config-posif-2/1)# write memory
```

Syntax: pos flag c2 | h1 | j0 <value(Hex)>

The **c2**, **h1**, and **j0** parameters specify the flag you are setting.

The <value(Hex)> parameter specifies the value you are assigning to the flag. The flag values are hexadecimal numbers.

Changing the Frame Type

Foundry POS interfaces support the following frame types:

- SDH (Synchronous Digital Hierarchy) – An international standard for optical digital transmission at rates from 155 Mbps (used for STM-1) to 2.5 Gbps (used for STM-16) and higher.
- SONET (Synchronous Optical Network) – An American National Standards Institute (ANSI) standard (T1.1051988) for optical digital transmission at rates from 51 Mbps (used for OC-1) to 2.5 Gbps (used for OC-48c) and higher.

To change the frame type, use the following method.

USING THE CLI

To change the frame type of POS interface 2/1 to SDH, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# pos framing sdh
BigIron(config-posif-2/1)# write memory
```

Syntax: pos framing sdh | sonet

The **sdh** and **sonet** parameters specify the framing type.

Enabling or Disabling ATM Scrambling

You can enable a POS interface to scramble the data in the Synchronous Payload Envelope (SPE), which is the data portion of ATM packets. Scrambling the data provides additional security.

NOTE: Both ends of the link must use the same scrambling algorithm.

To enable scrambling, use the following method.

USING THE CLI

To enable scrambling on POS interface 2/1, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# pos scramble-atm
BigIron(config-posif-2/1)# write memory
```

Syntax: pos scramble-atm

Configuring a POS Interface Using the Web Management Interface

To configure a POS port or view its current settings, select Configure->Port->POS from the tree view of options. The following panel is displayed.

POS Port Configuration													
Port	Speed	Encapsulation	MTU	Clock	Loop Back	Scramble-ATM	Framing	CRC	Keep Alive	C2	J0	H1	
12/1	622000	PPP	4470	Internal	None	Disable	SONET	32	10	cf	cc	00	Modify
12/2	622000	PPP	4470	Internal	None	Disable	SONET	32	10	cf	cc	00	Modify
Port	Speed	Encapsulation	MTU	Clock	Loop Back	Scramble-ATM	Framing	CRC	Keep Alive	C2	J0	H1	

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

To modify settings for a port, click Modify next to the row of information for the port. The following panel is displayed.

POS Port	
Slot: 12 Port: 2	
Name:	<input type="text"/>
Status:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Speed:	<input type="radio"/> 155000 <input checked="" type="radio"/> 622000 <input type="radio"/> 2488000
Encapsulation:	<input type="radio"/> HDLC <input checked="" type="radio"/> PPP
MTU:	<input type="radio"/> 1500 <input checked="" type="radio"/> 4470
Clock:	<input checked="" type="radio"/> Internal <input type="radio"/> Line
Loop Back:	<input type="radio"/> Line <input type="radio"/> Internal <input checked="" type="radio"/> None
CRC:	<input checked="" type="radio"/> 32 <input type="radio"/> 16
Framing:	<input checked="" type="radio"/> SONET <input type="radio"/> SDH
Scramble ATM:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Keep Alive:	<input type="text" value="10"/>
C2:	<input type="text" value="cf"/>
J0:	<input type="text" value="cc"/>
H1:	<input type="text" value="00"/>

[\[Show\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

For information about the parameters, see the equivalent CLI sections above.

Configuring POS for Frame Relay

You can configure a Foundry POS interface for Frame Relay. To configure the interface:

- Set the encapsulation type to Frame Relay and specify the Frame Relay type. You can specify IETF (RFC 1490) or use the default, Cisco-compatible.
- Specify the Frame Relay interface type.

NOTE: The current software release supports Data Terminal Equipment (DTE) only. The other end of the link must be configured as a DCE link.

- Specify the Data-Link Connection Identifier (DLCI). This is the circuit ID for the link and can be a number in the range from 1 – 1023. The circuit ID must be the same on both ends of the link.
- Set the Local Management Interface (LMI) type. You can specify ANSI, CCITT, or LMI (Cisco-compatible). The default is LMI.

In addition, the setting for POS keepalive messages must be the same on both ends of the link. The keepalive messages are enabled by default for Foundry POS interfaces. See “Disabling or Reenabling Keepalive Messages” on page 7-11.

NOTE: The current software release supports Data Terminal Equipment (DTE) only. The other end of the link must be configured as a DCE link. Also, the current release supports only point-to-point links, not point-to-multipoint. Both ends of the link must be configured for point-to-point.

Changing the Encapsulation Type

To configure POS interface 2/1 to use Frame Relay, enter the following commands:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# encapsulation frame-relay
```

Syntax: [no] encapsulation hdlc | ppp | frame-relay [ietf]

If you enter the command without the **ietf** parameter, the software uses the default Frame Relay type, Cisco-compatible.

Specifying the Frame Relay Interface Type

NOTE: DTE is the default. The current release supports DTE only, so you do not need to enter this command. The other end of the link must be DCE.

Syntax: [no] frame-relay intf-type dte

Specifying the DLCI

The DLCI identifies the circuit number for the POS link. You must use the same DLCI on both ends of the link. To specify the DLCI, enter a command such as the following:

```
BigIron(config-posif-2/1)# frame-relay interface-dlci 69
```

Syntax: [no] frame-relay interface-dlci <num>

The <num> parameter specifies the DLCI and can be a number from 1 – 1023.

Specifying the LMI Type

Foundry POS Frame Relay links use the Cisco-compatible LMI type by default. To specify a different LMI type, enter a command such as the following:

```
BigIron(config-posif-2/1)# frame-relay lmi-type ansi
```

Syntax: [no] frame-relay lmi-type ansi | ccitt | lmi

The default is **lmi**.

Verifying the Configuration

To verify the POS Frame Relay configuration, enter commands such as the following:

```
BigIron(config-posif-2/1)# show interface pos 2/1

POS2/1 is up, line protocol is up
  No port name
  Hardware is Packet over Sonet
  Internet address is 101.101.101.5/30
  MTU 4470 bytes, encapsulation FR(cisco), clock is line
  Framing is SONET, BW 155000Kbit, CRC 32
  Loopback not set, keepalive is set (10 sec), scramble disabled
  5 minute input rate: 112448 bits/sec, 135 packets/sec
  5 minute output rate: 2394880 bits/sec, 204 packets/sec
  7292298 packets input, 386668128 bytes, 0 no buffer
  Received 0 CRCs, 0 shorts, 0 giants, 0 alignments
  10029741 packets output, 14905847936 bytes, 0 underruns
```

The command in this example indicates that POS interface 2/1 is using the Cisco-compatible Frame Relay encapsulation and the keepalive is enabled and set to 10 seconds.

```
BigIron(config-posif-2/1)# show interface brief slot 2
Port Link State  Encap Clock Loop Speed  mtu  frame  scram  crc c2 j0 h1
2/1 up           FR    int  none  622  4470  sonet  no   32  cf  cc  00
2/2 up           FR    int  none  622  4470  sonet  no   32  cf  cc  00
```

This command shows brief interface information for the POS interfaces on the module in chassis slot 2.

Configuring POS for Layer 2 Switching

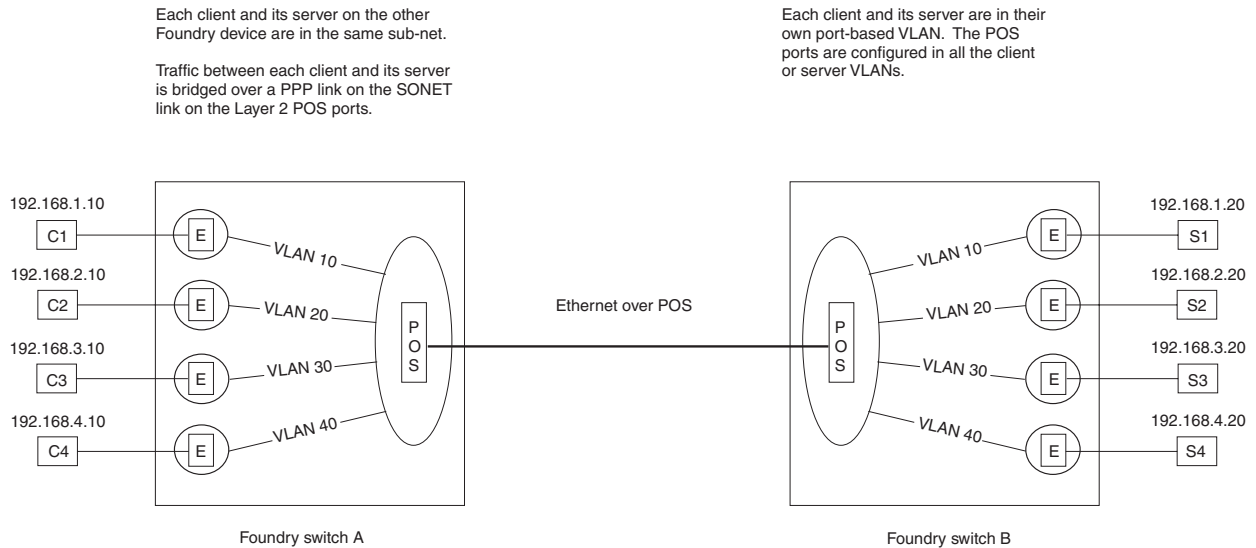
You can configure a POS port for Layer 2 switching. Layer 2 switching enables the devices on each end of a POS link to exchange Ethernet frames over POS. Layer 2 switching over POS is especially useful for Virtual Private Networking (VPN) applications in which IP hosts share a common subnet even though the hosts are geographically distant from one another.

NOTE: To use Layer 2 POS on an NPA module model N2P2488-A, install the L2P flash image. To use Layer 2 POS on an NPA module model N2P2488, install the O2R flash image. To use Layer 2 POS on a non-NPA module, install the P2R flash image.

NOTE: HDLC encapsulation is not supported on Layer 2 POS. You must use the default encapsulation, PPP.

Figure 7.2 shows an example of a Layer 2 switching configuration using Layer 2 POS.

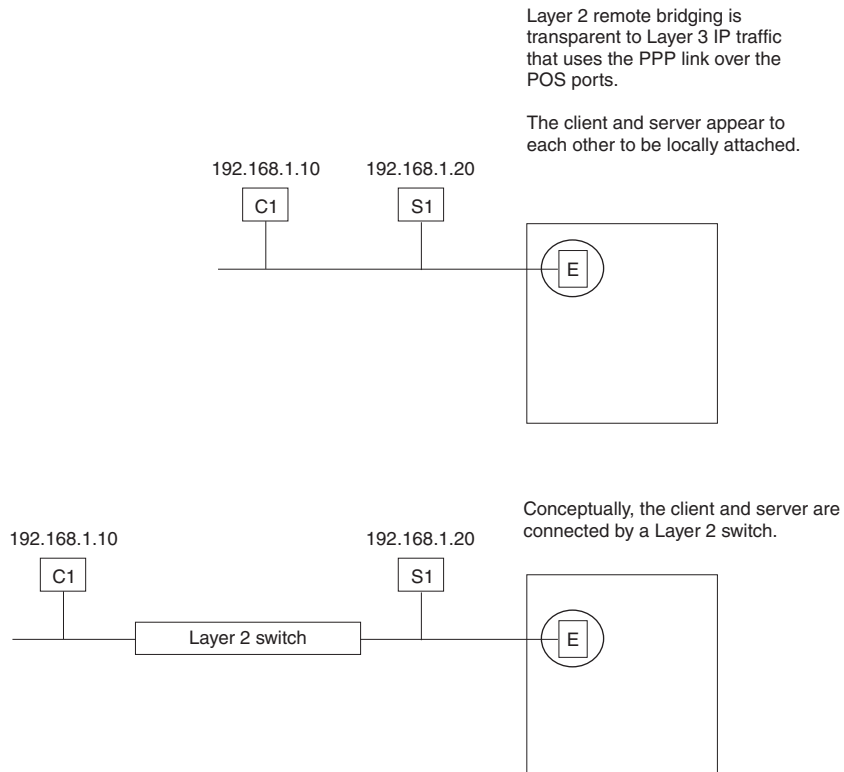
Figure 7.2 Basic POS Layer 2 switching configuration



This configuration shows four IP subnets. Each subnet has members on both sides of the POS link. For example, client C1 in the 192.168.1.x subnet can communicate at Layer 2 with server S1 in the same subnet, even though the two devices are on different sides of the POS link.

Figure 7.3 shows the network in Figure 7.2 from the perspective of client C1 and server S1. The devices appear to one another to be on the same Ethernet LAN segment. Conceptually, the devices are connected by a Layer 2 switch or bridge. The clients are not aware that the Layer 2 connection between them is a POS WAN link between two geographically distant devices.

Figure 7.3 IP Host perspective of the switching link



Notice that the devices in Figure 7.2 are configured with a separate port-based VLAN for each subnet. The separate port-based VLANs provide a private Layer 2 broadcast domain for each subnet. The separate port-based VLANs ensure that devices on one subnet do not receive broadcast traffic from devices on the other subnets. Partitioning the ports into separate port-based VLANs is especially useful for co-location implementations, when multiple clients lease ports on the same devices. The port-based VLANs ensure that one client's ports do not receive broadcast traffic from another client's ports.

In fact, to enable Layer 2 switching on a POS port, you must add the port as a tagged port to each of the port-based VLANs that contains the subnets you want to bridge. The ports are tagged so that they can properly multiplex traffic from the different VLANs for sending the traffic over the PPP link, and demultiplex the traffic at the other end of the link. For example, when client C2 sends traffic to server S2, the traffic is tagged with the VLAN ID (20). The POS ports on both devices also are members of VLAN 20 and can therefore forward traffic between C2 and S2 at Layer 2.

In addition, since the traffic is tagged, broadcast traffic on the 192.168.2.x subnet (in VLAN 20) goes only to other devices on the same subnet, because VLAN 20 is its own Layer 2 broadcast domain. Only the ports in VLAN 20 receive broadcasts from a host connected to a port in VLAN 20. For example, the POS ports forward a broadcast packet from S2 to client C2 but not to clients C1, C3, or C4.

Link Redundancy and Load Balancing

The configuration in Figure 7.2 on page 7-17 uses a single POS link between the two remote devices. However, you can provide link redundancy using either of the following methods:


- Enable the Spanning Tree Protocol (STP) on all the port-based VLANs of which the POS ports are members, then set the path costs to different values for different VLANs, so that STP prefers one POS port for some of the VLANs but not others.
- Configure the POS ports as a trunk group.

NOTE: Trunking is supported on POS OC-3 and OC-12 ports but not on OC-48 ports. Server trunking of POS ports is supported only for Layer 2 and requires software release 07.6.01 or later.

Using STP for Redundancy and Load Balancing

Figure 7.4 shows an example of link redundancy using STP. In this example, STP is enabled in port-based VLANs 10, 20, 30, and 40. The POS ports are members of each of these VLANs. Each VLAN runs a separate spanning tree (a separate instance of STP).

Figure 7.4 POS Layer 2 port redundancy using STP – default port path cost used

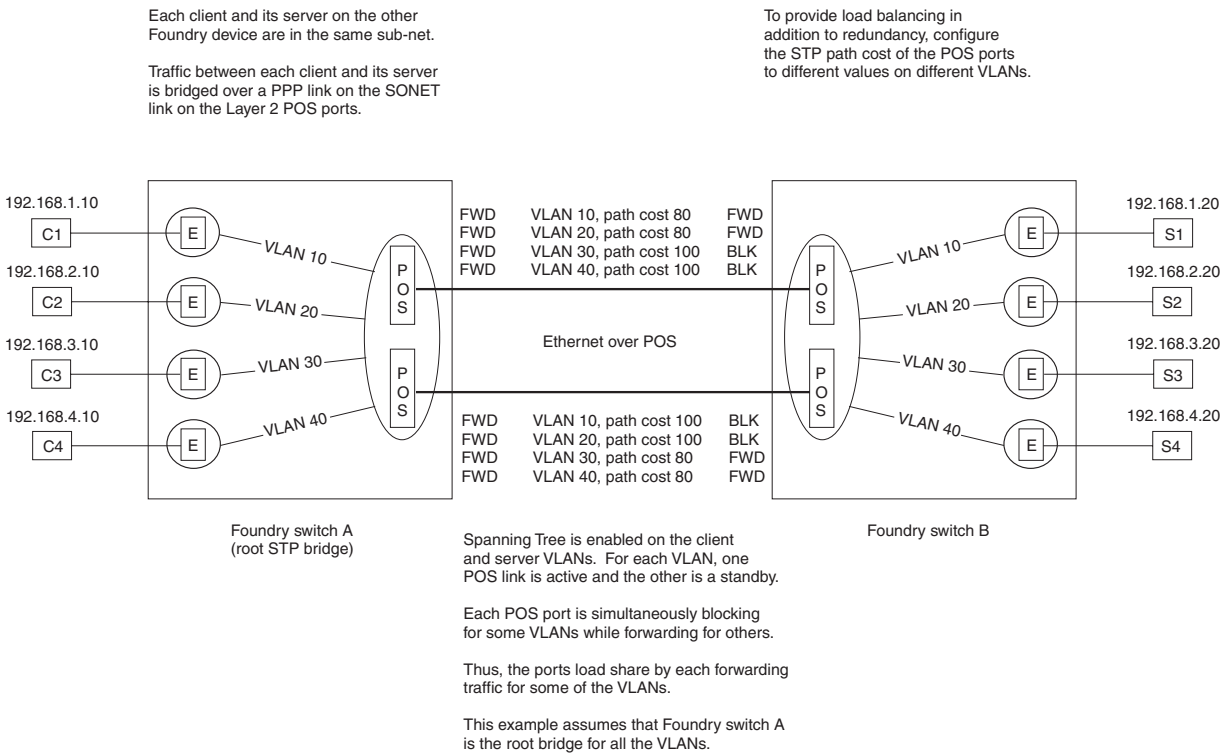


If you use the default settings for the STP parameters as shown in Figure 7.4, enabling STP on the POS ports for all the port-based VLANs provides redundancy but does not provide load sharing. In this example, STP blocks one of the POS ports for all the port-based VLANs.

Assuming the POS ports are operating at the same speed, the default STP path cost for each of the ports is the same. Therefore, if you use the default STP path cost for each port, the spanning tree in each of the port-based VLANs selects the same POS port for forwarding and blocks the other POS port.

To also provide load balancing, change the path cost of each the POS ports in some of the VLANs, but no others. For example, to configure the POS ports on a switch so that STP uses one port for forwarding VLANs 10 and 20 and the other port for forwarding VLANs 30 and 40, change the port path cost on the first POS port to a lower value for VLANs 10 and 20. On the other POS port, change the port path cost to the same lower value for the other VLANs, 30 and 40. Figure 7.5 shows this configuration.

Figure 7.5 POS Layer 2 port redundancy using STP – port path costs configured for load balancing

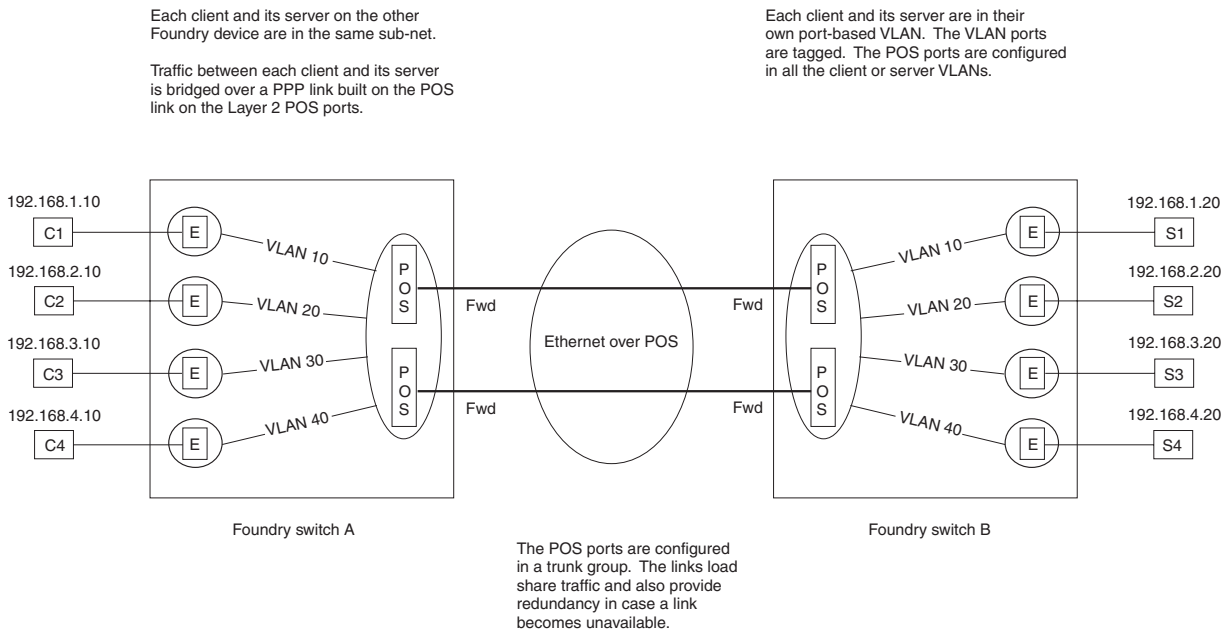


Notice that each POS port on switch B blocks two of the VLANs but forwards the other two VLANs. STP prefers the path that has the lower value as the path to the root bridge. In this example, STP in VLANs 10 and 20 on switch B prefers the first POS port as the path to the root bridge. STP in VLANs 30 and 40 prefer the second POS port as the path to the root bridge.

Using a Trunk Group for Redundancy and Load Balancing

Figure 7.6 shows an example of POS ports configured for redundancy by adding them to a trunk group. In addition to redundancy, a trunk group provides load sharing. The ports can each actively send traffic, thus providing additional bandwidth. When you use STP to provide redundancy, all ports in a trunk group have the same STP status. The status is based on the trunk group’s primary (lead) port.

NOTE: You can use a trunk group for Layer 2 POS redundancy and load balancing only on a BigIron or NetIron running Layer 2 Switch code, not Layer 3 Switch code.

Figure 7.6 POS Layer 2 port redundancy using a trunk group

This example shows a trunk group configured on a POS module in each chassis. In each module, both POS ports are configured together as a single link. The links load share and also provide redundancy. If a link in a trunk group becomes unavailable, the connection is maintained by the other link.

You can use STP with trunk group links. STP regards a trunk group as a single link and thus either forwards or blocks traffic on all the ports within the trunk group.

Configuration Procedures

To configure a Foundry device for POS Layer 2 switching:

- Change POS interface parameters, if you need to change a parameter from its default value.

NOTE: HDLC encapsulation is not supported on Layer 2 POS. You must use the default encapsulation, PPP.

- Configure each host subnet in a separate port-based VLAN.
- Add the POS port as a tagged port to all the host port-based VLANs.
- Optionally, configure redundancy by configuring a second POS port as above, then doing one of the following:
 - Configuring Spanning Tree Protocol (STP) parameters on the ports.
 - Adding the POS ports to a trunk group. If you add the ports to a trunk group, the ports load balance traffic in addition to providing link redundancy.

Configuring a POS Port for Layer 2 Switching

POS ports are configured for Layer 3 IP routing by default. To configure a POS port for Layer 2 switching, you must add the port as a tagged port to a port-based VLAN.

NOTE: Layer 2 POS ports must be tagged. You cannot add a POS port to a port-based VLAN without tagging the port.

A POS port is by default a routing-only port. That is, it is by default not a member of the default Layer 2 broadcast domain (VLAN 1). You can convert a POS port to a Layer-2 switching-only port by adding it to one or more VLANs as a tagged port. Only traffic from ports belonging to the same VLANs as the tagged POS port are forwarded.

NOTE: By default, POS ports are not members of the device's default VLAN (VLAN 1) or of any other VLAN.

NOTE: A POS port cannot be added as an untagged port to any VLAN.

A single POS link can multiplex and demultiplex traffic from different clients while keeping each client's traffic within the client's own Layer 2 broadcast domain (port-based VLAN). The POS port on the send side of the link sends packets with different tags to the POS port on the receive side. The port on the receive side forwards the packets based on the VLAN tag. Thus, on the receive side each packet is forwarded only to the Ethernet ports in the same VLAN as the client.

To configure a POS port for Layer 2 switching, use the following CLI method.

USING THE CLI

To configure a POS port for Layer 2 switching, enter a command such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# tagged pos 2/1
```

These commands change the CLI to the configuration level for port-based VLAN 10, then add POS port 2/1 to the VLAN as a tagged port. You must add the port as a tagged port. You cannot add the port as an untagged port.

Syntax: vlan <vlan-id> [by port]

Syntax: tagged pos <portnum> [to <portnum> | pos <portnum>]

NOTE: This example assumes that port-based VLAN 10 has already been configured. If the VLAN is not configured, use the command **vlan 10 by port**.

To verify the VLAN configuration for all the VLANs, enter the following command at any level of the CLI:

```
BigIron(config)# show vlans
```

To view all the VLANs configured for a specific POS port, enter a command such as the following at any level of the CLI:

```
BigIron(config)# show vlans pos 2/1
```

Configuring STP Parameters

The following STP features are supported on Layer 2 POS ports:

- Standard STP parameters – All the standard bridge and port parameters supported on Ethernet ports also are supported on POS ports.
- Single-instance STP – All ports that are members of VLANs that have STP enabled are members of a single BPDU broadcast domain, and thus share a common STP root bridge. However, the VLANs continue to be separate broadcast domains for other types of Layer 2 traffic.
- Fast Uplink Span – The Fast Uplink feature enhances STP performance for wiring closet switches with redundant uplinks. Convergence following a transition from an active link to a redundant link takes around four seconds, instead of the STP's default convergence time of 30 seconds. Convergence consists of both the listening and learning states.

NOTE: Fast Port Span applies only to ports that are connected to end hosts, not to Layer 2 POS configurations, and therefore is not supported.

Changing a POS Port's Path Cost

When selecting among multiple links to the root bridge, STP chooses the link with the lowest path cost and blocks the other paths. Each port type has a default STP path cost, listed in Table 7.3.

Table 7.3: Default STP Port Path Costs

Port Type	Default Path Cost
10 Mbps	100
100 Mbps	19
Gigabit	4
OC-3c	200
OC-12c	80
OC-48c	20

You can set a port's STP path cost to a value from 0 – 65535. If you want to bias STP's selection to favor one POS port over another of the same speed, use the following CLI method.

USING THE CLI

To change the STP path cost on a POS port, enter a command such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning-tree pos 2/1 path-cost 100
```

These commands change the CLI to the VLAN configuration level for VLAN 10, then change the POS port's path cost for that VLAN's spanning tree to 100.

Syntax: [no] spanning-tree ethernet | pos <portnum> path-cost <value> | priority <value>

The **ethernet | pos <portnum>** parameter specifies the interface.

The <num> specifies the path cost and can be from 0 – 65535.

Changing a POS Port's STP Priority

The STP priority of a port determines which port within a spanning tree STP prefers for forwarding traffic into and out of the spanning tree. For a port-based VLAN, the port priorities determine the port that STP selects within that VLAN to forward traffic out of the VLAN. STP prefers the port with the highest port priority. You can set a port's STP priority to a value from 0 – 255. The default for all port types is 128. The default depends on the port type. See Table 7.3.

To change a POS port's STP priority, use the following CLI method.

USING THE CLI

To change the STP priority for a POS port, enter a command such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning pos 2/1 priority 200
```

These commands change the CLI to the configuration level for the POS port, then change the POS port's priority for that VLAN's spanning tree to 200.

Syntax: [no] spanning-tree ethernet | pos <portnum> priority <num>

The **ethernet | pos <portnum>** parameter specifies the interface.

The <num> specifies the priority and can be from 0 – 255. The default for all port types is 128.

Enabling or Disabling STP on the Port

Use the following CLI method to enable or disable STP on a POS port.

USING THE CLI

To enable STP on a POS port, enter commands such as the following:

```
BigIron(config)# interface pos 2/1
BigIron(config-if-2/1)# spanning-tree
```

These commands change the CLI to the configuration level for the POS port, then enable STP on the VLAN.

Syntax: [no] spanning-tree

Enabling or Disabling Fast Uplink Span on the Port

To enable Fast Uplink on a pair of POS ports, use the following CLI method.

USING THE CLI

To configure a group of POS ports for Fast Uplink Span, enter a command such as the following:

```
BigIron(config)# fast uplink-span pos 2/1 to 2/2
```

This command configures POS ports 2/1 and 2/2 as a Fast Uplink Span group.

Syntax: [no] fast uplink-span [pos <portnum> [pos <portnum>... | to <portnum>]]

For more information about this feature, see “Configuring Spanning Tree Protocol (STP) and IronSpan Features” on page 10-1.

Displaying Spanning Tree Information

To display STP information for a VLAN, enter the **show span** command. See “Displaying STP Information” on page 10-8.

Configuring the POS Ports into a Trunk Group

Use the following CLI method to configure POS ports into a trunk group.

NOTE: You can use a trunk group for Layer 2 POS redundancy and load balancing only on a BigIron or NetIron running Layer 2 Switch code, not Layer 3 Switch code.

NOTE: Trunking is supported on POS OC-3 and OC-12 ports but not on OC-48 ports. Server trunking of POS ports is supported only for Layer 2 and requires software release 07.6.01 or later.

USING THE CLI

To configure two POS ports on a single POS module into a trunk group, enter a command such as the following:

```
BigIron(config)# trunk pos 1/1 to 1/2
```

This command configures the ports on a POS module in slot 1 into a trunk group. Port 1/1 is the primary port. To make configuration changes to the trunk group, make the changes on the primary port. The software automatically applies the changes to the other port(s).

To configure a multi-slot trunk group, enter commands such as the following:

```
BigIron(config)# trunk pos 1/1 to 1/2 pos 3/1 to 3/2
BigIron(config)# write memory
BigIron(config)# trunk deploy
```

The first command configures the POS ports on the modules in slots 1 and 3 into a single, multi-module trunk group. The lowest numbered port is always the primary port. The other commands save the configuration change to the startup-config file, then activate the trunk configuration change.

Syntax: [no] trunk [server | switch] pos <portnum> to <portnum> [pos <portnum> to <portnum>]

The **server | switch** parameter indicates the trunk group type. Since the POS ports in the trunk group will be connected to POS ports on another chassis, the type is always **switch**. This is the default, so you do not need to specify it.

The **pos <portnum> to <portnum>** parameter specifies the port range. Always specify the primary port first. To configure a multi-module trunk group, repeat the parameter for the port range on the second module.

The configuration rules for POS ports are the same as the rules for Ethernet ports.

- Each group consists of a primary port and consecutively numbered secondary ports. Always specify the lowest numbered port first (the primary port) followed by the other ports in ascending numerical order. On a two-port POS module, the first port on the module is the primary port and the second port is the secondary port. When you configure interface parameters for a trunk group, configure them on the primary port. The software then applies the changes to all the secondary ports automatically.
- You can configure ports on two POS modules in a chassis as a single, multi-module trunk group. In a multi-slot trunk group, the port with the lowest port number (the first port in the lower numbered chassis slot) is always the primary port. The trunk group must contain the same number of ports from each module. In addition, the modules must be the same type (for example, two 622 Mbps modules).
- You cannot mix POS ports and Ethernet ports in the same trunk group.
- All the ports in the trunk group must be connected to a single trunk group on the other chassis.

To verify the trunk configuration, enter the **show trunk** command.

Configuring Automatic Protection Switching (APS) for Layer 2 POS

Starting in release 07.6.03, Automatic Protection Switching (APS) is supported on Layer 2 POS links. Previous releases supported POS APS with Layer 3 features only. Foundry devices support POS APS in a 1 + 1 architecture, where a working interface is paired with a protect (backup) interface. If the working interface fails, the protect interface takes over the traffic load. See “Configuring Automatic Protection Switching (APS)” on page 7-35 for more information on APS.

APS operates transparently in a Layer 2 POS configuration. When a link switchover occurs, Layer 2 features, including known/unknown unicast, spanning tree, and MAC learning and aging, are not affected or disrupted.

All of the commands for Layer 2 POS and POS APS are still applicable in a Layer 2 POS APS configuration. To configure APS for Layer 2 POS, you specify working and protect interfaces, then add the working interface to a VLAN.

For example, to configure a working APS interface:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# spanning-tree
BigIron(config-posif-2/1)# aps working 1
BigIron(config-posif-2/1)# exit
```

To configure a protect APS interface:

```
BigIron(config)# interface pos 3/1
BigIron(config-posif-3/1)# aps protect 1 10.0.0.1
BigIron(config-posif-3/1)# exit
```

To add the working interface to a VLAN

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# tagged pos 2/1
```

Configuration Considerations for Layer 2 POS APS

The following configuration considerations apply when using APS with Layer 2 POS links:

- If you are configuring APS on a 15-slot Chassis device, both the working and the protect interface must be located in the same set of slots (slots 1 – 7 or 9 – 15). If the working interface is located in a slot from 1 – 7, the protect interface must also be located in a slot from 1 – 7; if the working interface is located in slot 9 – 15, the protect interface must also be located in a slot from 9 – 15. If you attempt to configure the working and

protect interfaces in different sets of slots, an error message is displayed. Neither the working nor the protect interface can reside in slot 8.

- The protect interface cannot have any Layer 2 commands in its configuration. If a POS interface has any Layer 2 commands in its configuration, then it cannot be specified as a protect interface. The protect interface inherits its Layer 2 configuration from the working interface.
- Layer 2 POS trunk groups and Layer 2 POS APS are mutually exclusive. An APS interface cannot be added to a trunk group, and a POS interface that is a member of a trunk group cannot be configured as an APS interface.
- APS for Layer 2 POS is supported only in single-device configurations, where the working and protect interfaces both reside on the same device. Multiple-device APS configurations are not supported.
- The **aps protect** command requires the IP address of the router where the working interface resides. If the device is running switch code, you specify the management IP address for this purpose.
- APS for Layer 2 POS is supported on OC-3, OC-12, and OC-48 POS modules. Both the working and protect ports must be on the same module type.

Displaying Layer 2 POS Port Information

To display detailed information for a POS port, enter a command such as the following at any level of the CLI:

```
BigIron(config)# show interface pos 2/1

POS2/1 is up
  No port name
  Hardware is Packet over Sonet
  MTU 4470 bytes, encapsulation PPP, clock is internal
  Framing is SONET, BW 622000Kbit, CRC 32
  Loopback not set, keepalive is set (10 sec), scramble disabled
  LCP state is open
  5 minute input rate: 0 bits/sec, 0 packets/sec
  5 minute output rate: 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 no buffer
  Received 0 CRCs, 0 shorts, 0 giants, 0 alignments
  0 packets output, 0 bytes, 0 underruns
  Line protocol is UP
Member of 4 L2 VLANs, port is tagged, port state is FORWARDING
STP configured to ON
```

The lines shown in bold type in this example appear only if the port is enabled for Layer 2 switching.

Displaying POS Information

You can display the following POS module information:

- Software version – see “Displaying the Software Version Running on the Module” on page 7-26 and “Displaying the Software Installed in Flash Memory” on page 7-28
- Displaying general module information – “Displaying General Module Information” on page 7-28
- Module status – see “Determining POS Module Status” on page 7-28
- Interface parameters – see “Displaying Interface Parameters” on page 7-30
- POS statistics – see “Displaying POS Statistics” on page 7-33

Displaying the Software Version Running on the Module

To display the software version running on the POS module, use either of the following methods.

USING THE CLI

To display the software version running on the POS module, enter the following command at any CLI level:

```
BigIron> show version

SW: Version 07.1.05T1 Copyright (c) 1996-1999 Foundry Networks, Inc.
   Compiled on Sep 29 2000 at 17:10:51 labeled as B2R07105
   (1357024 bytes) from Primary b2r07105.car
HW: Chassis 4000 Router, SYSIF version 21
=====
SL 3: B8GMR Fiber Management Module, ACTIVE
  2048 KB BRAM, SMC version 1, ICBM version 21
  512 KB PRAM(512K+0K) and 2048*8 CAM entries for DMA 8, version 0209
  512 KB PRAM(512K+0K) and shared CAM entries for DMA 9, version 0209
  512 KB PRAM(512K+0K) and 2048*8 CAM entries for DMA 10, version 0209
  512 KB PRAM(512K+0K) and shared CAM entries for DMA 11, version 0209
=====
SL 4: B2P155 POS Module
2048 KB BRAM, SMC version 1, ICBM version 21
=====
  240 MHz Power PC processor 603 (revision 7) 63 MHz bus
  512 KB boot flash memory
  8192 KB code flash memory
  256 KB SRAM
  128 MB DRAM
The system uptime is 45 seconds
The system : started=warm start   reloaded=by "reload"
```

Syntax: show version

The command shows all the software versions running on the device. The POS information is shown in this example in bold text.

USING THE WEB MANAGEMENT INTERFACE

You cannot display the POS module software versions using the Web management interface.

Displaying the Software Installed in Flash Memory

To display the software images installed in the POS modules' flash memory, enter the following command:

```
BigIron# show flash
POS module slot 6 CPU 1:
Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304
Boot Flash Type: AMD 29F010, Size: 8 * 16384 = 131072
Compressed Pri Code: size = 805873 Version 07.5.00b1T61 (p2r07500b1.bin)
Compressed Sec Code: size = 805873 Version 07.5.00b1T61 (p2r07500b1.bin)
Maximum Code Image Size Supported: 2096640 (0x001ffe00)
Boot Image size = 30492 Version 06.00.00
Maximum Boot Image Size Supported: 131072 (0x00020000)

POS module slot 6 CPU 2:
Code Flash Type: AMD 29F032B, Size: 64 * 65536 = 4194304
Boot Flash Type: AMD 29F010, Size: 8 * 16384 = 131072
Compressed Pri Code: size = 805873 Version 07.5.00b1T61 (p2r07500b1.bin)
Compressed Sec Code: size = 805873 Version 07.5.00b1T61 (p2r07500b1.bin)
Maximum Code Image Size Supported: 2096640 (0x001ffe00)
Boot Image size = 30492 Version 06.00.00
Maximum Boot Image Size Supported: 131072 (0x00020000)
```

Syntax: show flash

The lines shown in bold type list the software installed on the module.

Displaying General Module Information

To display general module information, use the following method.

USING THE CLI

To display general information for the POS module, enter the following command at any CLI level:

Syntax: show pos

This command displays the state of the POS module, the software version running on the module, the contents of the primary and secondary flash on the module, the system uptime, and the status of the CPUs on the module.

USING THE WEB MANAGEMENT INTERFACE

You cannot display general POS module information using the Web management interface.

Determining POS Module Status

You can determine the status of a POS module in the following ways:

- Status LEDs – Each POS port has LEDs that show link status, transmit and receive activity, and indicate whether an alarm condition has occurred.
- Module information in software – The module information displayed by the software indicates whether the module came up properly.

Status LEDs

You can determine the status of a POS port by observing its LEDs. Each POS port has the following LEDs.

Table 7.4: Port LED Indicators for POS Modules

LED	Position	State	Meaning
Link	Upper left	On	Port is connected.
		Off	No port connection exists.
Alarm	Upper right	On	At least one of the following SONET alarm conditions has been detected: <ul style="list-style-type: none"> • LOS – Loss of Signal • LOF – Loss of Frame • LOP – Loss of Pointer • AIS – Alarm Indication Signal
		Off	None of the alarm conditions listed above have been detected.
TxAct	Lower left	Blinking	The port is transmitting traffic.
RxAct	Lower right	Blinking	The port is receiving traffic.

Software

You can display status information for a POS module using either of the following methods.

NOTE:

- Slots on a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots on an 8-slot chassis are numbered 1 – 4, from top to bottom.
- Slots on a 15-slot chassis are numbered 1 – 15, from left to right.

USING THE CLI

To display the status of a POS module using the CLI, enter the following command at any CLI level:

```
BigIron> show module

Module                               Status  Ports Starting MAC
S1: B0GMR Management Module          ACTIVE    0
S2: BI POS 622M Module                OK        2  00e0.5281.eb20
S3: B24E Copper Switch Module        OK       24  00e0.5281.eb40
S4: B8G Fiber Switch Module           OK        8  00e0.5281.eb60
```

Syntax: show module

NOTE: The module descriptions do not distinguish between SX and LX ports.

The Status column shows the module status. A POS module can have one of the following statuses:

- FAILED – This status indicates that the host module failed to come up.

- OK – This status indicates that the module came up and is operating normally.

NOTE: Management modules have different status values.

USING THE WEB MANAGEMENT INTERFACE

1. Select the Home link to display the System configuration sheet, if not already displayed.
2. Select the Module link to display the Module panel. The Status column shows the module status. A POS module can have one of the following statuses:
 - FAILED – This status indicates that the host module failed to come up.
 - OK – This status indicates that the module came up and is operating normally.

Displaying Interface Parameters

To display the current settings for the POS interface parameters, use the following method.

USING THE CLI

To display interface information using the CLI, enter the following command at any CLI level:

```
BigIron> show interface brief slot 2

Port Link State  Encap Clock Loop Speed  mtu  frame  scram  crc  c2 j0 h1
2/1  up           ppp  int   none 622   1200  sonet  no    16  cf cc 00
2/2  up           ppp  int   none 155   1500  sonet  no    16  cf cc 00
```

Syntax: show interface brief [pos <portnum>] | [slot <slotnum>]

This command shows information for all the ports in the chassis. For simplicity, the example and syntax above show only the information relevant to POS ports.

The command shows the following information for POS ports.

Table 7.5: CLI Display of POS Interface Information

This Field...	Displays...
Port	The chassis slot and port number of the interface.
Link State	The state of the link, which can be one of the following: <ul style="list-style-type: none"> • down • up
Encap	The encapsulation type in use on the interface. The encapsulation type can be one of the following: <ul style="list-style-type: none"> • hdlc – High-Level Data Link Control • ppp – Point-to-Point Protocol. This is the default. To change this parameter, see “Changing the Encapsulation Type” on page 7-10.

Table 7.5: CLI Display of POS Interface Information (Continued)

This Field...	Displays...
Clock	<p>The clock source, which can be one of the following:</p> <ul style="list-style-type: none"> • int – The interface is using the clock on the POS module. • line – The interface is using the clock source supplied on the network. <p>To change this parameter, see “Changing the Clock Source” on page 7-10.</p>
Loop	<p>The loopback state of the interface. The loopback state can be one of the following:</p> <ul style="list-style-type: none"> • int – The loopback path consists only of the POS circuitry on this interface. • line – The loopback path consists of both this POS interface and the POS interface at the remote end of the link. • none – The interface is not operating in loopback mode. <p>To change this parameter, see “Changing the Loopback Path” on page 7-10.</p>
Speed	<p>The bandwidth of the interface, which can be one of the following:</p> <ul style="list-style-type: none"> • 155 • 622 • 2488 <p>To change this parameter on the B2P622 module, see “Changing the Bandwidth” on page 7-12.</p>
MTU	<p>The Maximum Transmission Unit (MTU) of packets sent on this interface. The MTU can be from 60 – 4770 bytes. The default is 4470 bytes.</p> <p>To change this parameter, see “Changing the Bandwidth” on page 7-12.</p>
Frame	<p>The frame type used on the interface. The frame type can be one of the following:</p> <ul style="list-style-type: none"> • sdh – Synchronous Digital Hierarchy. • sonet – Synchronous Optical Network. This is the default. <p>To change this parameter, see “Changing the Frame Type” on page 7-13.</p>
Scram	<p>The state of the ATM scramble mode, which can be one of the following:</p> <ul style="list-style-type: none"> • no – Scrambling is disabled. This is the default. • yes – Scrambling is enabled. <p>To change this parameter, see “Enabling or Disabling ATM Scrambling” on page 7-13.</p>

Table 7.5: CLI Display of POS Interface Information (Continued)

This Field...	Displays...
CRC	<p>The length of the CRC field in packets transmitted on the interface. The length can be one of the following:</p> <ul style="list-style-type: none"> • 16 – The field is 16 bits long. • 32 – The field is 8 bits long. This is the default. <p>To change this parameter, see “Changing the CRC Length” on page 7-11.</p>
c2	<p>The value of the c2 flag in the SONET headers of packets transmitted by the interface. The c2 flag identifies the payload type of the packets transmitted on this interface. The c2 flag is set to 0xcf by default. This value indicates that the payload is SONET or SDH.</p> <p>To change this parameter, see “Changing the POS Flags” on page 7-12.</p>
j0	<p>The value of the j0 flag in the SONET headers of packets transmitted by the interface. This flag sets the trace byte, which is used to trace the origin of an STS-1 frame on a SONET network. This flag is set to 0xcc by default.</p> <p>To change this parameter, see “Changing the POS Flags” on page 7-12.</p>
h1	<p>The value of the h1 flag in the SONET headers of packets transmitted by the interface. This flag sets the H1 pointer, which is used to indicate where the SPE (Synchronous Payload Envelope) starts within the packet. The SPE contains the packet’s payload.</p> <p>This flag can have one of the following values:</p> <ul style="list-style-type: none"> • 0x00 – The pointer for SONET frames. This is the default. • 0x02 – The pointer for SDH frames. <p>To change this parameter, see “Changing the POS Flags” on page 7-12.</p>

To display detailed information for a POS port, enter a command such as the following at any level of the CLI:

```
BigIron(config)# show interface pos 2/1

POS2/1 is up
  No port name
  Hardware is Packet over Sonet
  MTU 4470 bytes, encapsulation PPP, clock is internal
  Framing is SONET, BW 622000Kbit, CRC 32
  Loopback not set, keepalive is set (10 sec), scramble disabled
  LCP state is open
  5 minute input rate: 0 bits/sec, 0 packets/sec
  5 minute output rate: 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 no buffer
  Received 0 CRCs, 0 shorts, 0 giants, 0 alignments
  0 packets output, 0 bytes, 0 underruns
  Line protocol is UP
Member of 4 L2 VLANs, port is tagged, port state is FORWARDING
STP configured to ON
```

The "Line protocol is UP" line means that the Link Control Protocol (LCP) handshake was successful and the end-to-end connectivity is up.

NOTE: The lines shown in bold type in this example appear only if the port is enabled for Layer 2 switching.

Displaying POS Statistics

To display POS packet statistics, use the following method.

USING THE CLI

To display POS statistics for POS interface 2/1, enter the following command at any CLI level:

```
BigIron> show statistics pos 2/1

POS          Packets          Errors
Port  [Receive Transmit]  [Align  FCS  Giant  Short]
2/1      1475    12301      0   1378     3     0
```

Syntax: show statistics pos <portnum>

The command shows the following information for the specified POS port.

Table 7.6: CLI Display of POS Interface Statistics

This Field...	Displays...
Packet counters	
Receive	The number of packets received on this interface.
Transmit	The number of packets transmitted on this interface.
Packet Errors	
These fields show statistics for various types of packet errors. The Layer 3 Switch drops packets that contain one of these errors.	

Table 7.6: CLI Display of POS Interface Statistics (Continued)

This Field...	Displays...
Align	The number of packets that contained frame alignment errors.
FCS	The number of packets that contained Frame Check Sequence errors.
Giant	The number of packets that were longer than the configured MTU.
Short	The number of packets that were shorter than the minimum valid length.

Displaying POS Alarms and Error Conditions

SONET equipment detects alarms and error conditions at the three layers of the SONET protocol: section, line, and path. Other devices on the network are notified of these problems. To determine if any alarms and error conditions have been reported for POS, use the following method.

USING THE CLI

To display any alarm or error condition that have been logged for POS, enter the following command at any CLI level:

```
BigIron# show controller pos

BigIron# show controllers pos

POS 4/1
  BIP(B1): 0  BIP(B2): 0  BIP(B3): 0
  AIS: 0  RDI: 0  LOP: 0  LOF: 298  LOS: 298
POS 4/2
  BIP(B1): 0  BIP(B2): 0  BIP(B3): 0
  AIS: 0  RDI: 0  LOP: 0  LOF: 298  LOS: 298
```

Syntax: show controllers pos

Table 7.7: CLI Display of POS Interface Statistics

This Field...	Displays...
POS slot/port number	The slot number (if applicable) and port number of the POS interface.
BIP(B1)	The number of received frames that has parity errors at the section layer of the SONET link.
BIP(B2)	The number of received frames that has parity errors at the line layer of the SONET link.
BIP(B3)	The number of received frames that has parity errors at the path layer of the SONET link.
AIS	The number of Alarm Indicator Signals (AIS) that were received by the interface.
RDI	The number of Remote Defect Indicator (RDI) signals that were received by the interface.
LOP	Loss of pointer (LOP) condition that resulted from an invalid path pointer or if excessive number of new data flag were enabled.

Table 7.7: CLI Display of POS Interface Statistics (Continued)

This Field...	Displays...
LOF	How many times the interface experienced an out of frame alignment problems, which is also called a loss of frame (LOF) condition.
LOS	The number of times POS interfaces experienced a loss of signal (LOS). With LOS, incoming signals are all zeros during a 100 microsecond period.

Configuring Automatic Protection Switching (APS)

The Automatic Protection Switching (APS) feature provides redundancy for a POS link. Foundry's implementation of POS APS supports 1 + 1 protection architecture, where a **working** interface is paired with a **protect** or backup interface.

When the signal on the working interface degrades below a user-configurable threshold, or the working interface is manually taken down, the protect interface automatically takes over, becoming the working interface, while the previously working interface becomes the protect interface. When the previously working interface is available again, it can be switched back to working status manually or configured to automatically revert to working status after becoming available.

Typically the working and protect interfaces are on separate routers. Both interfaces are connected to a SONET add/drop multiplexer (ADM), which sends identical traffic through them. Switching is controlled by the K1 and K2 bytes of the line overhead (LOH) in a SONET frame.

Information on signal quality is exchanged between the working and protect interface using the APS Protect Group Protocol, running on top of UDP. This communication takes place on a channel independent of the working and protect interfaces themselves.

Configuring POS APS consists of the following steps:

1. Configuring the working interface
2. Configuring the protect interface
3. Setting optional parameters

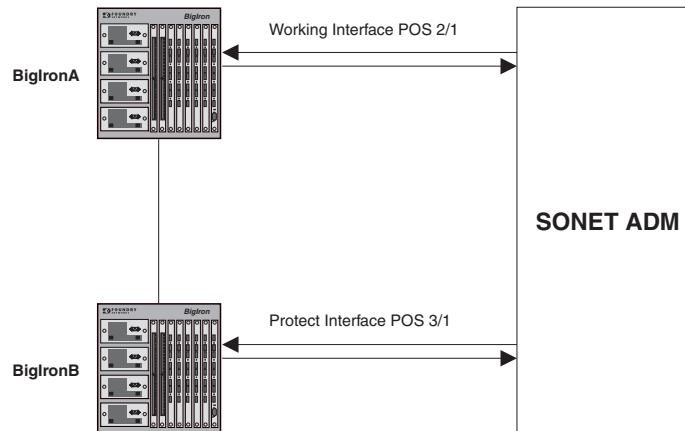
You can also display information about an APS configuration.

NOTE: Starting in release 07.6.03, APS is supported on Layer 2 POS links. See "Configuring Automatic Protection Switching (APS) for Layer 2 POS" on page 7-25 for more information.

Basic POS APS Configuration

Figure 7.7 shows a basic POS APS configuration, where a POS interface on one router serves as a protect interface for a working interface on another router.

Figure 7.7 Basic POS APS configuration



In this configuration, the ADM sends identical traffic through the working interface on BigIronA and the protect interface on BigIronB. If the working interface malfunctions (that is, experiences high bit error ratios (BERs), significant signal degradation, or signal failure) or is taken down manually, the protect interface automatically takes over.

The following commands configure the working interface on BigIronA:

NOTE: Foundry Networks recommends that you configure the working interface prior to configuring the protect interface, so that the protect interface does not inadvertently become the working interface.

```
BigIronA(config)# interface loopback 1
BigIronA(config-lbif-1)# ip address 10.0.0.1/24
BigIronA(config-lbif-1)# exit

BigIronA(config)# interface pos 2/1
BigIronA(config-posif-2/1)# aps working 1
BigIronA(config-posif-2/1)# exit
```

Syntax: `aps working <circuit-number>`

The **aps working** command establishes this interface as the working interface in circuit 1. This working interface corresponds to a protect interface on the BigIronB.

The following commands configure the protect interface on BigIronB:

```
BigIronB(config)# interface loopback 2
BigIronB(config-lbif-2)# ip address 10.0.0.2/24
BigIronB(config-lbif-2)# exit

BigIronB(config)# interface pos 3/1
BigIronB(config-posif-3/1)# aps group 1
BigIronB(config-posif-3/1)# aps protect 1 10.0.0.1
BigIronB(config-posif-3/1)# exit
```

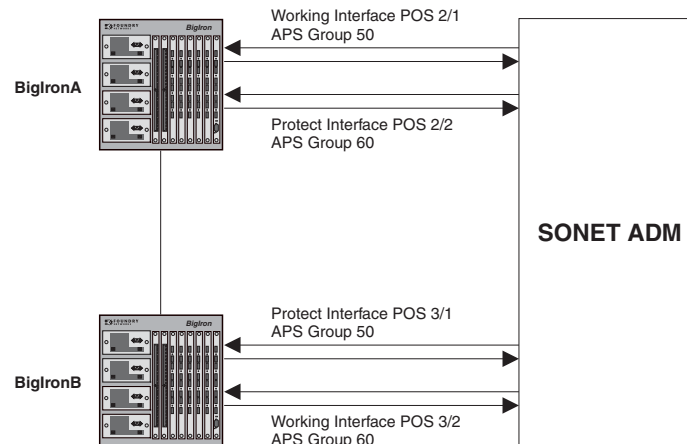
Syntax: `aps protect <circuit-number> <ip-addr>`

The **aps protect** command specifies the circuit this interface is protecting, as well as the IP address of the router where the working interface resides. This is normally the loopback address of the router. If the device is a switch, this is the management IP address.

Multi-Group APS Configuration

A router can have more than one protect or working interface. To configure more than one protect or working interface on a router, you assign each interface to a group using the **aps group** command. Figure 7.8 illustrates a configuration with two APS working/protect circuit pairs.

Figure 7.8 Configuration with multiple POS APS interfaces



In this configuration, interface 3/1 on BigIronB serves as the protect interface to working interface 2/1 on BigIronA, and interface 2/2 on BigIronA serves as the protect interface to working interface 3/2 on BigIronB. To implement this configuration, you place each APS working/protect circuit pair in a separate APS group.

The following commands configure the working interface for APS group 50 and protect interface for APS group 60 on BigIronA:

```
BigIronA(config)# interface loopback 1
BigIronA(config-lbif-1)# ip address 10.0.0.1/24
BigIronA(config-lbif-1)# exit

BigIronA(config)# interface pos 2/1
BigIronA(config-posif-2/1)# aps group 50
BigIronA(config-posif-2/1)# aps working 1
BigIronA(config-posif-2/1)# exit

BigIronA(config)# interface pos 2/2
BigIronA(config-posif-2/2)# aps group 60
BigIronA(config-posif-2/2)# aps protect 1 10.0.0.2
BigIronA(config-posif-2/2)# exit
```

The following commands configure the working and protect interfaces on BigIronB:

```
BigIronB(config)# interface loopback 2
BigIronB(config-lbif-2)# ip address 10.0.0.2/24
BigIronB(config-lbif-2)# exit

BigIronB(config)# interface pos 3/1
BigIronB(config-posif-3/1)# aps group 50
BigIronB(config-posif-3/1)# aps protect 1 10.0.0.1
BigIronB(config-posif-3/1)# exit

BigIronB(config)# interface pos 3/2
BigIronB(config-posif-3/2)# aps group 60
BigIronB(config-posif-3/2)# aps working 1
BigIronB(config-posif-3/2)# exit
```

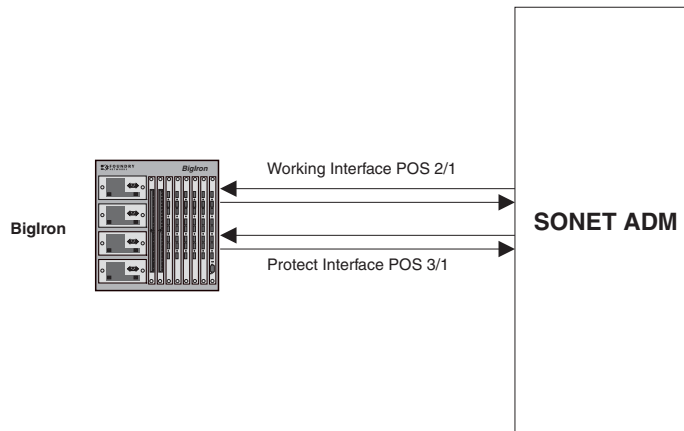
Syntax: `aps group <group-number>`

The **aps group** command allows more than one APS working/protect interface pair to be configured on the device. The default group number is 0. The **aps group** command is required on both the working and protect interfaces.

Single-Device APS Configuration

The working and protect interfaces can reside on the same device. Figure 7.9 illustrates this kind of configuration.

Figure 7.9 Single-device POS APS configuration



The working interface resides on POS interface 2/1, and the protect interface resides on POS interface 3/1. The working and protect interfaces can also reside on the same module.

The following commands configure the working and protect interfaces in Figure 7.9:

```
BigIron(config)# interface loopback 1
BigIron(config-lbif-1)# ip address 10.0.0.1/24
BigIron(config-lbif-1)# exit

BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# aps working 1
BigIron(config-posif-2/1)# exit

BigIron(config)# interface pos 3/1
BigIron(config-posif-3/1)# aps protect 1 10.0.0.1
BigIron(config-posif-3/1)# exit
```

Starting with release 07.5.00, the single-device POS APS configuration was enhanced to allow faster switchover when a protect interface becomes a working interface. The working interface and protect interface can be configured to behave as one logical interface. When a link switchover occurs, this logical interface stays up, and only the association of the logical interface to the physical interface changes. Before the link switchover, the logical interface is associated with the working interface; after the link switchover, the logical interface is associated with the protect interface. When the protect interface becomes active, it inherits the configuration of the working interface.

The following diagrams illustrate this kind of single-device POS APS configuration. Figure 7.10 depicts the configuration before the link switchover, and Figure 7.11 depicts the configuration after the link switchover.

Figure 7.10 Single-device POS APS configuration (before link switchover)





Figure 7.11 Single-device POS APS configuration (after link switchover)



In this example, the working APS interface on BigIron A, interface POS 2/1, has an IP address of 10.10.1.1/24. The IP address of the POS interface on the neighboring BigIron B is 10.10.1.2/24. No IP address is configured on the protect interface on BigIron A, interface POS 3/1.

Before a link switchover occurs (Figure 7.10), BigIron A and B form an adjacency and exchange routes. For traffic entering BigIron A that uses a route for which the next hop router is 10.10.1.2/24 (BigIron B), the outgoing interface is interface 2/1. On the receive side, the ADM forwards the same data signal onto both the working and the protect lines. Traffic arriving on the working interface is accepted, while traffic arriving on the protect interface is dropped.

After a link switchover occurs (Figure 7.11), BigIron A reacts by keeping the IP interface 10.10.1.1/24 up, so that the link switchover is transparent for neighboring BigIron B. The routing adjacency formed between BigIron A and BigIron B is not disturbed. Any traffic entering BigIron A that was using BigIron B as the next hop is now forwarded out interface 3/1. Also, on the receive side, BigIron A starts accepting packets on interface 3/1 and dropping packets on interface 2/1. In this way, network recovery is achieved very quickly, without the need for routing re-convergence.

To set up this APS configuration, on the working interface you configure the IP address, routing protocols, and encapsulation type. For example:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# ip address 10.10.10.1/24
BigIron(config-posif-2/1)# ip ospf area 1
BigIron(config-posif-2/1)# encap hdlc
BigIron(config-posif-2/1)# aps working 1
```

On the protect interface, you do not specify an IP address, routing protocol parameters, or encapsulation type. This information is inherited from the working interface when a link switchover occurs. All other non-APS configuration statements on the working interface must be duplicated on the protect interface, however.

```
BigIron(config)# interface pos 3/1
BigIron(config-posif-3/1)# no keepalive
BigIron(config-posif-3/1)# aps protect 1 1.1.1.1
BigIron(config-posif-3/1)# aps revert 1
```

In this APS configuration, there are two separate physical interfaces, but only one logical IP interface. If the encapsulation is PPP, there would be only one PPP interface, on which the line protocol would be based. This PPP interface would be tied to the working interface, so to check the status of the PPP interface, you examine the line protocol status on the working interface. The line protocol on the protect interface would always remain down.

When you enter the **show interface pos** command for the port, the port status would reflect the actual status of the physical port, and the line protocol would reflect the status of the PPP interface. For example, if POS 2/1 is the working interface and POS 3/1 is the protect interface, and POS 2/1 is up and the line protocol is up, the **show interface pos** command would display the following:

```
BigIron(config)# show interface pos 2/1
...
Port is up, line protocol is up
```

When port 2/1 goes down and the link is switched over to the to the protect line, the **show interface pos** command would display the following:

```
BigIron(config)# show interface pos 2/1
...
Port is down, line protocol is up
```

At this time, the **show interface pos** command for the protect interface would display the following:

```
BigIron(config)# show interface pos 3/1
...
Port is up, line protocol is down
```

Configuring Optional Parameters

You can configure optional POS APS parameters to do the following:

- Configure an authentication string for communication between the process controlling the working interface and the process controlling the protect interface
- Force a protect interface to take over as a working interface
- Prevent a protect interface from taking over from a working interface
- Manually cause a switchover from a working interface to a protect interface
- Configure the WTR (wait-to-restore) interval for a working interface
- Set the timers that the protect interface uses for sending hello packets and waiting for a response from the working interface
- Set thresholds for bit error rate, signal degradation, and signal failure on a POS interface

Configuring an Authentication String

The working and protect interfaces are synchronized using the APS Protect Group Protocol, which provides communication between the process controlling the working interface and the process controlling the protect interface. This communication takes place on a channel independent of the working and protect interfaces themselves.

You can specify an authentication string that must be part of each packet sent between the process controlling the working interface and the process controlling the protect interface. To do so, enter commands such as the following:

```
BigIronA(config)# interface pos 2/1
BigIronA(config-posif-2/1)# aps working 1
BigIronA(config-posif-2/1)# aps authenticate mulvaney
BigIronA(config-posif-2/1)# exit

BigIronB(config)# interface pos 3/1
BigIronB(config-posif-3/1)# aps protect 1 10.0.0.1
BigIronB(config-posif-3/1)# aps authenticate mulvaney
BigIronB(config-posif-3/1)# exit
```

NOTE: The same authentication string must be configured on both the working and protect interfaces.

Syntax: `aps authenticate <string>`

Forcing an APS Switchover

You can manually force a protect interface to take over as a working interface by entering the **aps force** command on the protect interface. This command is useful when you want to bring down a working interface for maintenance purposes.

For example:

```
BigIron(config)# interface pos 3/1
BigIron(config-posif-3/1)# aps protect 1 10.0.0.1
BigIron(config-posif-3/1)# aps force 1
BigIron(config-posif-3/1)# exit
```

Syntax: `aps force <circuit-number>`

The `<circuit-number>` is a valid POS APS circuit number. In addition, you can specify 0 as the `<circuit-number>` (**aps force 0**) to manually force traffic from the protect interface to the working interface.

The switchover takes place immediately after you enter the command. The **aps force** command is not saved if you write the active configuration to memory.

Preventing an APS Switchover

You can prevent a protect interface from taking over from a working interface by using the **aps lockdown** command. If the working interface malfunctions, no switchover to the protect interface takes place.

To prevent a protect interface from becoming a working interface, enter commands such as the following on the protect interface:

```
BigIron(config)# interface pos 3/1
BigIron(config-posif-3/1)# aps protect 1 10.0.0.1
BigIron(config-posif-3/1)# aps lockdown
BigIron(config-posif-3/1)# exit
```

Syntax: `aps lockdown`

Manually Causing an APS Switchover

To manually cause a switchover from a working interface to a protect interface, use the **aps manual** command. This command can be used when you want to bring a working interface down for maintenance, or if you want to bring an interface back up without waiting for the WTR (wait-to-restore) interval specified by the **aps revert** command.

For example, to cause traffic on a working interface to switch over to the protect interface, enter commands such as the following:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# aps working 1
BigIron(config-posif-2/1)# aps manual 1
BigIron(config-posif-2/1)# exit
```

Syntax: `aps manual <circuit-number>`

The `<circuit-number>` is a valid POS APS circuit number. In addition, you can specify 0 as the `<circuit-number>` (**aps manual 0**) to manually force traffic from the protect interface to the working interface.

Configuring the Wait-To-Restore Interval for a Working Interface

You can configure an interface to automatically revert to being a working interface after it has been available for a specified amount of time. For example, to cause an interface to switch back to being a working interface 3 minutes after becoming available, enter commands such as the following:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# aps protect 1
BigIron(config-posif-2/1)# aps revert 3
BigIron(config-posif-2/1)# exit
```

Syntax: `aps revert <minutes>`

Setting POS APS Timers

You can configure the interval at which the process controlling the protect interface sends hello packets to the process controlling the working interface, as well as how long the process controlling the protect interface waits for a response before declaring the working interface down.

For example, to configure the protect interface process to send hello packets every 3 seconds and wait a maximum of 6 seconds for a response, enter commands such as the following on the protect interface:

```
BigIron(config)# interface pos 3/1
BigIron(config-posif-3/1)# aps protect 1 10.0.0.1
BigIron(config-posif-3/1)# aps timers 3 6
BigIron(config-posif-3/1)# exit
```

Syntax: `aps timers <hello-timer> <response-timer>`

The `<hello-timer>` is the interval between hello packets. The default is 1 second.

The `<response-timer>` is the amount of time the protect interface process waits for a response from the working interface process before declaring the working interface down. The default is 3 seconds.

Setting Thresholds for POS APS

A switchover from the working interface to the protect interface can occur due to poor signal quality, bit rate errors, or link failure. You can set the thresholds for when an APS switchover occurs.

For example, to set thresholds for bit error rate, signal degradation, and signal failure on a POS interface:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# pos threshold b1-tca 9
BigIron(config-posif-2/1)# pos threshold sd-ber 5
BigIron(config-posif-2/1)# pos threshold sf-ber 5
BigIron(config-posif-2/1)# exit
```

Syntax: `pos threshold <alarm> <threshold>`

The `<alarm>` parameter can be one of the following:

- b1-tca** B1 bit error rate alarm.
- b2-tca** B2 bit error rate alarm.
- b3-tca** B3 bit error rate alarm.
- sd-ber** Signal degradation alarm.
- sf-ber** Signal failure alarm.

The `<threshold>` parameter is the bit error rate from 3 – 9. For all alarms except **sf-ber**, the default is 6 (10e-6). For the **sf-ber** alarm, the default is 3 (10e-3).

Displaying POS APS Information

To display information about a POS APS configuration, use the `show aps` command. For example:

```
BigIron# show aps
POS2/1 working group 1 channel 1 Enabled Selected
```

Syntax: `show aps`

In this example, the output indicates that POS interface 2/1 is the working interface for channel 1 in APS group 1, and the interface is active. If there is a tilde next to Selected (for example, `~Selected`) it means the interface is not active.

Path Trace (J1 Byte) Field Support

The Path Trace (J1 byte) field in the SONET payload envelope (SPE) transmits a 64-byte, fixed-length string that the receiving Path Terminating Equipment can use to verify its connection to the device that sent the SPE. In this release, you can configure the string a POS interface transmits in the Path Trace field. You can also display Path Trace strings received and configured on the POS interface.

To specify a string to be transmitted in the Path Trace field for POS interface 2/1:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# pos flag j1 Foundry_BigIron
```

Syntax: `[no] pos flag j1 <string>`

To clear the user-configured string:

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# no pos flag j1
```

Syntax: no pos flag j1

The string can be up to 62 bytes long. If you do not explicitly configure a Path Trace string, the Foundry device transmits a Path Trace string consisting of a concatenation of the hostname of the device, interface number, and IP address of the interface.

To display the Path Trace strings configured and received on POS interface 2/1:

```
BigIron(config)# show interface pos 2/1
POS2/1 is up, line protocol is up
  No port name
  Hardware is Packet over Sonet
  Internet address is 12.1.1.1/24
  Peer Internet address is 0.0.0.0
  MTU 4470 bytes, encapsulation PPP, clock is internal
  Framing is SONET, BW 155000Kbit, CRC 16
  Loopback not set, keepalive is set (10 sec), scramble enabled
  LCP state is opened, IPCP state is opened
  5 minute input rate: 80 bits/sec, 0 packets/sec
  5 minute output rate: 184 bits/sec, 0 packets/sec
  4444 packets input, 83530 bytes, 0 no buffer
  Received 14 CRCs, 0 shorts, 1 giants, 0 alignments
  4373 packets output, 164880 bytes, 0 underruns
Configured Path Trace String : Foundry_BigIron
Received Path Trace String : fdrycanuread
```

Syntax: show interface pos <portnum>

The **show interface pos** <portnum> command displays the configured (as opposed to transmitted) and received Path Trace strings. When a Path Trace string has been explicitly configured, the configured and transmitted Path Trace strings are the same. However, when a Path Trace string has not been explicitly configured, the Foundry device transmits a string that is the concatenation of the hostname, interface number, and the IP address of the interface. In this case, the configured and transmitted Path Trace strings are different.

XENPAK WAN PHY Transceiver

The XENPAK WAN PHY transceiver enables a 10 Gbps ethernet port to use SONET/SDH for Layer 1 transport across a WAN transport backbone. The XENPAK WAN PHY transceiver can be directly linked to an OC192 SONET-based port. This allows for the extension of ethernet links across a WAN transport backbone.

The XENPAK WAN PHY transceiver can be set-up for either WAN or LAN mode. The default state is LAN mode. To determine the mode of an individual transceiver, use the **show media** command as shown below. XENPAKs of the type "10G-LW," (as shown for port 12/2) are configured for WAN mode. XENPAKs of the type "10G-LR" (as shown for ports 12/1) are configured for LAN mode.

```
BigIron#show media
Port 12/1:10G-LR (XENPAK) 12/2: 10G-LW (XENPAK)
```

Setting the WAN PHY XENPAK to WAN Mode

The **x10g-phy-wan** command allows you to change to WAN mode from LAN mode.

EXAMPLE:

If you want to change the mode of a WAN PHY XENPAK to WAN mode, use the following command:

```
BigIron#interface ethernet 6/3
BigIron#(config-if-e10000-6/3)# x10g-phy-wan
```

Syntax: x10g-phy-wan

This command changes the mode of a WAN PHY XENPAK to WAN mode.

Setting WAN PHY Link Fault Signaling

WAN PHY link fault signaling can be enabled for 10 Gbps ethernet and SONET using the **link-fault-signaling** command as described in the following.

EXAMPLE:

Use the following command to enable link fault signaling for 10 Gbps ethernet and SONET:

```
BigIron#interface ethernet 6/3
BigIron#(config-if-e10000-6/3)# link-fault-signaling sonet
```

Syntax: link-fault-signaling <sonet>

This command enables 802.3ae link fault signaling for 10 Gbps ethernet.

sonet enables link fault signaling for 10 Gbps ethernet and SONET.

Chapter 8

Using Asynchronous Transfer Mode Modules

This chapter describes the Foundry Asynchronous Transfer Mode (ATM) modules and how to configure and manage them.

Overview


Foundry supports the following types of ATM modules. You can use the modules in a NetIron Internet Backbone router chassis or a BigIron chassis running Layer 3 Switch software:

- Model B2N155 – Contains two OC-3 ports providing 155 Mbps each.
- Model B4N155 – Contains four OC-3 ports providing 155 Mbps each.

NOTE: You can use the ATM modules only in a device running Layer 3 Switch software and using a Management 2 module or higher.

Figure 8.1 shows the front panel of a B2N155 ATM module.

Figure 8.1 ATM Module



Foundry's implementation of ATM is based on RFC 2684, which describes how to encapsulate multi-protocol data over ATM Adaptation Layer 5 (AAL5). Foundry's implementation supports encapsulation of IP over ATM, using the Logical Link Control (LLC) method described in RFC 2684. Foundry's implementation also supports User-Network Interface (UNI) version 3.1.

Foundry's implementation supports fragmentation. Foundry ATM PVCs support Maximum Transmission Unit (MTU) sizes from 1500 – 9180 bytes, in accordance with RFC 1626, "Default IP MTU for use over ATM AAL5". The default MTU is 4470 bytes. ATM fragments packets above 1500 bytes of payload. Each fragment except the first and last has a payload length of 1440 bytes with an extra 20-byte IP header.

NOTE: The MTUs on both sides of an ATM PVC must match.

NOTE: The Foundry implementation does not support virtual channel (VC) MUXing or ATM encapsulation of non-IP protocols such as AppleTalk or IPX.

Foundry ATM Interface Specifications

Table 8.1 lists the fiber specifications for Foundry ATM interfaces.


Table 8.1: ATM Fiber Specifications

Transceiver	Power Budget	Launch Window	Transmit Power	Receive Power	Maximum Distance
OC-3c interfaces					
Single-mode short-reach	13 dB	1270 to 1380 nm	-28 to -8 dBm	-31 to -8 dBm	9.75 miles (15 Km)
Single-mode intermediate-reach	29 dB	1280 to 1335 nm	-5 to 0 dBm	-34 to -8 dBm	26 miles (40 Km)
Multimode	11.5 dB	1270 to 1380 nm	-18 to -14 dBm	-30 to -14 dBm	1.3 miles (2 Km)

Virtual Channel (VC) Support

Foundry ATM modules provide ATM interfaces with other ATM devices through virtual channels (VCs). A VC is a dedicated circuit between the two end points of an ATM connection. A VC consists of a virtual path ID (VPI) and a virtual channel ID (VCI). Each VC must be unique for a given ATM port, but you can use the same VC (combination of VPI and VCI) on other ATM ports on the same module or in the same chassis. Figure 8.2 shows an example of a VC.

Figure 8.2 Example of a VC



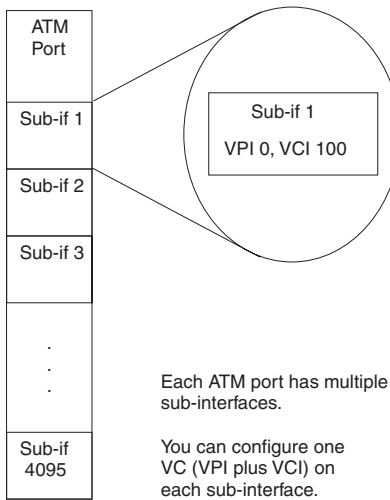
Each VPI contains multiple VCIs.

A VC consists of a VPI and VCI.

When you configure an ATM port, you configure VCs within sub-interfaces on the ATM port. Each sub-interface is a separate point-to-point or point-to-multipoint router interface. You can run IP, RIP, OSPF, IS-IS, and BGP4 over the sub-interfaces.

NOTE: The current release does not support point-to-multipoint links or Non-Broadcast Multi-Access (NBMA).

Figure 8.3 shows the configuration layers on a Foundry ATM port.

Figure 8.3 Configuration hierarchy for ATM interfaces

As shown in Figure 8.3, the following hierarchy applies to configuration of ATM interfaces on a Foundry device:

- Each chassis can contain multiple ATM modules.
- Each ATM module has two or four ATM ports.
- Each ATM port has multiple sub-interfaces.
- Each sub-interface has one VC (one VPI, VCI pair).
- Each sub-interface has one IP interface, associated with one VC on the port.

Table 8.2 lists the default maximum number of VCs and related parameters you can configure.

Table 8.2: ATM Interface Parameters – Maximums

Parameter	Maximum
sub-interface	up to 4095 per port, with a per chassis maximum of 4095
VPI	4 per port
VCI	1024 per VPI
IP subnet address	1 per sub-interface; thus up to 4095 per port, with a per chassis maximum of 4095
ATM router interface	1 per sub-interface; thus up to 4095 per port, with a per chassis maximum of 4095 ^a

a. An ATM router interface is a completely configured interface including the VC information, Class of Service (CoS) information, and IP subnet address.

You can change the maximum number of VCs per VP if needed. See “Changing the Number of Virtual Channels (VCs) per Virtual Path (VP)” on page 8-10.

Class of Service (CoS) Support

When you configure a VC on a Foundry ATM port, you also specify the Class of Service (CoS) for the VC. The ATM CoS specifies how the ATM device classifies and prioritizes outbound traffic on the device. Each CoS type is designed to provide optimal service to a specific type of traffic (voice, data, streaming media, and so on).

ATM provides the following types of CoS:

- Constant Bit Rate (CBR)
- Variable Bit Rate (VBR)
- Unspecified Bit Rate (UBR)
- Available Bit Rate (ABR)

You configure CoS on an individual basis for each VC. Thus, each VC has its own CoS settings.

Foundry ATM modules support CBR, VBR, and UBR. These CoS methods are described in the following sections. ABR is not supported and thus is not described.

CBR

CBR is the highest-priority service you can configure. CBR provides a steady, unchanging amount of bandwidth to the VC. CBR is especially suitable for voice-over-IP and other traffic that requires steady, guaranteed throughput.

When you configure CBR, you specify the Peak Cell Rate (PCR). The PCR is the maximum number of bits of data (in kilobits per second) the device sends on the VC each second. The device allocates the corresponding fraction of the available cells per second to the PVC.

If the PVC does not have enough traffic to use its allocated bandwidth, the device sends null cells in the interface to maintain the CBR.

For CBR, the PCR is also the constant data rate since the port sends all the data at the PCR.

VBR

VBR is the second-highest priority service you can configure. VBR provides bandwidth only as needed up to a specified maximum rate, and can provide throughput in excess of the specified rate for short periods of time to accommodate bursts of heavy traffic. VBR is especially suitable for LANs and other environments that have a lot of “bursty” data traffic. “Bursty” traffic is traffic that typically consists of short bursts of irregular duration and byte count. LAN traffic tends to be bursty because the traffic consists mainly of data files of various sizes sent at various times.

VBR also is well suited for compressed video and audio traffic, which is bursty in a cyclic way. This type of traffic consists of bursts with regular intervals in between. End stations buffer traffic from the bursts to present a constant stream of video or audio to the end user.

NOTE: Real-time VBR is not supported.

When you configure VBR, you specify the PCR as well as the Sustained Cell Rate (SCR), and Maximum Burst Size (MBS). The PCR is described in the previous section. The SCR is the average data rate for the VC (expressed in bits per second), during normal traffic loads. The MBS is the maximum number of bits over the SCR that the device can send on the VC before exceeding the limitations of the VBR definition for the VC.

When the PVC starts, it transmits cells at the PCR until the PVC has sent MB cells at the PCR. Once the MB is reached, the PVC transmission rate is limited to the SCR. However, when the PVC is not transmitting at PCR, and has an opportunity to transmit a cell but does not use it because there is no traffic, the PCR burst size is allowed to recover, up to the configured MBS. Thus, VBR supports bursts within the shared bandwidth resources of the port.

UBR

UBR is the lowest-priority service you can configure. UBR does not have any traffic parameters. Thus, you do not specify the bandwidth for the PVC. Instead, when you enable a PVC for UBR, the PVC uses the cells that are left over after CBR or VBR has used the cells they need. Otherwise, the only limit to UBR is the total bandwidth of the

port. If a port has multiple PVCs that are configured for UBR, the PVCs divide the available bandwidth evenly among themselves.

NOTE: In software release 07.6.01, a PVC configured for UBR makes its unused bandwidth available to other PVCs configured for UBR. In earlier releases, unused bandwidth on a PVC configured for UBR is not available to other PVCs.

Access Control List Support

You can apply IP ACLs to an ATM sub-interface.

NOTE: Filters configured using the **ip access-policy** command are not supported. To filter traffic on an ATM sub-interface, use ACLs.

Installing an ATM Module

To install an ATM module, perform the following tasks:

- Configure the chassis slot to receive the module.
- Insert the module.

Configuring the Chassis to Receive the Module

When you plan to insert a module into a chassis slot, you first must configure the slot to receive the module unless the slot already contains the same type of module.

NOTE: If you are swapping out another module, you must disable the module before removing it from the Chassis device. See “Removing the Old Module” on page 2-40.

USING THE CLI

To prepare slot 1 to receive a 2-port ATM module, enter the following commands at the global CONFIG level:

```
BigIron(config)# module 1 bi-atm-2-port-155m-module
BigIron(config)# write memory
```

Syntax: module <slot-num> <module-type>

The <slot-num> parameter specifies the chassis slot:

- Slots on a four-slot chassis are numbered 1 – 4, from top to bottom.
- Slots on an eight-slot chassis are numbered 1 – 8, from left to right.
- Slots on an fifteen-slot chassis are numbered 1 – 15, from left to right.

In the current software release, the <module-type> for an ATM module can be one of the following:

- bi-atm-2-port-155m-module
- bi-atm-4-port-155m-module

Upgrading ATM Software from a TFTP Server

The ATM modules contain their own flash memory from which they can boot. To upgrade the flash code (system software) on an ATM module, copy the upgrade onto a TFTP server to which the Layer 3 Switch has access, then download the code from the TFTP server to the ATM modules in the chassis.

By default, the flash code on all the ATM modules in the chassis is upgraded. If you want to upgrade only a particular module, you can specify the module's slot number.

To upgrade the ATM software, use the following CLI methods.

Upgrading the Boot Code

To upgrade the ATM boot code from a TFTP server, enter a command such as the following:

```
BigIron# atm copy tftp flash 109.157.22.26 P2B06000.bin boot
```

This command upgrades the boot code on all ATM modules in the chassis.

Syntax: atm copy tftp flash <ip-addr> <image-file-name> boot

Upgrading the Flash Code

To upgrade the software in the primary flash on all ATM modules in a BigIron Chassis device, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# atm copy tftp flash 109.157.22.26 A2R07205.bin primary
```

Syntax: atm copy tftp flash <tftp-server-ip-addr> <atm-image-file-name> primary | secondary [<slot>]

The **primary** and **secondary** parameters identify either the primary or secondary flash on the module. For each command, the parameter specifies the destination of the copy operation.

The **slot** parameter specifies a chassis slot. This parameter is optional. If you specify a slot number, the upgrade affects only the module in the slot you specify. If you do not specify a slot, the upgrade affects all the ATM modules in the chassis.

- Slots in a four-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an eight-slot chassis are numbered 1 – 8, from left to right.
- Slots in a fifteen-slot chassis are numbered 1 – 15, from left to right.

To upgrade the software on the secondary flash on the ATM module in chassis slot 6 only, enter the following command:

```
BigIron# atm copy tftp flash 109.157.22.26 A2R07205.bin secondary 6
```

To verify that the download is successful, display the contents of the ATM flash modules by entering the following command at any CLI prompt:

```
BigIron# show flash
```

The versions of flash code contained in the redundant management module flash and the ATM module flash are listed.

Configuring ATM Boot Parameters

The ATM module has its own system software and boots after the management module boots. By default, an ATM module boots from the software image in its own primary flash. You can configure an ATM module to boot from one of the following sources:

- ATM module's primary flash
- ATM module's secondary flash

To boot an ATM module from a TFTP server, you must use the interactive boot mode, then enter the **atm boot tftp...** command after the module comes up.

Changing the Boot Source

To change the boot source for the ATM module, use either of the following methods.

USING THE CLI

To change the boot source from the ATM module's primary flash to its secondary flash, enter the following commands:

```
BigIron(config)# atm boot secondary
BigIron(config)# write memory
```

Syntax: atm boot interactive | primary | secondary

The **primary** and **secondary** parameters identify either the primary or secondary flash on the ATM module.

The **interactive** parameter enables you to enter a separate command after the module comes up to boot the module from a TFTP server. If you use this method, you also need to use the **atm boot tftp...** command to boot the module after the module comes up. See the following section.

Booting the Module from TFTP

To boot the ATM module from a TFTP server, you must use the interactive boot method, then use the following method to load the software after the module comes up.

To boot the ATM module from a TFTP server, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# atm boot tftp 209.157.22.26 A2R07205.bin
```

Syntax: atm boot tftp <tftp-server-ip-addr> <atm-image-file-name>

The <tftp-server-ip-addr> parameter specifies the IP address of the TFTP server.

The <atm-image-file-name> parameter lists the name of the image file you want the module to boot from the TFTP server.

Reloading an Individual ATM Module

You can reload (boot) the software on an individual ATM module, without also reloading the management module. To reload an ATM module, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# reload atm 2
```

This command reloads the ATM module in slot 2. Messages are displayed in the CLI to show the status of the reload. The management module is not also reloaded and thus continues to operate while the ATM module is being reloaded.

Syntax: reload atm <slotnum>

Copying an ATM Image File from a Flash Card to an ATM Module's Flash Memory

To copy an ATM image file from a flash card to an ATM module's flash memory, use the following method.

USING THE CLI

To copy an ATM image file from a flash card onto all the ATM modules in the chassis, enter a command such as the following:

```
BigIron# atm copy slot1 flash A2R07205.bin primary
```

Syntax: atm copy slot1 | slot2 flash <atm-image-file-name> primary | secondary [slot]

The command in this example copies an ATM image file named A2R07205.bin from the flash card in slot 1 to all the ATM modules in the chassis.

To copy an ATM image file from a flash card onto a specific ATM module, enter a command such as the following:

```
BigIron# atm copy slot1 A2R07205.bin flash primary 4
```

The command in this example copies the specified image file onto the ATM module in chassis slot 4 only, but does not copy the file to other ATM modules in the chassis.

The following command copies an ATM image file from a TFTP server to flash memory.

Syntax: atm copy tftp flash <ip-addr> <atm-image-file-name> primary | secondary [slot]

Configuring ATM Interfaces

You can configure an ATM interface at the physical port level or the sub-interface level. Table 8.3 lists the parameters you can configure at each level.

Table 8.3: ATM Port and Sub-Interface Parameters

Parameter	Description	Default	See page...
Physical port parameters			8-10
Maximum number of VCs per VP	The maximum number of VCs each VP on the port can have.	4 VPs 1024 VCs on each VP	8-10
Port state	Whether the port is enabled or disabled.	Enabled	8-11
Loopback path	The extent to which the Foundry device sends ATM cells when testing the circuitry of a port. You can specify one of the following loopback paths: <ul style="list-style-type: none"> The port circuitry alone The port circuitry and the circuitry of the port at the other end of the link Loopback is disabled by default and must be off to use the port for normal traffic.	Loopback testing is disabled by default. Loopback is enabled when you specify a loopback path.	8-12
Clock source	The source the port uses for timing functions. An ATM port can use one of the following clock sources: <ul style="list-style-type: none"> Internal – The port's uses the clock located within it's its own circuitry. Line – The port uses an external source located on another device somewhere on the network. 	Internal	8-12

Table 8.3: ATM Port and Sub-Interface Parameters (Continued)

Parameter	Description	Default	See page...
SONET scramble mode	The port can scramble long sequences of zeros or ones at the SONET level to provide security and clock recovery. The SONET scramble mode can be set to one of the following states: <ul style="list-style-type: none"> Disabled Enabled 	Enabled	8-12
ATM scramble mode	The port can scramble long sequences of zeros or ones at the ATM level. The ATM scramble mode can be set to one of the following states: <ul style="list-style-type: none"> Disabled Enabled 	Enabled	8-13
Cyclic Redundancy Check (CRC) for ATM header checksums	The port can add a CRC to the header checksum for each ATM cell. The CRC mode can be set to one of the following states: <ul style="list-style-type: none"> Disabled Enabled 	Enabled	8-13
Sub-interface parameters			8-15
Link type	Whether the port is for a point-to-point interface or a point-to-multipoint interface. Note: On a point-to-multipoint interface, you can use Inverse ARP for address resolution.	Point-to-point	8-15
Permanent Virtual Circuit (PVC) Class of Service (CoS)	A PVC is a user-configured dedicated circuit used by a specific IP interface over the ATM port. A PVC consists of the following: <ul style="list-style-type: none"> Virtual Path Identifier (VPI) – A number that identifies a virtual path. Virtual Channel Identifier (VCI) – A number that identifies a virtual channel within the virtual path. The CoS is a policy the Foundry device uses for sending outbound traffic on the interface. You can configure the following types of CoS for ATM: <ul style="list-style-type: none"> Constant Bit Rate (CBR) – The device always sends outbound traffic on the PVC at the same rate. Variable Bit Rate (VBR) – The device accommodates bursty traffic by temporarily allowing the PVC to send traffic at a higher rate than the rate specified for the PVC's normal operation. Unspecified Bit Rate (UBR) – The device uses bandwidth as it becomes available. Note: Foundry does not support ABR.	No PVCs configured	8-18

Table 8.3: ATM Port and Sub-Interface Parameters (Continued)

Parameter	Description	Default	See page...
MTU	The MTU is the maximum packet size supported on a PVC. You can configure the MTU to a value from 1500 – 9180 bytes.	4470 bytes	8-18
IP address	An IP subnet interface. You can configure one IP address on a PVC.	None configured	8-19
IP ACL	A standard or extended IP ACL.	None configured	8-20
Route parameters			8-20
static IP route		None configured	8-20

Configuring Port Parameters

Use the following procedures to configure the ATM port parameters listed in Table 8.3 on page 8-8.

NOTE: The defaults for the ATM port parameters are appropriate for most configurations. Unless you are sure you need to change a port parameter, go to “Configuring Sub-Interface Parameters”.

Changing the Number of Virtual Channels (VCs) per Virtual Path (VP)

By default, an ATM port can have up to four VPs and each VP can have up to 1024 VCs. In previous releases, the maximum number of VCs allowed in a VP is not configurable. Software release 07.5.00 enables you to change, on an individual port basis, the maximum number of VCs allowed in a VP to one of the following values:

- 4096
- 2048
- 1024 (the default)
- 256

When you change the maximum number of VCs a VP can have, the maximum number of VPs the port can have also is changed.

Maximum Number of VCs per VP	Maximum Number of VPs
4096	1
2048	2
1024	4
256	16

NOTE: The total number of VCs an ATM port can have is still 4096.

NOTE: When you change the maximum number of VCs per VP, the software deletes all configured PVCs from the port.

To change the maximum number of VCs per VP, enter commands such as the following:

```
BigIron(config)# interface atm 3/1
BigIron(config-atmif-3/1)# atm vc-per-vp 256
Changing the number of VCs per VP will remove all current PVCs configured on this
atm port.
Do you want to continue?(enter 'y' or 'n'):
```

These commands change the CLI to the configuration level for ATM port 3/1, and change the maximum number of VCs per VP to 256. The software asks you to verify that you want to make the change, since all PVCs already configured on the port will be removed.

Syntax: [no] atm vc-per-vp 256 | 1024 | 2048 | 4096

NOTE: In earlier versions of the software, supported configurable number of VCs per VP is; 16, 1024, 2048, and 4096.

To display how many VCs a VP can have on a port, enter a command such as the following:

```
BigIron(config)# show interface atm 3/1
No port name
Hardware is ATM
Encapsulation llcsnap, clock is internal
Framing is SONET, BW 155000Kbit
Loopback not set, keepalive not set, scramble enabled
Each virtual path contains 4096 virtual channels
300 second input rate: 0 bits/sec, 0 packets/sec
300 second output rate: 0 bits/sec, 0 packets/sec
0 packets input, 0 bytes, 0 no buffer
Received 0 CRCs, 0 shorts, 0 giants, 0 alignments
0 packets output, 0 bytes, 0 underruns
```

The line shown in bold type indicates how many VCs each VP on the port can have.

Changing the Interface State

The ATM interfaces are enabled by default. To disable or re-enable an interface, use the following method.

USING THE CLI

To disable ATM interface 4/1, enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# disable
```

Syntax: [no] disable

To re-enable ATM interface 4/1, enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# enable
```

Syntax: [no] enable

Changing the Loopback Path

Foundry ATM interfaces can use the following loopback configurations for self tests:

- Disabled – The port is not in loopback mode. Loopback must be disabled to use the port for normal traffic. This is the default.
- Internal – Packets that the router transmits on the interface are looped back to the interface's ATM framer. The internal loopback configuration is useful for checking the ATM circuitry.
- Line – The interface's transmit and receive fibers are logically linked so that packets received on the receive fiber are sent back out on the transmit fiber. Use this mode on the ATM interfaces on both ends of a link to test the interfaces along with the link.

NOTE: Loopback must be disabled if you want to use the port for normal traffic. Do not specify a loopback source (which automatically enables loopback) unless you are testing the interface.

USING THE CLI

To configure ATM interface 4/1 for internal loopback, enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# loop internal
```

Syntax: [no] loop internal | line

The **internal** and **line** parameters specify the path for the loopback. The **internal** parameter loops packets transmitted on the interface back to the framer. The **line** parameter loops packets that are received on the receive fiber of the port back out on the transmit fiber.

To disable loopback again so you can use the port for normal traffic, enter a command such as the following:

```
BigIron(config-atmif-4/1)# no loop internal
```

Changing the Clock Source

By default, Foundry ATM interfaces use clock information from the ATM port itself as the clocking source. You can change the clock source to "line", in which case the port receives its clock information from the network.

If you are connecting two Foundry ATM modules back-to-back or the ATM interfaces are connected by a fiber link that has no clocking information on it, use the internal clock source. Otherwise, use the network (line) as the source.

To change the clock source, use the following method.

USING THE CLI

To change the clock source for ATM interface 4/1 to line (the network), enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# clock line
```

Syntax: [no] clock internal | line

The **internal** and **line** parameters specify whether the clock source is on the ATM module (internal) or on the network (line). The default is **internal**.

Disabling or Re-Enabling SONET Scramble Mode

By default, Foundry devices perform bit scrambling at the SONET level. Scrambling helps ensure that long strings of zeros or ones are converted into more random-appearing bit sequences. This is useful for ensuring clock recovery and security. Most ATM links have this feature enabled, so the default is appropriate for most ATM networks. The setting of the SONET scramble mode affects both send and receive traffic on the port.

To disable or re-enable the SONET scramble mode, use the following method.

USING THE CLI

To disable the SONET scramble mode for ATM interface 4/1, enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# no sonet-scram
```

Syntax: [no] sonet-scram

To re-enable the feature, enter the following command:

```
BigIron(config-atmif-4/1)# sonet-scram
```

Disabling or Re-Enabling ATM Scramble Mode

By default, Foundry ATM ports perform bit scrambling at the ATM cell level. Most ATM scramble mode affects both send and receive traffic on the port.

To disable or re-enable the ATM scramble mode, use the following method.

USING THE CLI

To disable the ATM scramble mode for ATM interface 4/1, enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# no atm-scram
```

Syntax: [no] atm-scram

To re-enable the feature, enter the following command:

```
BigIron(config-atmif-4/1)# atm-scram
```

Disabling or Re-Enabling the Cyclic Redundancy Check for ATM Header Checksums

By default, Foundry ATM ports add a CRC for the header checksum in an ATM cell. Some ATM switches add the CRC by default while others do not. Check the documentation for your ATM switch to determine the setting for your switch. The setting of the CRC mode affects both send and receive traffic on the port.

To disable or re-enable the CRC for header checksums, use the following method.

USING THE CLI

To disable the CRC for ATM interface 4/1, enter the following commands:

```
BigIron(config)# interface atm 4/1
BigIron(config-atmif-4/1)# no atm-hcsadd
```

Syntax: [no] atm-hcsadd

To re-enable the feature, enter the following command:

```
BigIron(config-atmif-4/1)# atm-hcsadd
```

Using the Web Management Interface to Configure an ATM Port

To configure an ATM port or view its current settings, select Configure->Port->ATM from the tree view of options. The following panel is displayed.

ATM Port Configuration

Port	Speed	Encapsulation	MTU	Clock	Loop Back	Scramble-ATM	Framing	CRC	Keep Alive	
15/1	155000	SNAP	4470	Internal	None	Enable	SONET	32	0	Modify
15/2	155000	SNAP	4470	Internal	None	Enable	SONET	32	0	Modify
15/3	155000	SNAP	4470	Internal	None	Enable	SONET	32	0	Modify
15/4:123	155000	SNAP	4470	Internal	None	Enable	SONET	32	0	Modify
15/1.10	0	SNAP	4470	Internal	none	Disable	SONET	32	0	Modify
15/3.1023	0	SNAP	4270	Internal	none	Disable	SONET	32	0	Modify
Port	Speed	Encapsulation	MTU	Clock	Loop Back	Scramble-ATM	Framing	CRC	Keep Alive	

[\[Home\]](#)[\[Site Map\]](#)[\[Logout\]](#)[\[Save\]](#)[\[Frame Enable\]](#)[\[Disable\]](#)[\[TELNET\]](#)

To modify settings for a port, click Modify next to the row of information for the port. The following panel is displayed.

PVC Port

Port:15/1.3	
Name:	<input type="text"/>
Status:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Speed:	<input checked="" type="radio"/> 155000 <input type="radio"/> 622000 <input type="radio"/> 2488000
Encapsulation:	<input checked="" type="radio"/> SNAP
Mtu:	<input type="text" value="4470"/>
Clock:	<input checked="" type="radio"/> Internal <input type="radio"/> Line
Loop Back:	<input type="radio"/> Line <input type="radio"/> Internal <input checked="" type="radio"/> None
CRC:	<input checked="" type="radio"/> 32 <input type="radio"/> 16
Framing:	<input checked="" type="radio"/> SONET <input type="radio"/> SDH
Scramble PVC:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Keep Alive:	<input type="text" value="0"/>

[\[Show\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

For information about the parameters, see the equivalent CLI sections above.

Configuring Sub-Interface Parameters

The following sections describe how to configure a point-multipoint interface, PVC and CoS parameters, and IP addresses. You must configure PVC, CoS, and IP address parameters to use an ATM interface to send and receive traffic. The default link type is point-to-point.

Specifying the Link Type

When you configure an ATM sub-interface, you can create a point-to-point sub-interface or a point-to-multipoint sub-interface.

- Point-to-point – The sub-interface contains one PVC, and thus represents a point-to-point link between one local IP interface (associated with the sub-interface) and one remote IP interface. This is the only type of ATM sub-interface supported in previous releases.

NOTE: If you plan to use a point-to-point link, you can skip the rest of this section. The CLI creates a point-to-point link by default when you enter the **interface atm** <slot>/<port>.<subif> command without also specifying a link type.

- Point-to-multipoint – The sub-interface contains multiple PVCs. Each PVC uses the same local IP interface (the one associated with the sub-interface) but has a unique remote IP interface. Support for this type of ATM sub-interface is new in software release 07.5.00.

When you configure a PVC on a point-to-multipoint sub-interface, you must specify the source of the remote IP address of the PVC. You can specify the address explicitly or enable the Layer 3 Switch to dynamically obtain the address using Inverse ARP.

Inverse ARP

When a point-to-multipoint sub-interface comes up, the software sends an Inverse ARP request on each PVC enabled for Inverse ARP. The software maps the IP address received in an Inverse ARP reply to the PVC on which the reply was received.

The software regularly updates the remote IP addresses mapped to PVCs by sending a new Inverse ARP at regular intervals. By default, the software sends an Inverse ARP request every 10 minutes. If the device does not receive a response, the software retries once a minute indefinitely until a reply is received. You can change the Inverse ARP interval to a value from 1 – 60 minutes.

The Foundry implementation of Inverse ARP is based on the following RFCs:

- RFC 1293 – Inverse Address Resolution Protocol
- RFC 2225 – Classical IP and ARP over ATM

Configuring a Point-To-Multipoint Sub-Interface

To configure a point-to-multipoint ATM interface, enter commands such as the following:

```
BigIron(config)# interface atm 4/1.1 multipoint
BigIron(config-subif-4/1.1)# atm pvc 1 1 cbr 10000 ip inarp
BigIron(config-subif-4/1.1)# atm pvc 1 2 cbr 20000 ip inarp
BigIron(config-subif-4/1.1)# atm pvc 1 3 cbr 25000 ip inarp
BigIron(config-subif-4/1.1)# ip address 10.10.10.4
```

The first command configures a point-to-multipoint ATM interface. The remaining commands configure PVCs on the interface. In this example, all three PVCs use Inverse ARP to resolve the remote IP addresses.

Syntax: [no] interface atm <slot>/<port>.<subif> [multipoint | point-to-point]

Syntax: [no] atm pvc <vpi>< vci> cbr |ubr | vbr ip <remote-ip-addr> | inarp [<mins>]

The **multipoint** | **point-to-point** parameter specifies the type of ATM interface.

- **multipoint** – The sub-interface can have more than one PVC.
- **point-to-point** – The sub-interface can have only one PVC. This is the default.

NOTE: Once you create the sub-interface, you cannot change the interface type.

The **ip <remote-ip-addr> | inarp <mins>** parameter specifies how the PVC's remote IP address is obtained.

- <remote-ip-addr> – The address is obtained statically, when you specify the address using this parameter.
- **inarp** [<mins>] – The address is obtained dynamically using Inverse ARP. The <mins> parameter specifies how often the software sends a new Inverse ARP to refresh the remote IP address mapped to the PVC. You can specify from 1 – 60 minutes. The default is 10 minutes.

You must specify an IP address or **inarp**. There is no default.

The other parameters are supported in previous releases as well as the current release. See the “Using Asynchronous Transfer Mode Modules” chapter of *Foundry Switch and Router Installation and Basic Configuration Guide*.

To configure a point-to-multipoint interface with PVCs that use remote IP address that you specify, enter commands such as the following:

```
BigIron(config)# interface atm 4/1.1 multipoint
BigIron(config-subif-4/1.1)# atm pvc 1 1 ubr ip 10.10.10.1
BigIron(config-subif-4/1.1)# atm pvc 1 2 ubr ip 10.10.10.2
BigIron(config-subif-4/1.1)# atm pvc 1 3 ubr ip 10.10.10.3
BigIron(config-subif-4/1.1)# ip address 10.10.10.4
```

Displaying the Point-To-Multipoint Mappings

To display the point-to-multipoint mappings, enter the following command:

```
BigIron(config)# show atm map
      IP Address      VCI - VPI      Type      Age      Port
1     10.1.1.1       0 - 202       Static    None     4/2
2     10.1.1.2       0 - 204       Static    None     4/2
3     10.1.1.3       0 - 205       Static    None     4/2
```

Syntax: show atm map

This command shows the following information.

Table 8.4: CLI Display of ATM Point-To-Multipoint Mappings

This Field...	Displays...
Entry number	The number of this entry (row) in the mapping table.
IP Address	The remote IP address mapped to this PVC.
VCI - VPI	The VCI and VPI of the PVC.
Type	The type of mapping, which can be one of the following: <ul style="list-style-type: none"> Dynamic – The mapping was created using an address received through Inverse ARP. Static – The mapping was created using an IP address that you specified when you configured the PVC.
Age	The number of minutes since the last time a dynamic mapping was refreshed. For static mappings, the field value is "None".
Port	The port on which the PVC is configured.

Clearing the Dynamically Learned Point-To-Multipoint Mappings

To clear the learned point-to-multipoint mappings and refresh the interfaces that use Inverse ARP, enter the following command:

```
BigIron# clear atm map
```

Syntax: clear atm map

This command does not affect statically configured mappings.

Displaying Diagnostic Information

To display diagnostic information for ATM Inverse ARP, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# debug atm multipoint
ATM_MULTIPPOINT: INARP ATM length = 20
ATM_MULTIPPOINT: Tx INATMARP Request packet:
source ip 1.1.1.2
target ip 0.0.0.0
inatmarp_pkt->src_atm_number_t1.length 0
inatmarp_pkt->src_atm_subaddress_t1.length 0
inatmarp_pkt->target_atm_number_t1.length 0
inatmarp_pkt->target_atm_subaddress_t1.length 0
```

```
BigIron Router#ATM_MULTIPPOINT: INARP ATM length = 20
ATM_MULTIPPOINT: Rx INATMARP packet:
  source ip 1.1.1.1
  target ip 1.1.1.2
  inatmarp_pkt->src_atm_number_t1.length 0
  inatmarp_pkt->src_atm_subaddress_t1.length 0
  inatmarp_pkt->target_atm_number_t1.length 0
  inatmarp_pkt->target_atm_subaddress_t1.length 0
```

Syntax: debug atm multipoint

Note Regarding Static Route Support

You cannot configure an IP static route on an ATM point-to-multipoint interface. If you accidentally try to configure such an interface, the CLI displays an error message, as shown in the following example:

```
BigIron(config)# ip route 5.5.5.5/24 atm 2/1.1
Error - static interface routing not allowed on atm multipoint subinterfaces
```

Configuring PVC Parameters

To configure a PVC, change the CLI to the configuration level for a sub-interface on an ATM port, then specify the following:

- The VPI and VCI
- The CoS type and the values for the applicable CoS parameters

Configure the parameters to match the VC and CoS settings on the ATM switch at the other end of the link.

The following sections describe how to perform these configuration tasks.

USING THE CLI

To configure a PVC and specify CoS parameters for the PVC, first enter a command such as the following to change the CLI to the ATM sub-interface configuration level:

```
BigIron(config)# interface atm 4/1.1
BigIron(config-subif-4/1.1)#
```

The command in this example changes the CLI to the configuration level for sub-interface 1 on port 1 of the ATM module in chassis slot 4. The CLI prompt changes to indicate the configuration level.

Syntax: interface atm <slot>/<port>.<subif>

The <slot>/<port> parameter specifies the chassis slot and the port number on the ATM module in the specified slot. The <subif> parameter specifies a sub-interface on the ATM port. You can configure up to 4095 sub-interfaces on an ATM port, with a per chassis maximum of 4095 total. Specify a number from 1 – 4095 for the sub-interface.

After changing to the configuration level for a sub-interface, enter a command such as the following:

```
BigIron(config-subif-4/1.1)# atm pvc 1 200 cbr 10000
```

This command adds a PVC with virtual path 1 and virtual channel 200, and configures CoS parameters for CBR. The **cbr** parameter indicates that the CoS method is CBR. The value following **cbr** indicates the Peak Cell Rate (PCR), in this case 10,000 kilobits per second (10 Mbps).

The following command configures a PVC for VBR:

```
BigIron(config-subif-4/1.1)# atm pvc 1 300 vbr 10000 5000 200
```

This example specifies that the maximum data rate on ATM VC 1, 300 is 10000 kilobits, while the average data rate is 5000 kilobits. The VC can accommodate a maximum burst size of 200 cells.

Syntax: [no] atm pvc <vpi> <vci> cbr <pcr>

or

Syntax: [no] atm pvc <vpi> <vci> vbr <pcr> <scr> <mbs>

or

Syntax: [no] atm pvc <vpi> <vci> ubr

The **pvc** <vpi> <vci> parameter specifies the virtual path and virtual channel of the PVC. Each PVC has a unique combination of virtual path and virtual channel. You can specify a number from 0 – 3 for the VPI. You can specify a number from 0 – 1023 for the VCI. You can use a given VPI and VCI combination only once on a given ATM port. However, you can use the same combination on other ATM ports on the same module or in the same chassis.

The **cbr** <pcr> parameter specifies CBR as the CoS method and specifies the PCR. The <pcr> can be from 1 – 155000 kilobits.

The **vbr** <pcr> <scr> <mbs> parameter specifies VBR as the CoS method, and specifies the PCR, Sustained Cell Rate (SCR), and Maximum Burst Size (MBS).

- The <pcr> can be from 1 – 155000 kilobits.
- The <scr> can be from 1 – 155000 kilobits.
- The <mbs> can be from 2 – 255.

NOTE: The <scr> value must be at least one less than the <pcr> value.

The **ubr** parameter specifies UBR as the CoS method.

NOTE: In software release 07.6.01, a PVC configured for UBR makes its unused bandwidth available to other PVCs configured for UBR. In earlier releases, unused bandwidth on a PVC configured for UBR is not available to other PVCs.

NOTE: The CLI checks the values you enter for the <pcr> and <scr> parameters to make sure that the combined cell rates on all the PVCs configured on the port do not exceed the bandwidth of the port.

NOTE: ABR is not supported and no support is planned.

Changing the MTU

The MTU specifies the maximum packet size supported on the PVC. The default MTU is 4470. You can configure the MTU to a value from 1500 – 9180. ATM fragments packets above 1500 bytes of payload. Each fragment except the first and last has a payload length of 1440 bytes with an extra 20-byte IP header.

NOTE: The MTU on both sides of an ATM PVC must match.

USING THE CLI

To change the MTU on ATM sub-interface 4/1.1, enter the following commands:

```
BigIron(config)# interface atm 4/1.1
BigIron(config-atmif-4/1.1)# mtu 9180
```

Syntax: [no] mtu <num>

The <num> parameter specifies the MTU and can be a value from 1500 – 9180. The default is 4470.

Configuring an IP Address

You can add one IP address to a PVC.

USING THE CLI

To configure an IP address on an ATM interface, enter a command such as the following:

```
BigIron(config-atmif-4/1.1)# ip address 192.168.2.9 255.255.255.0
```

Syntax: [no] ip address <ip-addr> <ip-mask>

or

Syntax: [no] ip address <ip-addr>/<mask-bits>

The <ip-addr> parameter specifies the subnet address.

Foundry devices support both classical IP network masks (Class A, B, and C subnet masks, and so on) and Classless Interdomain Routing (CIDR) network prefix masks. The <ip-mask> parameter specifies the subnet mask. Alternatively, you can use the /<mask-bits> parameter to indicate the number of significant bits in the network mask. Here is an example: **ip address 192.168.2.9/24**

Applying an IP ACL

You can apply IP ACLs to a sub-interface. The ACLs filter traffic sent or received in the sub-interface.

USING THE CLI

To configure an IP ACL and apply it to an ATM sub-interface, enter commands such as the following:

```
BigIron(config)# access-list 1 deny host 209.157.22.26 log
BigIron(config)# access-list 1 permit any
BigIron(config)# int atm 4/1.1
BigIron(config-atmif-4/1.1)# ip access-group 1 out
```

Syntax: [no] ip access-group <num> in | out

The **in** | **out** parameter specifies whether the ACL applies to incoming traffic or outgoing traffic on the sub-interface.

Configuring Route Parameters

The Layer 3 Switch can route IP traffic between ATM and Ethernet or POS based on the IP addresses configured on the interfaces.

You also can configure a static IP route that uses an ATM interface's IP address or sub-interface as its next-hop path.

Configuring a Static IP Route that Uses ATM as its Next-Hop Path

You can configure a static IP route on an ATM port's IP address or an individual sub-interface. When you configure a static IP route, you must specify the following parameters:

- The IP address and network mask for the route's destination network.
- The route's path, which can be one of the following:
 - The IP address of a next-hop gateway
 - A port. For ATM, the path can be a specific IP address on the port, or even a specific sub-interface on a PVC.
 - A virtual interface (a routing interface used by VLANs for routing Layer 3 protocol traffic among one another)
 - A "null" interface. The Layer 3 Switch drops traffic forwarded to the null interface.

You also can specify the following optional parameters:

- The route's metric – The value the Layer 3 Switch uses when comparing this route to other routes in the IP route table to the same destination. The metric applies only to routes that the Layer 3 Switch has already placed in the IP route table. The default metric for static IP routes is 1.
- The route's administrative distance – The value that the Layer 3 Switch uses to compare this route with routes from other route sources to the same destination before placing a route in the IP route table. This parameter does not apply to routes that are already in the IP route table. The default administrative distance for static IP routes is 1.

USING THE CLI

To configure a static IP route that uses an ATM port as its next-hop path, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# ip route 5.5.5.0/24 6.6.6.1
```

This command configures a static route to network 5.5.5.0/24, using IP interface 6.6.6.1 as the next-hop path to the network. In this case, the next-hop IP address 6.6.6.1 is configured on an ATM sub-interface. Therefore, the PVC configured on that sub-interface is used as the route's next-hop path.

Here is how the route appears in the IP route table:

```
BigIron(config)# show ip route 5.5.5.0
Destination      NetMask      Gateway      Port      Cost      Type
5.5.5.0          255.255.255.0  6.6.6.1      3/1.1     1         S
```

The Port field lists the ATM port and sub-interface on which the next-hop address (6.6.6.1) is configured. The route uses the PVC configured on that sub-interface.

Syntax: ip route <dest-ip-addr> <dest-mask>
<next-hop-ip-addr> | atm <slotnum>/<portnum>.<subif>
[<metric>] [distance <num>]

or

Syntax: ip route <dest-ip-addr>/<mask-bits>
<next-hop-ip-addr> | atm <slotnum>/<portnum>.<subif>
[<metric>] [distance <num>]

The <dest-ip-addr> is the route's destination. The <dest-mask> is the network mask for the route's destination IP address. Alternatively, you can specify the network mask information by entering a forward slash followed by the number of bits in the network mask. For example, you can enter 192.0.0.0 255.255.255.0 as 192.0.0.0/24. You can enter multiple static routes for the same destination for load balancing or redundancy.

The <next-hop-ip-addr> is the IP address of the next-hop router (gateway) for the route.

If you do not want to specify a next-hop IP address, you can instead specify a port or interface number on the Layer 3 Switch. If you specify an ATM port and sub-interface, the <portnum> is the port's number (including the slot number). In this case, the Layer 3 Switch forwards packets destined for the static route's destination network to the specified interface. Conceptually, this feature makes the destination network like a directly connected network, associated with a specific Layer 3 Switch interface.

NOTE: The port or virtual interface you use for the static route must have at least one IP address configured on it. The address does not need to be in the same subnet as the destination network.

The <metric> parameter can be a number from 1 – 16. The default is 1.

NOTE: If you specify 16, RIP considers the metric to be infinite and thus also considers the route to be unreachable.

The **distance** <num> parameter specifies the administrative distance of the route. When comparing otherwise equal routes to a destination, the Layer 3 Switch prefers lower administrative distances over higher ones, so make sure you use a low value for your default route. The default is 1.

You can also assign the default router as the destination by entering 0

To specify the ATM port and sub-interface instead of the IP address, enter a command such as the following:

```
BigIron(config)# ip route 5.5.5.0/24 atm 3/1.1
```

Here is how this route appears in the IP route table:

```
BigIron(config)# show ip route 5.5.5.0
Destination      NetMask          Gateway          Port   Cost   Type
5.5.5.0          255.255.255.0   0.0.0.0         3/1.1   1     S
```

Notice that the gateway address for the first route is 6.6.6.1, whereas the gateway address for the second route is 0.0.0.0. Since the second route's next hop is configured to be a sub-interface instead of an IP address, the route is not associated with a specific next-hop IP address but instead uses the address configured on the sub-interface.

Configuring OSPF on an ATM Interface

You can configure an OSPF interface on an ATM point-to-point interface or a point-to-multipoint interface.

Point-To-Multipoint

In OSPF, a point-to-multipoint interface means point-to-point neighbor adjacencies are established between the local router and the remote routers as long as the underlying PVCs are configured and the corresponding IP addresses are defined. A fully-meshed ATM cloud is not required. However, all devices that want to establish an adjacency with the local router must have point-to-multipoint interfaces.

To configure an OSPF interface on an ATM point-to-multipoint interface, enter commands such as the following:

```
BigIron(config)# interface atm 1/2.10 multipoint
BigIron(config-subif-1/2.10)# atm pvc 1 1 ubr ip 10.10.10.1
BigIron(config-subif-1/2.10)# atm pvc 1 2 ubr ip 10.10.10.2
BigIron(config-subif-1/2.10)# atm pvc 1 3 ubr ip 10.10.10.3
BigIron(config-subif-1/2.10)# ip address 10.10.10.9/28
BigIron(config-subif-1/2.10)# ip ospf area 0
BigIron(config-subif-1/2.10)# disable
BigIron(config-subif-1/2.10)# ip ospf network point-to-multipoint
BigIron(config-subif-1/2.10)# enable
```

The first four commands configure a point-to-multipoint ATM subinterface with three PVCs. The last five commands configure an IP address on the subinterface, disable the interface, enable OSPF on the subinterface, configure the OSPF network type for the interface to point-to-multipoint, then re-enable the interface. You must disable the interface to change the interface type.

Syntax: [no] ip ospf network [point-to-multipoint]

The default OSPF network type is point-to-point.

NOTE: You must enter the **ip ospf network point-to-multipoint** command before the interface forms OSPF adjacencies. The command is not valid once an adjacency is formed.

Point-To-Point

To configure an OSPF interface on an ATM point-to-point interface, enter commands such as the following:

```
BigIron(config)# interface atm 1/2.10
BigIron(config-subif-1/2.10)# atm pvc 3 34 ubr
BigIron(config-subif-1/2.10)# ip address 10.10.10.9/28
BigIron(config-subif-1/2.10)# ip ospf area 0
BigIron(config-subif-1/2.10)# disable
BigIron(config-subif-1/2.10)# ip ospf network
BigIron(config-subif-1/2.10)# enable
```

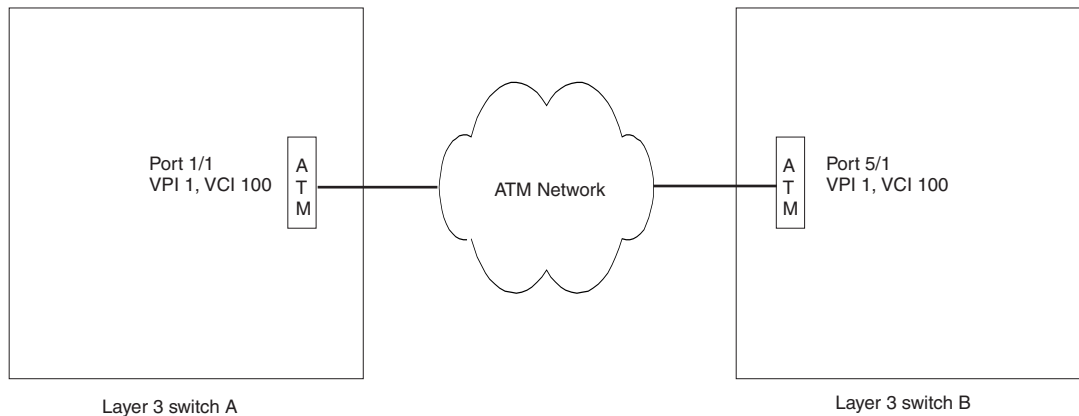
Verifying PVC Configuration

After you configure a Permanent Virtual Circuit (PVC) in your ATM network, you can test end-to-end PVC connectivity by pinging the port number, VP ID, and VC ID of the other end of a link. The ATM ping facility is based on the ATM OAM loopback standard.

Figure 8.4 shows an example of two Foundry devices configured with ATM interfaces. Each of the devices is configured for two PVCs.

Figure 8.4 Example of a VC loopback test

1. Ping is sent from Port 1/1, VPI 1, VCI 100.
2. The port at the other end of the link for VPI 3, VCI 300 receives the ping and responds.
3. The port and VC (VPI and VCI) that sent the ping receive the reply and display the result.



To test Layer 3 Switch A's configuration for VPI 1, VCI 100, send an ATM ping cell from port 1/1, VPI 1, VCI 100. If the PVC is properly configured in the network, then the VC at the other end of the link sends a reply to the ping. The VC that sent the ping receives the reply and displays the ping results.

Notice that the PVC crosses an ATM cloud, which can contain numerous ATM devices. The ping tests the entire circuit.

To verify the successful configuration of a PVC, use the following CLI method.

USING THE CLI

To verify configuration of a PVC, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# ping atmvc 1/1 1 100
ATM: OAM: sending loopback request on ATM port 1/1, vpi 1, vci 100
BigIron#
ATM: OAM :loopback response received on 1/1, 1,100
```

This example shows the output that results from a successful ping. The command sends the ping from ATM port 1/1 on VPI 1, VCI 100. The device at the other end of the link is configured for the same VPI and VCI and responds to the ping.

Syntax: ping atmvc <portnum> <vpi> <vci>

The <portnum> parameter specifies the port on which the VC is configured on the local device (the device on which you are entering the command).

The <vpi> and <vci> parameters specify the virtual path and virtual channel, which together make the VC.

Mapping PVCs to VLANs

Foundry devices support mapping ATM Permanent Virtual Circuits (PVCs) to VLANs, based on the ATM bridging standards of RFC 2684. This feature allows an ATM PVC to be configured as a bridging interface and used in conjunction with a VLAN.

Figure 8.5 illustrates the how the Foundry device adds a VLAN ID and tag to packets subject to PVC-VLAN mapping. In this example, PVC = 3 is mapped to VLAN 20. The Foundry device adds VLAN ID 20 and tag 8100 to packets from PVC = 3.

Figure 8.5 Adding a VLAN ID and tag to packets from a PVC


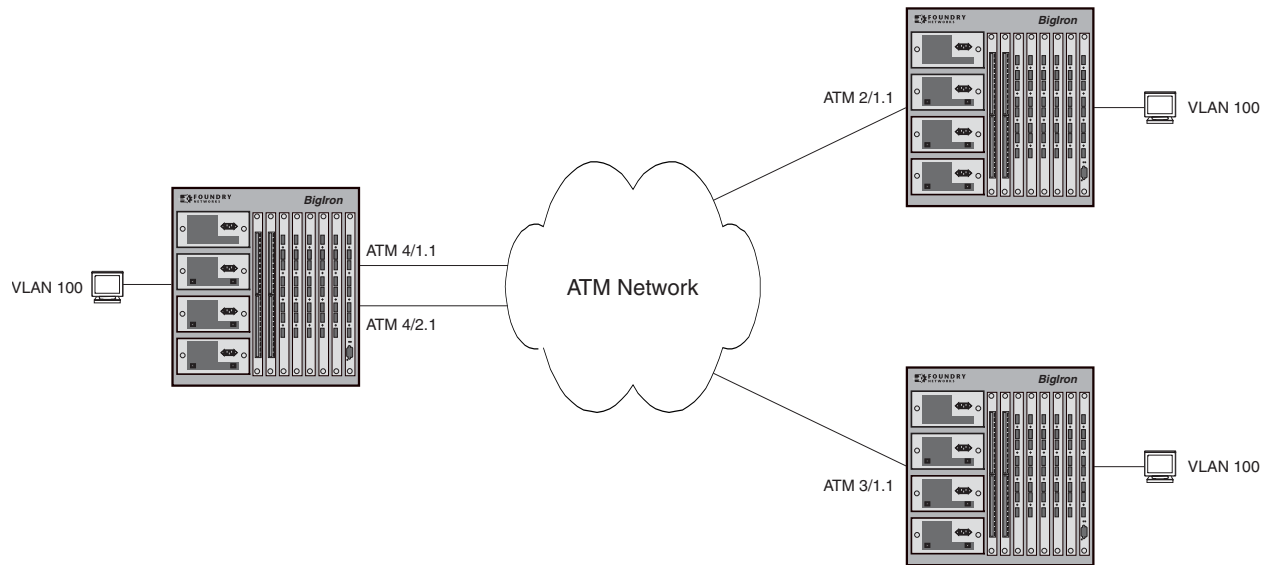


Figure 8.6 shows an example of a VLAN with two PVCs mapped to it.

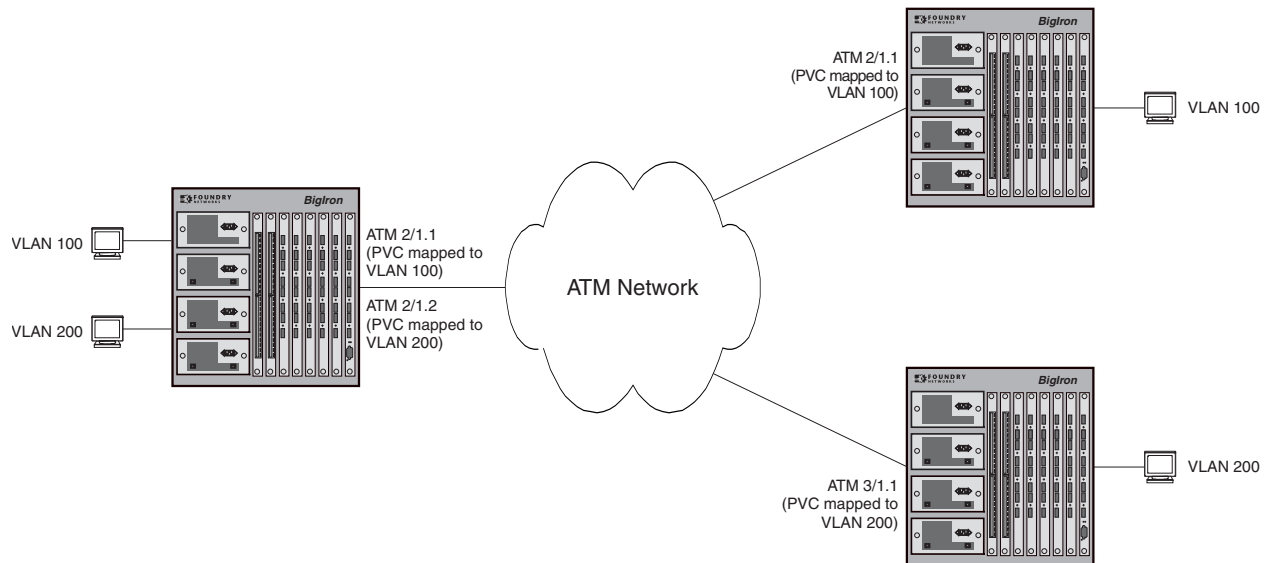
Figure 8.6 PVC-VLAN mapping – hub-and-spoke configuration with two PVCs mapped to one VLAN



In this example, PVCs on ATM sub-interfaces 4/1.1 and 4/2.1 have been mapped to VLAN 100. These PVCs serve as bridging interfaces for traffic on VLAN 100.

Traffic for more than one VLAN can be forwarded over PVCs on multiple sub-interfaces on the same ATM port. Figure 8.7 shows an example of this kind of configuration. In this configuration, a PVC on sub-interface 2/1.1 is mapped to VLAN 100, and a PVC on sub-interface 2/1.2 is mapped to VLAN 200.

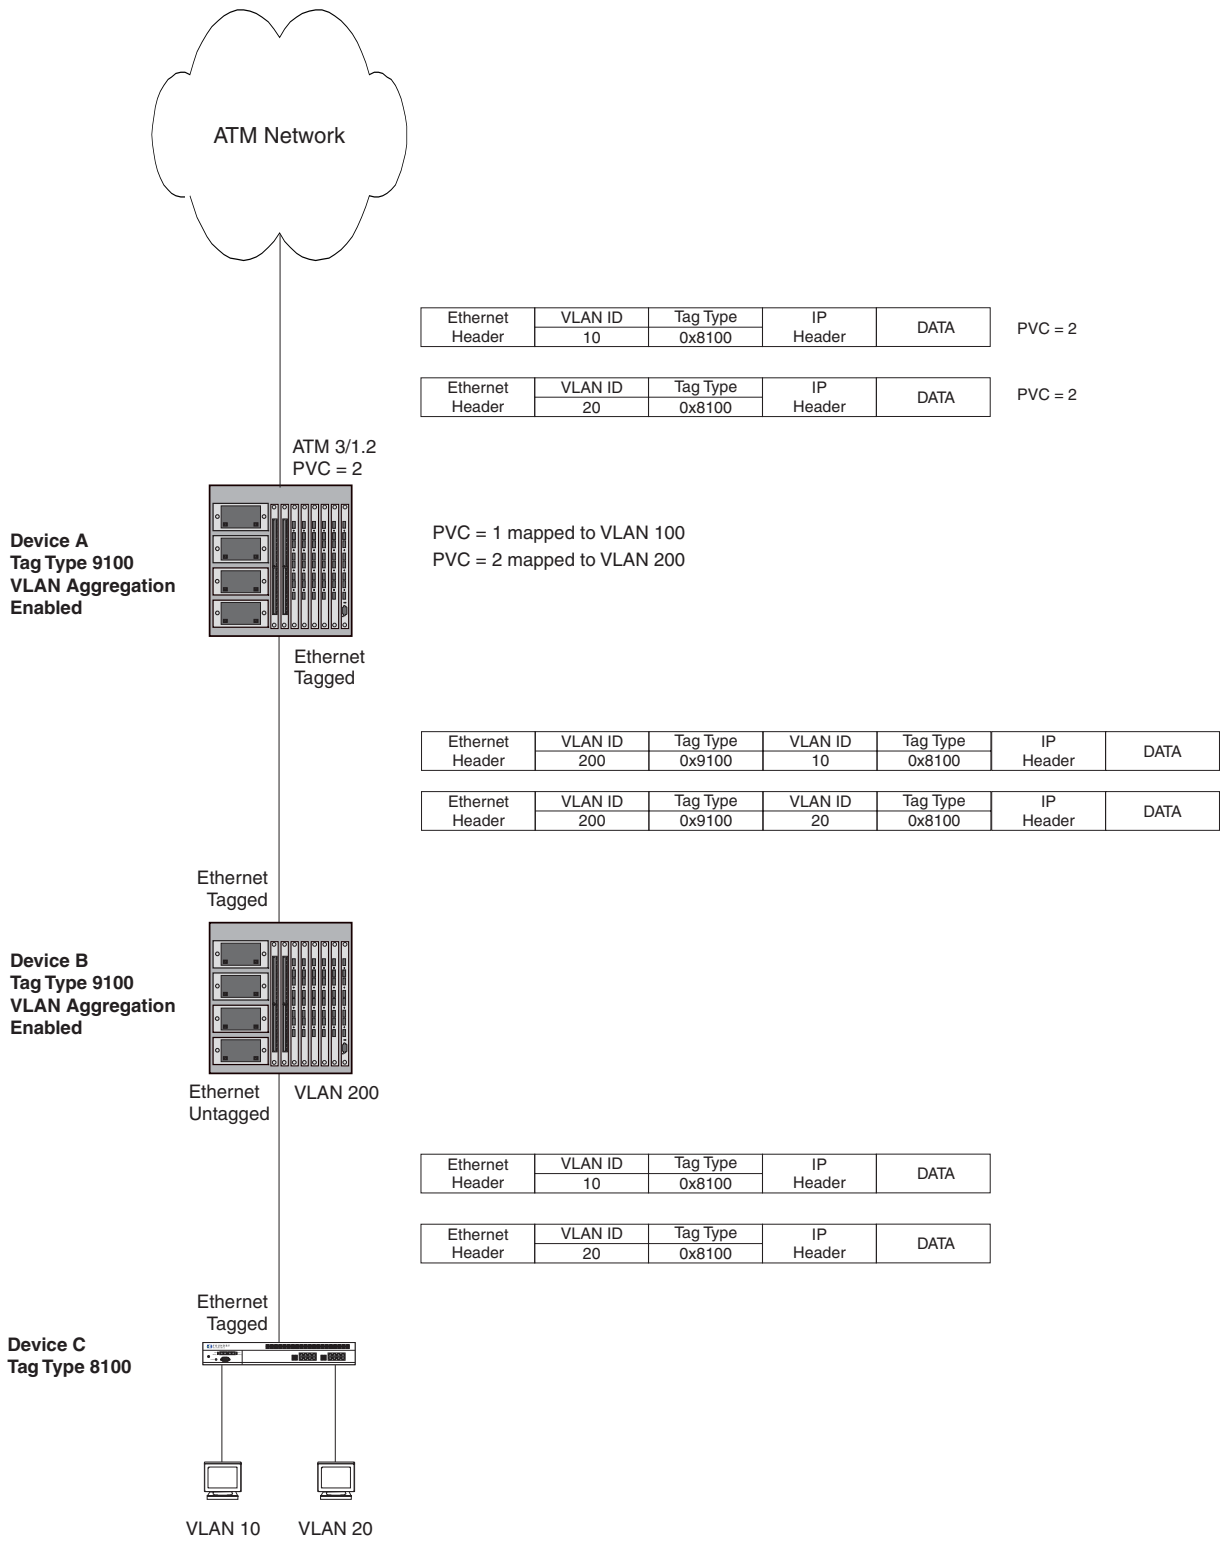
Figure 8.7 PVC-VLAN mapping – hub-and-spoke configuration with two PVCs mapped to two VLANs



You can configure multiple PVCs on the same port to be mapped to multiple VLANs. However, multiple PVCs on the same physical port cannot be part of the same VLAN. For example, in the configuration above, you could not map PVCs on both 2/1.1 and 2/1.2 to VLAN 100.

The PVC-VLAN mapping feature is designed to be used in conjunction with the Super Aggregated VLAN Application. Figure 8.8 illustrates this kind of configuration.

Figure 8.8 PVC-VLAN mapping used with the Super Aggregated VLAN application



In this example, PVC = 2 is mapped to VLAN 200. When the Super Aggregated VLAN application is enabled, the Foundry device aggregates VLANs 10 and 20 into a single VLAN, appending an additional VLAN tag (in this

example, 9100) to the packets. The edge device at the other end of the core removes this VLAN tag and separates the aggregated VLANs into individual VLANs before forwarding the traffic.

For details on configuring the Super Aggregated VLAN application, see “Configuring Super Aggregated VLANs” on page 22-50

Notes

- ATM sub-interfaces can only be assigned to a VLAN as untagged.
- If you are using the N2P software image, you must enable the ports you plan to use for the PVC-VLAN mapping. Ports are disabled by default in the N2P image. Also, you must enable Layer 2 switching (**no route-only** command). In the N2P image, Layer 2 switching is disabled by default.
- A VLAN can have multiple PVCs mapped to it. A PVC can be mapped to only one VLAN.
- Multiple PVCs on the same port cannot be part of the same VLAN.
- This feature supports LLC encapsulation only for bridged Ethernet/802.3 packets. 802.4 and 802.5 frames are rejected.
- This feature supports both FCS (Frame Check Sequence) and non-FCS frames.
- Unlike Ethernet ports, which belong to the default VLAN by default, ATM ports are not part of any VLAN by default. You must explicitly add ATM sub-interfaces to a VLAN.
- Per-VLAN spanning tree only is supported.
- When an ATM sub-interface is configured as a bridging interface, the MTU of the payload is 1528 bytes (compared to a default MTU of 4470 bytes and maximum MTU of 9180 bytes when the sub-interface is configured as a router interface). The 1528-byte length accommodates super-aggregated VLANs.
- A PVC can act as either a bridging component or a routing component, but not both at the same time. On a Layer 3 Switch, an ATM sub-interface is a route-only interface by default, and is converted to a Layer 2 interface when explicitly added to a VLAN. When a PVC is mapped to a VLAN, normal routing packets are discarded.
- Only port-based VLANs are supported. Protocol- and subnet-based VLANs are not supported.
- A PVC on a point-to-multipoint sub-interface cannot be mapped to a VLAN.
- VTP (VLAN Trunking Protocol) is not supported.
- SuperSpan is supported.

Adding an ATM Sub-Interface to a VLAN

To add an ATM sub-interface to a VLAN:

```
BigIron(config)# vlan 60
BigIron(config-vlan-60)# untagged atm 3/1.2
```

Syntax: [no] untagged atm <slot/port.subinterface>

If a PVC has not already been configured on the specified ATM sub-interface, an error message is displayed.

If the ATM sub-interface is configured with an IP address, an error message is displayed. You must remove the IP address from the sub-interface in order to use the sub-interface as a bridging interface.

If any static routes are configured on the specified sub-interface, a message is displayed indicating the static routes should be removed in order to allow the PVC-VLAN mapping.

The VLAN must have already been created prior to assigning an ATM sub-interface to it.

Removing an ATM Sub-Interface from a VLAN

To remove an ATM sub-interface from a VLAN:

```
BigIron(config)# vlan 60
BigIron(config-vlan-60)# no untagged atm 3/1.2
```

Syntax: no untagged atm <slot/port.subinterface>

Enabling Spanning Tree

STP is supported on ATM modules local to the VLAN. Per-VLAN spanning tree is supported.

To enable spanning tree on a VLAN:

```
BigIron(config)# vlan 60
BigIron(config-vlan-60)# spanning-tree
```

Syntax: [no] spanning-tree

Displaying ATM Information

You can display the following ATM module information:

- General module information – see “Displaying General Module Information” on page 8-28
- Module status – see “Determining ATM Module Status” on page 8-29
- Interface parameters – see “Displaying Interface Parameters” on page 8-30
- Port statistics – see “Displaying Port Statistics” on page 8-32
- VC statistics – see “Displaying VC Statistics” on page 8-33

Displaying General Module Information

You can use the following method to display general information about the ATM modules in the chassis, including:

- Module location (slot number)
- Boot and flash code information
- Module uptime

USING THE CLI

To display general ATM module information, enter the following command at any CLI level:

```
BigIron> show atm-state
=====
ATM MODULE (6) App CPU in running mode:
      CPU 0 in state of ATM_STATE_RUNNING
-----
Module 6 App CPU 1, SW: Version 07.2.05T91
Compiled on Apr 25 2001 at 18:16:41 labeled as A2R07205
DRAM 134M, BRAM 33554K, FPGA Version 0000
Code Flash 4M: Primary (429266 bytes, 07.2.05T91),
                Secondary (426573 bytes, 07.2.05T91)
Boot Flash 131K, Boot Version 06.00.00
The system uptime is 0 day 0 hour 31 minute 6 second
General Status: 4 ipc msg rec, 2 ipc msg sent
```

Syntax: show atm-state

USING THE WEB MANAGEMENT INTERFACE

You cannot display general ATM module information using the Web management interface.

Determining ATM Module Status

You can determine the status of an ATM module in the following ways:

- Status LEDs – Each ATM port has LEDs that show link status, transmit and receive activity, and indicate whether an alarm condition has occurred.
- Module information in software – The module information displayed by the software indicates whether the module came up properly.

Status LEDs

You can determine the status of an ATM port by observing its LEDs. Each ATM port has the following LEDs.

Table 8.5: Port LED Indicators for ATM Modules

LED	Position	State	Meaning
Link	Upper left	On	The port is connected.
		Off	The port is not connected.
Alarm	Upper right	On	At least one of the following SONET alarm conditions has been detected: <ul style="list-style-type: none"> • LOS – Loss of Signal • LOF – Loss of Frame • LOP – Loss of Pointer • AIS – Alarm Indication Signal
		Off	None of the alarm conditions listed above have been detected.
TxAct	Lower left	Blinking	The port is transmitting traffic.
RxAct	Lower right	Blinking	The port is receiving traffic.

Software

You can display status information for an ATM module using either of the following methods.

NOTE:

- Slots on a four-slot chassis are numbered 1 – 4, from top to bottom.
- Slots on a eight-slot chassis are numbered 1 – 8, from top to bottom.
- Slots on an fifteen-slot chassis are numbered 1 – 15, from left to right.

USING THE CLI

To display the status of an ATM module using the CLI, enter the following command at any CLI level:

```
BigIron> show module

      Module                               Status  Ports Starting MAC
S1: B8GMR Fiber Management Module, SYSIF M2, ACTIV   8  00e0.5291.5400
S2:
S3:
S4: B24E Copper Switch Module                 OK      24  00e0.5291.5460
S5:
S6: ATM 2 Port 155M Module, SYSIF II      OK       2
S7: Configured as B24E Copper Switch Module
S8:
```

Syntax: show module

In this example, slot 6 contains a two-port ATM module with 155 Mbps ports. This display shows the following information.

Table 8.6: CLI Display of ATM Interface Information

This Field...	Displays...
Slot number	The S<num> value in the left column indicates the chassis slot.
Module	The module type. In this example, slot 6 contains a 2-port ATM module with 155 Mbps (OC-3c) ports.
Status	Shows the module status. An ATM module can have one of the following statuses: <ul style="list-style-type: none"> • FAILED – This status indicates that the host module failed to come up. • OK – This status indicates that the module came up and is operating normally. <p>Note: Management modules have different status values.</p>
Ports	The number of physical interfaces (ports) on the module.
Starting MAC	This column does not apply to ATM modules.

USING THE WEB MANAGEMENT INTERFACE

1. Select the [Home](#) link to display the System panel, if not already displayed.
2. Select the [Module](#) link to display the Module panel. The Status column shows the module status. An ATM module can have one of the following statuses:
 - FAILED – This status indicates that the host module failed to come up.
 - OK – This status indicates that the module came up and is operating normally.

Displaying Interface Parameters

You can display brief information for all interfaces, including ATM interfaces. In addition, you can display ATM interface parameters for individual slots or ports.

To display ATM interface parameters, use the following methods.

USING THE CLI

To display brief information for all interfaces, including ATM interfaces, enter the following command at any CLI level:

```
BigIron> show interfaces brief
```

```
Port Link State   Encap Clock Loop Speed  frame  scram  total vc
6/1  up             llcsnap int  none  155   sonet  yes    5
6/2  up             llcsnap int  none  155   sonet  yes    1
```

Syntax: show interfaces brief

This command shows information for all the ports in the chassis. For simplicity, the example above shows only the information for ATM ports.

The command shows the following information for ATM ports.

Table 8.7: CLI Display of ATM Interface Information

This Field...	Displays...
Port	The chassis slot and port number of the interface.
Link State	The state of the link, which can be one of the following: <ul style="list-style-type: none"> down up To change the link state, see “Changing the Interface State” on page 8-11.
Encap	The encapsulation type in use on the interface. The encapsulation type for a Foundry ATM interface is always llcsnap, which means LLC SNAP encapsulation.
Clock	The clock source, which can be one of the following: <ul style="list-style-type: none"> int – The interface is using the clock on the ATM module. line – The interface is using the clock source supplied on the network. To change this parameter, see “Changing the Clock Source” on page 8-12.
Loop	The loopback state of the interface. The loopback state can be one of the following: <ul style="list-style-type: none"> int – The loopback path consists only of the ATM circuitry on this interface. line – The loopback path consists of both this ATM interface and the ATM interface at the remote end of the link. none – The interface is not operating in loopback mode. To change this parameter, see “Changing the Loopback Path” on page 8-12.
Speed	The bandwidth of the interface.

Table 8.7: CLI Display of ATM Interface Information (Continued)

This Field...	Displays...
Frame	The frame type used on the interface. The frame type is always SONET (shown as "sonet") and is not configurable.
Scram	The state of the ATM scramble mode, which can be one of the following: <ul style="list-style-type: none"> no – Scrambling is disabled. yes – Scrambling is enabled. This is the default.
Total VC	The number of VCs configured on the interface.

Displaying ATM Information for Individual Slots or Ports

To display ATM information for a slot, enter a command such as the following at any level of the CLI:

```
BigIron(config-subif-4/1.1)# show interface atm 6/1
ATM6/1 is up, line protocol is up
  No port name
  Hardware is ATM
  Encapsulation llcsnap, clock is internal
  Framing is SONET, BW 155000Kbit
  Loopback not set, keepalive not set, scramble enabled
  Each virtual path contains 4096 virtual channels
  5 minute input rate: 120360 bits/sec, 167 packets/sec
  5 minute output rate: 90176 bits/sec, 126 packets/sec
  265370 packets input, 23798000 bytes, 0 no buffer
  Received 0 CRCs, 0 shorts, 0 giants, 0 alignments
  205417 packets output, 18306822 bytes, 0 underruns
```

Syntax: show interface atm <slotnum>/<portnum> [to <slotnum>/<portnum>]

Displaying Port Statistics

Use the following method to display byte and packet statistics for ATM ports.

USING THE CLI

To display ATM port statistics, enter a command such as the following at any CLI level:

```
BigIron(config-subif-4/1.1)# show statistics atm 6/1 to 6/2

ATM          Packets          Errors
Port  [Receive Transmit] [Align  FCS  Giant  Short  CRC]
6/1    299582   231143    0      0      0      0      0
  5 minute input rate: 118344 bits/sec, 165 packets/sec
  5 minute output rate: 88880 bits/sec, 125 packets/sec

ATM          Packets          Errors
Port  [Receive Transmit] [Align  FCS  Giant  Short  CRC]
6/2    1051     524      0      0      0      0      0
  5 minute input rate: 952 bits/sec, 0 packets/sec
  5 minute output rate: 96 bits/sec, 0 packets/sec
```

The command shows the following statistics for ATM ports.

Table 8.8: CLI Display of ATM Interface Statistics

This Field...	Displays...
Port	The chassis slot and port number of the interface.
Receive	The total number of packets received on the port.
Transmit	The total number of packets sent on the port.
Align	The number of alignment errors.
FCS	The number of Frame Check Sequence (FCS) errors.
Giant	The number of giant packets.
Short	The number of short packets.
CRC	The number of packets with Cyclic Redundancy Check (CRC) errors.
5 minute input rate	The average receive rate of the port during the last five minute interval.
5 minute output rate	The average send rate of the port during the last five minute interval.

Displaying VC Statistics

Use the following method to display byte and packet statistics for individual VCs.

USING THE CLI

To display ATM VC statistics, enter the following command at any CLI level:

```
BigIron(config-subif-4/1.1)# show atm vc
```

Port	VPI	VCI	InPkts	OutPkts	InBytes	OutBytes
6/1.1	1	100	69626	1129	6330289	68211
6/1.2	1	101	214784	307	19170286	28584
6/1.3	2	200	4910	353	439675	34612
6/1.4	2	201	296	221596	28529	19768722
6/1.5	2	202	220	269	23252	29349
6/2.1	1	13	1040	515	389135	45016

Syntax: show atm vc [slot <slotnum> | interface <portnum> | <slotnum>/<portnum>.sub-interface <subif>]

You can use the command's optional parameters to refine the display request.

- The **slot** <slotnum> parameter displays VC statistics for the ATM interfaces on the ATM module in the specified chassis slot.
- The **interface** <portnum> parameter displays VC statistics for the specified ATM interface.
- The **sub-interface** <subif> parameter displays VC statistics for the specified sub-interface number.

The **show atm vc** commands display the following information.

Table 8.9: CLI Display of ATM VC Statistics

This Field...	Displays...
Port	The chassis slot and port number of the interface, and the number of the sub-interface.
VPI	The VP ID.
VCI	The VC ID.
InPkts	<p>The number of packets received on this VC.</p> <p>Note: The statistic is only for the VC within the VP and on the sub-interface indicated. Thus, if the same chassis has more than one VC with the same VC ID, each of those VCs is a unique interface and the software lists the statistics for each of them separately.</p>
OutPkts	The number of packets sent on this VC.
InBytes	The number of bytes received on this VC.
OutBytes	The number of bytes sent on this VC.

Chapter 9

Configuring Basic Features

This chapter describes how to configure basic, non-protocol features on Foundry devices using the CLI and Web management interface.

This chapter contains procedures for configuring the following parameters:

- Basic system parameters – see “Configuring Basic System Parameters” on page 9-3
- Basic port parameters – see “Configuring Basic Port Parameters” on page 9-21
- Basic Layer 2 parameters – see “Configuring Basic Layer 2 Parameters” on page 9-30
- Basic Layer 3 parameters – see “Enabling or Disabling Routing Protocols” on page 9-47
- System defaults and table sizes – see “Displaying and Modifying System Parameter Default Settings” on page 9-48
- Temperature sensor parameters – see “Using the Temperature Sensor” on page 9-52
- Mirror ports (for traffic diagnosis and troubleshooting) – see “Assigning a Mirror Port and Monitor Ports” on page 9-55

Foundry devices are configured at the factory with default parameters that allow you to begin using the basic features of the system immediately. However, many of the advanced features such as VLANs or routing protocols for the router must first be enabled at the system (global) level before they can be configured.

- If you use the Command Line Interface (CLI) to configure system parameters, you can find these system level parameters at the Global CONFIG level of the CLI.
- If you use the Web management interface, you enable or disable system level parameters on the System configuration panel, which is displayed by default when you start a management session. Figure 9.1 shows an example of the System configuration panel on a BigIron Layer 3 Switch.

NOTE: Before assigning or modifying any router parameters, you must assign the IP subnet (interface) addresses for each port.

NOTE: This chapter does not describe how to configure Virtual LANs (VLANs) or link aggregation. For VLAN configuration information, see “Configuring Virtual LANs (VLANs)” on page 22-1. For link aggregation information, see “Configuring Trunk Groups and Dynamic Link Aggregation” on page 17-1.

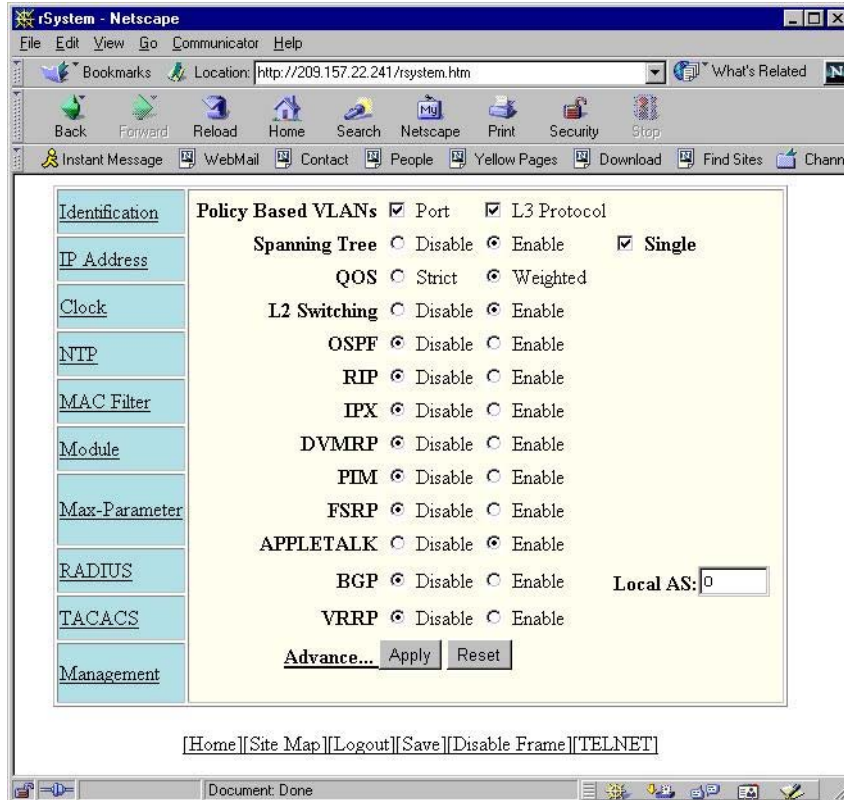
NOTE: For information about configuring IP addresses, DNS resolver, DHCP assist, and other IP-related parameters, see the “Configuring IP” chapter of the *Foundry Enterprise Configuration and Management Guide*.

For information about the Syslog buffer and messages, see “Using Syslog” on page A-1.

Using the Web Management Interface for Basic Configuration Changes

The Web management interface enables you to easily make numerous configuration changes by entering or changing information on configuration panels such as the one shown in Figure 9.1. This example is for a Layer 3 Switch. Layer 2 Switches do not have routing options but do have some additional options not available on Layer 3 Switches.

Figure 9.1 System configuration panel for a Foundry Layer 3 Switch



You can perform the following configuration tasks from the System configuration panel:

- Enter system administration information.
- Review or modify the IP, mask, and gateway addresses (Layer 2 Switches only).
- Assign IP subnet (interface) addresses and masks (Layer 3 Switches only).
- Assign DHCP gateway lists for DHCP Assist operation (Layer 2 Switches only).
- Configure Domain Name Server (DNS) Resolver.
- Define a MAC address filter.
- Set the system clock.
- Configure the device to use a Simple Network Time Protocol (SNTP) server.
- Enable port-based and/or Layer 3 protocol VLANs.
- Enable IP and IPX Layer 3 router acceleration (Stackable backbone Layer 2 Switches only).
- Enable or disable IP Multicast Traffic Reduction (Layer 2 Switches only).

- Enable or disable IGMP (Layer 2 Switches only).
- Enable or disable protocol—OSPF, RIP, IPX, DVMRP, PIM, FSRP, VRRP, BGP4, AppleTalk (Layer 3 Switches only).
- Assign Layer 4 QoS Priority (Layer 2 Switches only).

NOTE: Layer 4 priority for routers is set using the IP policy command found at the global CONFIG level of the CLI and the IP configuration sheet for the Web management interface.

- Enable or disable Spanning Tree Protocol.
- Enable or disable SNMP operation and configure SNMP community strings, trap receivers, and other parameters.
- Enable or disable IEEE 802.1q VLAN tagging.
- Enable or disable Layer 2 switching (Layer 3 Switches only).
- Enable or disable Telnet.
- Change the aging period (switch age time) for entries in the address table.
- Assign a mirror port.
- Modify system parameters.
- Add or delete modules (Chassis devices only).
- Modify tag type.
- Modify telnet timeout period.
- Modify broadcast limit.
- Enable or disable management using the Web management interface.
- Apply base (system) default values (Layer 2 Switches only).
- Configure redundant management module parameters (NetIron or BigIron Layer 3 Switch with Management 2 or higher modules only).

The procedures in this chapter describe how to configure these parameters.

Configuring Basic System Parameters

The procedures in this section describe how to configure the following basic system parameters:

- System name, contact, and location – see “Entering System Administration Information” on page 9-4
- SNMP trap receiver, trap source address, and other parameters – see “Configuring Simple Network Management (SNMP) Parameters” on page 9-5
- Single source address for all Telnet packets – “Configuring an Interface as the Source for All Telnet Packets” on page 9-11
- Single source address for all TFTP packets – “Configuring an Interface as the Source for All TFTP Packets” on page 9-12
- System time using a Simple Network Time Protocol (SNTP) server or local system counter – see “Specifying a Simple Network Time Protocol (SNTP) Server” on page 9-12 and “Setting the System Clock” on page 9-14
- Default Gigabit negotiation mode (for Chassis devices) – see “Changing the Default Gigabit Negotiation Mode” on page 9-16
- Broadcast, multicast, or unknown-unicast limits, if required to support slower third-party devices – see “Limiting Broadcast, Multicast, or Unknown-Unicast Rates” on page 9-18

- Banners that are displayed on users' terminals when they enter the Privileged EXEC CLI level or access the device through Telnet – see “Configuring CLI Banners” on page 9-20.
- Terminal display length – see “Configuring Terminal Display” on page 9-21.

NOTE: For information about the Syslog buffer and messages, see “Using Syslog” on page A-1.

Entering System Administration Information

You can configure a system name, contact, and location for a Foundry Layer 2 Switch or Layer 3 Switch and save the information locally in the configuration file for future reference. This information is not required for system operation but is suggested. When you configure a system name, the name replaces the default system name in the CLI command prompt. For example, if the system is a BigIron 8000, the system name you configure replaces “BigIron” in the command prompt.

The name, contact, and location each can be up to 32 alphanumeric characters.

NOTE: If you install Layer 2 Switch code on a Layer 3 Switch, the CLI command prompt begins with “SW-” to indicate the software change. This is true even if you change the system name.

USING THE CLI

Here is an example of how to configure a Layer 2 Switch or Layer 3 Switch name, system contact, and location:

```
BigIron(config)# hostname home
home(config)# snmp-server contact Suzy Sanchez
home(config)# snmp-server location Centerville
home(config)# end
home# write memory
```

Syntax: hostname <string>

Syntax: snmp-server contact <string>

Syntax: snmp-server location <string>

The text strings can contain blanks. The SNMP text strings do not require quotation marks when they contain blanks but the host name does.

NOTE: The **chassis name** command does not change the CLI prompt. Instead, the command assigns an administrative ID to the device.

USING THE WEB MANAGEMENT INTERFACE

Here is an example of how to configure a Layer 2 Switch or Layer 3 Switch name, system contact, and location:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the Identification link to display the following panel.

Identification

Name:	BigIron Router
Contact:	Suzy Creamcheese
Location:	Centerville

[Home](#) | [Site Map](#) | [Logout](#) | [Save](#) | [Frame Enable](#) | [Disable](#) | [TELNET](#)

3. Edit the value in the Name field to change the device name. The name can contain blanks.
4. Enter the name of the administrator for the device in the Contact field. The name can contain blanks.
5. Enter the device's location in the Location field. The location can contain blanks.
6. Click the Apply button to save the change to the device's running-config file.
7. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You also can access the dialog for saving configuration changes by clicking on the plus sign next to Command in the tree view, then clicking on [Save to Flash](#).

Configuring Simple Network Management (SNMP) Parameters

Use the procedures in this section to perform the following configuration tasks:

- Specify an SNMP trap receiver.
- Specify a source address and community string for all traps sent by the device.
- Change the holddown time for SNMP traps
- Disable individual SNMP traps. (All traps are enabled by default.)
- Disable traps for CLI access that is authenticated by a local user account, a RADIUS server, or a TACACS/TACACS+ server.

NOTE: To add and modify “get” (read-only) and “set” (read-write) community strings, see the *Foundry Security Guide*.

Specifying an SNMP Trap Receiver

You can specify a trap receiver to ensure that all SNMP traps sent by the Foundry device go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. When you specify the host, you also specify a community string. The Foundry device sends all the SNMP traps to the specified host(s) and includes the specified community string. Administrators can therefore filter for traps from a Foundry device based on IP address or community string.

When you add a trap receiver, the software automatically encrypts the community string you associate with the receiver when the string is displayed by the CLI or Web management interface. If you want the software to show the community string in the clear, you must explicitly specify this when you add a trap receiver. In either case, the software does not encrypt the string in the SNMP traps sent to the receiver.

To specify the host to which the device sends all SNMP traps, use one of the following methods.

USING THE CLI

To add a trap receiver and encrypt the display of the community string, enter commands such as the following:

To specify an SNMP trap receiver and change the UDP port that will be used to receive traps, enter a command such as the following:

```
BigIron(config)# # snmp-server host 2.2.2.2 0 mypublic port 200
BigIron(config)# write memory
```

Syntax: snmp-server host <ip-addr> [0 | 1] <string> [port <value>]

The <ip-addr> parameter specifies the IP address of the trap receiver.

The **0 | 1** parameter specifies whether you want the software to encrypt the string (**1**) or show the string in the clear (**0**). The default is **0**.

The <string> parameter specifies an SNMP community string configured on the Foundry device. The string can be a read-only string or a read-write string. The string is not used to authenticate access to the trap host but is instead a useful method for filtering traps on the host. For example, if you configure each of your Foundry devices

that use the trap host to send a different community string, you can easily distinguish among the traps from different Foundry devices based on the community strings.

The command in the example above adds trap receiver 2.2.2.2 and configures the software to encrypt display of the community string. When you save the new community string to the startup-config file (using the **write memory** command), the software adds the following command to the file:

```
snmp-server host 2.2.2.2 1 <encrypted-string>
```

To add a trap receiver and configure the software to encrypt display of the community string in the CLI and Web management interface, enter commands such as the following:

```
BigIron(config)# snmp-server host 2.2.2.2 0 BigIron-12
BigIron(config)# write memory
```

The **port <value>** parameter allows you to specify which UDP port will be used by the trap receiver. This parameter allows you to configure several trap receivers in a system. With this parameter, IronView Network Manager Network Manager and another network management application can coexist in the same system. Foundry devices can be configured to send copies of traps to more than one network management application.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access.
2. Click the Management link to display the Management configuration panel.
3. Click the Trap Receiver link to display the Trap Receiver panel.

Trap Receiver

IP Address	UDP Port Number	Community String	Encrypt		
207.95.6.75	200	public	no	Delete	Modify
207.95.6.75	162	public	no	Delete	Modify
207.95.6.75	300	mypublic	no	Delete	Modify
IP Address	UDP Port Number	Community String	Encrypt		

[Add Trap Receiver]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

4. Click Add Trap Receiver link to add a new trap receiver and display the following panel.

Trap Receiver

IP Address:	<input type="text"/>
UDP Port Number:	<input type="text"/>
Community String:	<input type="text"/>
Encrypt:	<input type="checkbox"/>

[Show]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

5. Enter the IP address of the receiver in the IP Address field.
6. Enter the UDP port number that will be used to receive traps. If no port number is entered, then UDP port 162 will be used by trap receivers.
7. Enter the community string you want the Layer 3 Switch to send in traps sent to this host in the Community String field.
8. Select the Encrypt checkbox to remove the checkmark if you want to disable encryption of the string display. Encryption prevents other users from seeing the string in the CLI or Web management interface. If you disable encryption, other users can view the community string. Encryption is enabled by default.

To re-enable encryption, select the checkbox to place a checkmark in the box.
9. Click Add to apply the change to the device's running-config file.
10. Select the [Save](#) link at the bottom of the panel. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Specifying a Single Trap Source

You can specify a single trap source to ensure that all SNMP traps sent by the Foundry device use the same source IP address. When you configure the SNMP source address, you specify the Ethernet port, POS port, loopback interface, or virtual routing interface that is the source for the traps. The Foundry device then uses the lowest-numbered IP address configured on the port or interface as the source IP address in the SNMP traps sent by the device.

Identifying a single source IP address for SNMP traps provides the following benefits:

- If your trap receiver is configured to accept traps only from specific links or IP addresses, you can use this feature to simplify configuration of the trap receiver by configuring the Foundry device to always send the traps from the same link or source address.
- If you specify a loopback interface as the single source for SNMP traps, SNMP trap receivers can receive traps regardless of the states of individual links. Thus, if a link to the trap receiver becomes unavailable but the receiver can be reached through another link, the receiver still receives the trap, and the trap still has the source IP address of the loopback interface.

To specify a port, loopback interface, or virtual routing interface whose lowest-numbered IP address the Foundry device must use as the source for all SNMP traps sent by the device, use the following CLI method.

USING THE CLI

To configure the device to send all SNMP traps from the first configured IP address on port 4/11, enter the following commands:

```
BigIron(config)# snmp-server trap-source ethernet 4/11
BigIron(config)# write memory
```

Syntax: snmp-server trap-source loopback <num> | ethernet <portnum> | pos <portnum> | ve <num>

The <num> parameter is a loopback interface or virtual routing interface number. If you specify an Ethernet or POS port, the <portnum> is the port's number (including the slot number, if you are configuring a Chassis device).

To specify a loopback interface as the device's SNMP trap source, enter commands such as the following:

```
BigIron(config)# int loopback 1
BigIron(config-lbif-1)# ip address 10.0.0.1/24
BigIron(config-lbif-1)# exit
BigIron(config)# snmp-server trap-source loopback 1
```

The commands in this example configure loopback interface 1, assign IP address 10.00.1/24 to the loopback interface, then designate the interface as the SNMP trap source for this Layer 3 Switch. Regardless of the port the Foundry device uses to send traps to the receiver, the traps always arrive from the same source IP address.

The following commands configure an IP interface on a POS port and designate the address as the SNMP trap source for a Layer 3 Switch. The Foundry device always sends traps through the POS port and the source IP address of the traps is always the lowest-numbered IP address configured on the POS port.

```
BigIron(config)# interface pos 2/1
BigIron(config-posif-2/1)# ip address 209.157.22.26/24
BigIron(config-posif-2/1)# exit
BigIron(config)# snmp-server trap-source pos 2/1
```

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a trap source using the Web management interface.

Setting the SNMP Trap Holddown Time

When a Foundry device starts up, the software waits for Layer 2 convergence (STP) and Layer 3 convergence (OSPF) before beginning to send SNMP traps to external SNMP servers. Until convergence occurs, the device might not be able to reach the servers, in which case the messages are lost.

By default, a Foundry device uses a one-minute holddown time to wait for the convergence to occur before starting to send SNMP traps. After the holddown time expires, the device sends the traps, including traps such as “cold start” or “warm start” that occur before the holddown time expires.

You can change the holddown time to a value from one second to ten minutes.

USING THE CLI

To change the holddown time for SNMP traps, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# snmp-server enable traps holddown-time 30
```

The command in this example changes the holddown time for SNMP traps to 30 seconds. The device waits 30 seconds to allow convergence in STP and OSPF before sending traps to the SNMP trap receiver.

Syntax: [no] snmp-server enable traps holddown-time <secs>

The <secs> parameter specifies the number of seconds and can be from 1 – 600 (ten minutes). The default is 60 seconds.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure the parameter using the Web management interface.

Disabling SNMP Traps

Foundry Layer 2 Switches and Layer 3 Switches come with SNMP trap generation enabled by default for all traps. You can selectively disable one or more of the following traps.

NOTE: By default, all SNMP traps are enabled at system startup.

Layer 2 Switch Traps

The following traps are generated on the Layer 2 Switches:

- SNMP authentication keys
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- Module insert (applies only to Chassis devices)

- Module remove (applies only to Chassis devices)

Layer 3 Switch Traps

The following traps are generated on the Layer 3 Switches:

- SNMP authentication key
- Power supply failure
- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Locked address violation
- Module insert
- Module remove
- BGP4
- OSPF
- FSRP
- VRRP
- VRRPE

ServerIron Traps

See the *Foundry ServerIron Installation and Configuration Guide*.

USING THE CLI

To stop link down occurrences from being reported, enter the following:

```
BigIron(config)# no snmp-server enable traps link-down
```

Syntax: [no] snmp-server enable traps <trap-type>

NOTE: For a list of the trap values, see the *Foundry Switch and Router Command Line Interface Reference*.

USING THE WEB MANAGEMENT INTERFACE

To enable or disable individual SNMP traps:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the Management link to display the Management panel.
3. Click on the Trap link to display the list of traps that you can enable or disable.

NOTE: The panel lists different traps for Layer 2 Switches and Layer 3 Switches.

4. Select the Disable or Enable button next to the trap you want to disable or enable.
5. Click the Apply button to save the change to the device's running-config file.
6. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Disabling Syslog Messages and Traps for CLI Access

Foundry devices send Syslog messages and SNMP traps when a user logs into or out of the User EXEC or Privileged EXEC level of the CLI. The feature applies to users whose access is authenticated by an authentication-method list based on a local user account, RADIUS server, or TACACS/TACACS+ server.

NOTE: The Privileged EXEC level is sometimes called the “Enable” level, because the command for accessing this level is **enable**.

The feature is enabled by default.

Examples of Syslog Messages for CLI Access

When a user whose access is authenticated by a local user account, a RADIUS server, or a TACACS/TACACS+ server logs into or out of the CLI's User EXEC or Privileged EXEC mode, the software generates a Syslog message and trap containing the following information:

- The time stamp
- The user name
- Whether the user logged in or out
- The CLI level the user logged into or out of (User EXEC or Privileged EXEC level)

NOTE: Messages for accessing the User EXEC level apply only to access through Telnet. The device does not authenticate initial access through serial connections but does authenticate serial access to the Privileged EXEC level. Messages for accessing the Privileged EXEC level apply to access through the serial connection or Telnet.

The following examples show login and logout messages for the User EXEC and Privileged EXEC levels of the CLI:

```
BigIron(config)# show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
Buffer logging: level ACDMEINW, 12 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Oct 15 18:01:11:info:dg logout from USER EXEC mode
Oct 15 17:59:22:info:dg logout from PRIVILEGE EXEC mode
Oct 15 17:38:07:info:dg login to PRIVILEGE EXEC mode
Oct 15 17:38:03:info:dg login to USER EXEC mode
```

Syntax: show logging

The first message (the one on the bottom) indicates that user “dg” logged in to the CLI's User EXEC level on October 15 at 5:38 PM and 3 seconds (Oct 15 17:38:03). The same user logged into the Privileged EXEC level four seconds later.

The user remained in the Privileged EXEC mode until 5:59 PM and 22 seconds. (The user could have used the CONFIG modes as well. Once you access the Privileged EXEC level, no further authentication is required to access the CONFIG levels.) At 6:01 PM and 11 seconds, the user ended the CLI session.

Disabling the Syslog Messages and Traps

Logging of CLI access is enabled by default. If you want to disable the logging, use the following method.

USING THE CLI

To disable logging of CLI access, enter the following commands:

```
BigIron(config)# no logging enable user-login
BigIron(config)# write memory
BigIron(config)# end
BigIron# reload
```

Syntax: [no] logging enable user-login

USING THE WEB MANAGEMENT INTERFACE

You cannot disable logging of CLI access using the Web management interface.

Configuring an Interface as the Source for All Telnet Packets

You can designate the lowest-numbered IP address configured on an interface as the source IP address for all Telnet packets from the Layer 3 Switch. Identifying a single source IP address for Telnet packets provides the following benefits:

- If your Telnet server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the Telnet server by configuring the Foundry device to always send the Telnet packets from the same link or source address.
- If you specify a loopback interface as the single source for Telnet packets, Telnet servers can receive the packets regardless of the states of individual links. Thus, if a link to the Telnet server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS/TACACS+, and RADIUS packets. You can configure a source interface for one or more of these types of packets.

To specify an interface as the source for all Telnet packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the interface as the source IP address for Telnet packets originated by the device.

USING THE CLI

To specify the lowest-numbered IP address configured on a virtual routing interface as the device's source for all Telnet packets, enter commands such as the following:

```
BigIron(config)# int loopback 2
BigIron(config-lbif-2)# ip address 10.0.0.2/24
BigIron(config-lbif-2)# exit
BigIron(config)# ip telnet source-interface loopback 2
```

The commands in this example configure loopback interface 2, assign IP address 10.0.0.2/24 to the interface, then designate the interface as the source for all Telnet packets from the Layer 3 Switch.

Syntax: ip telnet source-interface atm <portnum>.<subif> | ethernet <portnum> | loopback <num> | ve <num>

The following commands configure an IP interface on an Ethernet port and designate the address port as the source for all Telnet packets from the Layer 3 Switch.

```
BigIron(config)# interface ethernet 1/4
BigIron(config-if-1/4)# ip address 209.157.22.110/24
BigIron(config-if-1/4)# exit
BigIron(config)# ip telnet source-interface ethernet 1/4
```

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a single Telnet source using the Web management interface.

Canceling an Outbound Telnet Session

If you want to cancel a Telnet session from the console to a remote Telnet server (for example, if the connection is frozen), you can terminate the Telnet session by doing the following:

1. At the console, press Ctrl-^ (Ctrl-Shift-6).
2. Press the X key to terminate the Telnet session.

Pressing Ctrl-^ twice in a row causes a single Ctrl-^ character to be sent to the Telnet server. After you press Ctrl-^, pressing any key other than X or Ctrl-^ returns you to the Telnet session.

Configuring an Interface as the Source for All TFTP Packets

You can configure the device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, Ethernet port or POS port as the source for all TFTP packets from the device. The software uses the lowest-numbered IP address configured on the interface as the source IP address for the packets.

For example, to specify the lowest-numbered IP address configured on a virtual routing interface as the device's source for all TFTP packets, enter commands such as the following:

```
BigIron(config)# int ve 1
BigIron(config-vif-1)# ip address 10.0.0.3/24
BigIron(config-vif-1)# exit

BigIron(config)# ip tftp source-interface ve 1
```

The commands in this example configure virtual routing interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface's address as the source address for all TFTP packets

Syntax: [no] ip tftp source-interface atm <portnum>.<subif> | ethernet <portnum> | loopback <num> | ve <num>

The default is the lowest-numbered IP address configured on the port through which the packet is sent. The address therefore changes, by default, depending on the port.

Specifying a Simple Network Time Protocol (SNTP) Server

You can configure Layer 2 Switches and Layer 3 Switches to consult SNTP servers for the current system time and date.

NOTE: Foundry Layer 2 Switches and Layer 3 Switches do not retain time and date information across power cycles. Unless you want to reconfigure the system time counter each time the system is reset, Foundry Networks recommends that you use the SNTP feature.

USING THE CLI

To identify an SNTP server with IP address 208.99.8.95 to act as the clock reference for a Layer 2 Switch or Layer 3 Switch, enter the following:

```
BigIron(config)# sntp server 208.99.8.95
```

Syntax: sntp server <ip-addr> | <hostname> [<version>]

The <version> parameter specifies the SNTP version the server is running and can be from 1 – 4. The default is 1. You can configure up to three SNTP servers by entering three separate **sntp server** commands.

By default, the Layer 2 Switch or Layer 3 Switch polls its SNTP server every 30 minutes (1800 seconds). To configure the Layer 2 Switch or Layer 3 Switch to poll for clock updates from a SNTP server every 15 minutes, enter the following:

```
BigIron(config)# sntp poll-interval 900
```

Syntax: [no] sntp poll-interval <1-65535>

To display information about SNTP associations, enter the following command:

```
BigIron# show sntp associations
  address      ref clock      st  when  poll  delay  disp
~207.95.6.102  0.0.0.0        16  202   4    0.0    5.45
~207.95.6.101  0.0.0.0        16  202   0    0.0    0.0
* synced, ~ configured
```

Syntax: show sntp associations

The following table describes the information displayed by the **show sntp associations** command.

Table 9.1: Output from the show sntp associations command

This Field...	Displays...
(leading character)	One or both of the following: * Synchronized to this peer ~ Peer is statically configured
address	IP address of the peer
ref clock	IP address of the peer's reference clock
st	NTP stratum level of the peer
when	Amount of time since the last NTP packet was received from the peer
poll	Poll interval in seconds
delay	Round trip delay in milliseconds
disp	Dispersion in seconds

To display information about SNTP status, enter the following command:

```
BigIron# show sntp status
Clock is unsynchronized, stratum = 0, no reference clock
precision is 2**0
reference time is 0 .0
clock offset is 0.0 msec, root delay is 0.0 msec
root dispersion is 0.0 msec, peer dispersion is 0.0 msec
```

Syntax: show sntp status

The following table describes the information displayed by the **show sntp status** command.

Table 9.2: Output from the show sntp status command

This Field...	Indicates...
unsynchronized	System is not synchronized to an NTP peer.
synchronized	System is synchronized to an NTP peer.
stratum	NTP stratum level of this system

Table 9.2: Output from the show sntp status command (Continued)

This Field...	Indicates...
reference clock	IP Address of the peer (if any) to which the unit is synchronized
precision	Precision of this system's clock (in Hz)
reference time	Reference time stamp
clock offset	Offset of clock to synchronized peer
root delay	Total delay along the path to the root clock
root dispersion	Dispersion of the root path
peer dispersion	Dispersion of the synchronized peer

USING THE WEB MANAGEMENT INTERFACE

To identify a reference SNTP server for the system:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the [NTP](#) link to display the NTP panel.
3. Optionally change the polling time by editing the value in the Polling Time field, then click Apply to save the change in the device's running-config file. You can specify a number from 1 – 65535.
4. Select the [NTP Server](#) link to display the NTP Server panel.

NOTE: If you have already configured an SNTP server, the server information is listed; otherwise, select the [Add NTP Server](#) link at the bottom of the panel to add a new SNTP server.

5. Enter the IP address of the SNTP server.
6. Select the SNTP version the server is running from the version field's pulldown menu. The default version is 1.
7. Click the Add button to save the change to the device's running-config file.
8. Repeat steps 5 – 7 up to two more times to add a total of three SNTP servers.
9. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Setting the System Clock

In addition to SNTP support, Foundry switches and routers also allow you to set the system time counter. The time counter setting is not retained across power cycles and is not automatically synchronized with an SNTP server. The counter merely starts the system time and date clock with the time and date you specify.

NOTE: You can synchronize the time counter with your SNTP server time by entering the **sntp sync** command from the Privileged EXEC level of the CLI.

NOTE: Unless you identify an SNTP server for the system time and date, you will need to re-enter the time and date following each reboot.

For more details about SNTP, see “Specifying a Simple Network Time Protocol (SNTP) Server” on page 9-12.

USING THE CLI

To set the system time and date to 10:15:05 on October 15, 1999, enter the following command:

```
BigIron# clock set 10:15:05 10-15-99
```

Syntax: [no] clock set <hh:mm:ss> <mm-dd-yy> | <mm-dd-yyyy>

By default, Foundry switches and routers do not change the system time for daylight savings time. To enable daylight savings time, enter the following command:

```
BigIron# clock summer-time
```

Syntax: clock summer-time

Although SNTP servers typically deliver the time and date in Greenwich Mean Time (GMT), you can configure the Layer 2 Switch or Layer 3 Switch to adjust the time for any one-hour offset from GMT or for one of the following U.S. time zones:

- US Pacific (default)
- Alaska
- Aleutian
- Arizona
- Central
- East-Indiana
- Eastern
- Hawaii
- Michigan
- Mountain
- Pacific
- Samoa

The default is US Pacific.

To change the time zone to Australian East Coast time (which is normally 10 hours ahead of GMT), enter the following command:

```
BigIron(config)# clock timezone gmt gmt+10
```

Syntax: clock timezone gmt gmt | us <time-zone>

You can enter one of the following values for <time-zone>:

- US time zones (**us**): alaska, aleutian, arizona, central, east-indiana, eastern, hawaii, michigan, mountain, pacific, samoa.
- GMT time zones (**gmt**): gmt+12, gmt+11, gmt+10...gmt+01, gmt+00, gmt-01...gmt-10, gmt-11, gmt-12.

USING THE WEB MANAGEMENT INTERFACE

To set the local time for the system:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.

2. Select the [Clock](#) link to display the Clock panel, shown below.

Clock

Time Zone:	GMT+00
Daylight Saving Time:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Date (mm-dd-yyyy):	10 23 1999
Time (hh:mm:ss):	1 0 32 PM

[Home](#)
[Site Map](#)
[Logout](#)
[Save](#)
[Frame Enable/Disable](#)
[TELNET](#)

3. Select the time zone by selecting the offset from Greenwich Mean Time that applies to your time zone. For example, to set your device to California time, select GMT-08, which means Greenwich Mean Time minus eight hours.

NOTE: You do not need to adjust for Daylight Savings Time. You enable or disable Daylight Savings Time separately in the following step.

4. Select Disable or Enable next to Daylight Saving Time to enable or disable it.
5. Enter the month, day, and year in the Date fields. You must enter the year as four digits.
6. Enter the hour, minute, and seconds in the Time fields.
7. Select AM or PM.
8. Click Apply to save the changes to the device's running-config file.
9. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Changing the Default Gigabit Negotiation Mode

You can configure the default Gigabit negotiation mode to be one of the following:

- Negotiate-full-auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default for Chassis devices (as well as the Turbolron/8).
- Auto-Gigabit – The port tries to perform a handshake with the other port to exchange capability information.
- Negotiation-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

NOTE: This feature applies only to Chassis devices and the Turbolron/8. For Stackable devices, the default behavior is auto-Gigabit, but you can configure individual Gigabit ports on Stackable devices for negotiation-off. See "Changing the 802.3x Gigabit Negotiation Mode" on page 9-28.

Although the standard for 100BaseTX ports provides an option for a negotiating port to link with a non-negotiating port, the 802.3x standard for Gigabit ports does not provide this option. As a result, unless the ports at both ends of a Gigabit Ethernet link use the same mode (either auto-Gigabit or negotiation-off), the ports cannot establish a link. An administrator must intervene to manually configure one or both sides of the link to enable the ports to establish the link.

Foundry Chassis software provides a solution by changing the default negotiation behavior for Gigabit Ethernet ports on Chassis devices. The new default behavior allows a port to establish a link with another port whether the other port is configured for auto-Gigabit or negotiation-off. By default, Gigabit Ethernet ports first attempt auto-

Gigabit. If auto-Gigabit does not succeed (typically because the port at the other end is not configured for auto-Gigabit), the port switches to negotiation-off.

Backward Compatibility

When you upgrade a Layer 3 Switch that is running software older than 05.2.00, the new software makes modifications to the running-config and startup-config files to ensure that the negotiation settings remain unchanged for the installed device. For new devices running 05.2.00, the default for all Gigabit Ethernet ports is negotiate-full-auto.

To provide the backward compatibility, the software places a line in the running-config file to identify the software version that generated the file. For software release 05.2.00, the version line is as follows: "version 05.2.00". When you save configuration changes to the startup-config file, the software assumes, based on the presence of the version line in the running-config file, that the device is running software release 05.2.00 or later, which contains the change to the Gigabit Ethernet negotiation default.

If the device already has a startup-config file when you update to software release 05.2.00, the software adds the following command to the startup-config file: **gig-default neg-off**. This command sets the global negotiation mode to negotiation-off, the default behavior in software releases earlier than 05.2.00. By setting the default mode to negotiation-off, the new software ensures that the device's Gigabit Ethernet links continue to operate as before. (Although you cannot set a global default for Gigabit Ethernet negotiation in software releases earlier than 05.2.00, the implicit default behavior is negotiation-off.)

If the startup-config file contains the **auto-gig** command to configure individual ports for auto-Gigabit, the command is changed to the new format, **gig-default auto-gig**. Thus, the ports continue to use the auto-Gigabit setting.

NOTE: Software release 05.2.00 and later also adds a version line to the running-config file on Stackable devices. However, the command syntax and default behavior for Gigabit Ethernet ports on Stackable devices is unchanged from earlier software releases.

Changing the Negotiation Mode

You can change the negotiation mode globally and for individual ports. Use either of the following methods.

USING THE CLI

To change the mode globally, enter a command such as the following:

```
BigIron(config)# gig-default neg-off
```

This command changes the global setting to negotiation-off. The global setting applies to all Gigabit Ethernet ports except those for which you set a different negotiation mode on the port level.

To change the mode for individual ports, enter commands such as the following:

```
BigIron(config)# int ethernet 4/1 to 4/4
BigIron(config-mif-4/1-4/4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gigabit for ports 4/1 – 4/4.

Here is the syntax for globally changing the negotiation mode.

Syntax: gig-default neg-full-auto | auto-gig | neg-off

Here is the syntax for changing the negotiation mode on individual ports.

Syntax: gig-default neg-full-auto | auto-gig | neg-off

USING THE WEB MANAGEMENT INTERFACE

To change the global default:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the [Advance](#) link to display the advanced System parameters panel.

3. Select one of the following values from the Gig Port Default field's pulldown menu:
 - Neg-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.
 - Auto-Gig – The port tries to perform a handshake with the other port to exchange capability information.
 - Neg-Full-Auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information).
4. Click Apply to save the changes to the device's running-config file.
5. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

To override the global negotiation mode for an individual port:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link for the port type you want to change (for example, [Ethernet](#), [ATM](#), and others) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select one of the following values from the Gig Port Default field's pulldown menu:
 - Default – The port uses the negotiation mode that was set at the global level.
 - Neg-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.
 - Auto-Gig – The port tries to perform a handshake with the other port to exchange capability information.
 - Neg-Full-Auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information).
7. Click Apply to save the changes to the device's running-config file.
8. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Limiting Broadcast, Multicast, or Unknown-Unicast Rates

Foundry devices can forward all traffic at wire speed. However, some third-party networking devices cannot handle high forwarding rates for broadcast, multicast, or unknown-unicast packets. You can limit the number of broadcast, multicast, or unknown-unicast packets a Foundry device forwards each second using the following methods.

The limits are individually configurable for broadcasts, multicasts, and unknown-unicasts. You can configure limits globally and on individual ports. The valid range is 1 – 4294967295 packets per second. If you specify 0, limiting is disabled. Limiting is disabled by default.

NOTE: By default, IP Multicast (including IGMP) is disabled. You can enable it using the **ip multicast passive | active** command. As long as IP Multicast is enabled (regardless of whether it is passive or active), no IP Multicast packets (not even IGMP packets) are limited. See "Configuring IP Multicast Traffic Reduction" on page 24-1.

Limiting Broadcasts

To limit the number of broadcast packets a Foundry device can forward each second, use the following CLI method.

USING THE CLI

To globally limit the number of broadcast packets a BigIron Layer 3 Switch forwards to 100,000 per second, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# broadcast limit 100000
BigIron(config)# write memory
```

To limit the number of broadcast packets sent on port 1/3 to 80,000, enter the following commands:

```
BigIron(config)# int ethernet 1/3
BigIron(config-if-1/3)# broadcast limit 80000
BigIron(config-if-1/3)# write memory
```

Syntax: broadcast limit <number>

NOTE: On BigIron MG8 and NetIron 40G, the broadcast limit is configured at the global level, but the value you enter applies to each management module (slot) installed on the device.

USING THE WEB MANAGEMENT INTERFACE

You cannot perform this procedure using the Web management interface.

Limiting Multicasts

To limit the number of multicast packets a Foundry device can forward each second, use the following CLI method.

USING THE CLI

To globally limit the number of multicast packets a BigIron Layer 3 Switch forwards to 120,000 per second, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# multicast limit 120000
BigIron(config)# write memory
```

To limit the number of multicast packets sent on port 3/6 to 55,000, enter the following commands:

```
BigIron(config)# int ethernet 3/6
BigIron(config-if-3/6)# multicast limit 55000
BigIron(config-if-3/6)# write memory
```

Syntax: multicast limit <number>

NOTE: On BigIron MG8 and NetIron 40G, the multicast limit is configured at the global level, but the value you enter applies to each management module (slot) installed on the device. .

USING THE WEB MANAGEMENT INTERFACE

You cannot perform this procedure using the Web management interface.

Limiting Unknown Unicasts

To limit the number unknown unicast packets a Foundry device can forward each second, use the following CLI method.

USING THE CLI

To globally limit the number of unknown unicast packets a BigIron Layer 3 Switch forwards to 110,000 per second, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# unknown-unicast limit 110000
BigIron(config)# write memory
```

To limit the number of unknown unicast packets sent on port 4/2 to 40,000, enter the following commands:

```
BigIron(config)# int ethernet 4/2
BigIron(config-if-4/2)# unknown-unicast limit 40000
BigIron(config-if-4/2)# write memory
```

Syntax: unknown-unicast limit <number>

NOTE: Only BigIron MG8 and NetIron 40G, the unknown-unicast limit is configured on the global level, but the value you enter applies to each management module (slot) installed on the device. .

USING THE WEB MANAGEMENT INTERFACE

You cannot perform this procedure using the Web management interface.

Configuring CLI Banners

Foundry devices can be configured to display a greeting message on users' terminals when they enter the Privileged EXEC CLI level or access the device through Telnet. In addition, a Foundry device can display a message on the Console when an incoming Telnet CLI session is detected.

Setting a Message of the Day Banner

You can configure the Foundry device to display a message on a user's terminal when he or she establishes a Telnet CLI session. For example, to display the message "Welcome to BigIron!" when a Telnet CLI session is established:

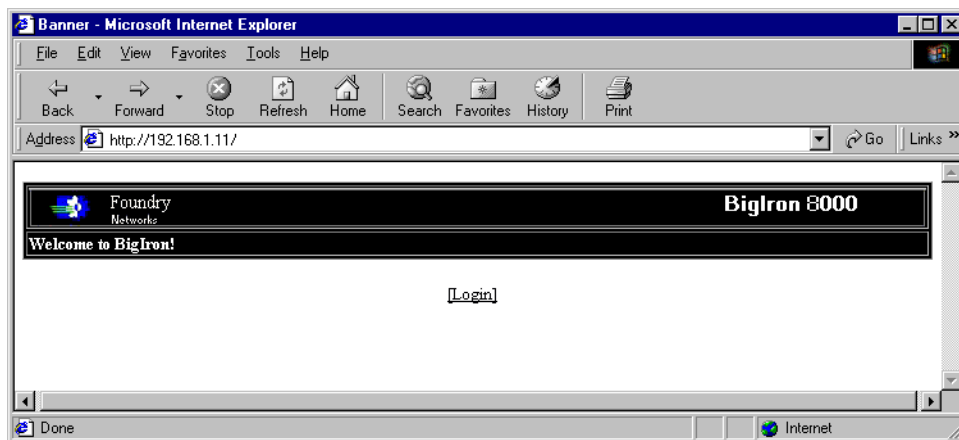
```
BigIron(config)# banner motd $ (Press Return)
Enter TEXT message, End with the character '$'.
Welcome to BigIron!! $
```

A delimiting character is established on the first line of the **banner motd** command. You begin and end the message with this delimiting character. The delimiting character can be any character except " (double-quotation mark) and cannot appear in the banner text. In this example, the delimiting character is \$ (dollar sign). The text in between the dollar signs is the contents of the banner. The banner text can be up to 2048 characters long and can consist of multiple lines. To remove the banner, enter the **no banner motd** command.

Syntax: [no] banner <delimiting-character> | [motd <delimiting-character>]

NOTE: The **banner <delimiting-character>** command is equivalent to the **banner motd <delimiting-character>** command.

When you access the Web management interface, the banner is displayed:



Setting a Privileged EXEC CLI Level Banner

You can configure the Foundry device to display a message when a user enters the Privileged EXEC CLI level. For example:

```
BigIron(config)# banner exec_mode # (Press Return)
Enter TEXT message, End with the character '#'.

```

```
You are entering Privileged EXEC level
Don't foul anything up! #
```

As with the **banner motd** command, you begin and end the message with a delimiting character; in this example, the delimiting character is # (pound sign). To remove the banner, enter the **no banner exec_mode** command.

Syntax: [no] banner exec_mode <delimiting-character>

Displaying a Message on the Console When an Incoming Telnet Session Is Detected

You can configure the Foundry device to display a message on the Console when a user establishes a Telnet session. This message indicates where the user is connecting from and displays a configurable text message.

For example:

```
BigIron(config)# banner incoming $ (Press Return)
Enter TEXT message, End with the character '$'.
Incoming Telnet Session!! $
```

When a user connects to the CLI using Telnet, the following message appears on the Console:

```
Telnet from 209.157.22.63
Incoming Telnet Session!!
```

Syntax: [no] banner incoming <delimiting-character>

To remove the banner, enter the **no banner incoming** command.

Configuring Terminal Display

You can configure and display the number of lines displayed on a terminal screen during the current CLI session.

The **terminal length** command allows you to determine how many lines will be displayed on the screen during the current CLI session. This command is useful when reading multiple lines of displayed information, especially those that do not fit on one screen.

To specify the maximum number of lines displayed on one page, enter a command such as the following:

```
BigIron(config)# terminal length 15
```

Syntax: terminal length <number-of-lines>

The <number-of-lines> parameter indicates the maximum number of lines that will be displayed on a full screen of text during the current session. If the displayed information requires more than one page, the terminal pauses. Pressing the space bar displays the next page.

The default for <number-of-lines> is 24. Entering a value of 0 prevents the terminal from pausing between multiple output pages:

Checking the Length of Terminal Displays

The **show terminal** command specifies the number of lines that will be displayed on the screen as specified by the **terminal length**, **page display**, and **skip-page-display** commands. It also shows if the **enable skip-page-display** command has been configured. The **enable skip-page-display** command allows you to use the **skip-page-display** to disable the configured page-display settings.

```
BigIron(config)# show terminal
Length: 24 lines
Page display mode (session): enabled
Page display mode (global): enabled
```

Configuring Basic Port Parameters

The procedures in this section describe how to configure the following port parameters:

- Name – see “Assigning a Port Name” on page 9-24
- Speed – see “Modifying Port Speed” on page 9-25
- Mode (half-duplex or full-duplex) – see “Modifying Port Mode” on page 9-26
- Status – see “Disabling or Re-Enabling a Port” on page 9-26
- Flow control – see “Disabling or Re-Enabling Flow Control” on page 9-27
- Gigabit negotiate mode – see “Changing the 802.3x Gigabit Negotiation Mode” on page 9-28
- QoS priority – see “Modifying Port Priority (QoS)” on page 9-30

NOTE: To modify Layer 2, Layer 3, or Layer 4 features on a port, see the appropriate section in this chapter or other chapters. For example, to modify Spanning Tree Protocol (STP) parameters for a port, see “Modifying STP Bridge and Port Parameters” on page 9-31.

NOTE: To configure trunk groups or dynamic link aggregation, see “Configuring Trunk Groups and Dynamic Link Aggregation” on page 17-1.

All Foundry ports are pre-configured with default values that allow the device to be fully operational at initial startup without any additional configuration. However, in some cases, changes to the port parameters may be necessary to adjust to attached devices or other network requirements.

The current port configuration for all ports is displayed when you select the [Port](#) link from the Configure tree. You can easily determine a port’s state by observing the color in the Port field.

- Red – indicates there is no link.
- Green – indicates the link is good.

This example shows the port states for a BigIron Layer 3 Switch that has not yet been connected to the rest of the network.

[Port Attribute][Port Statistic][Port Utilization][Relative Utilization]

Port Configuration

Port	Speed	QOS	Monitor	Mode	Lock Addr	Tag	STP	Fast STP	Fast Uplink STP	Flow Ctrl	Gig Default	Trunk	
1/1	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/2	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/3	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/4	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/5	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/6	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/7	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
1/8	1Gbps	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
3/1	Auto	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify
3/2	Auto	0	Disable	Full Duplex	Disable	Disable	Enable	Enable	Disable	Enable	Default	None	Modify

Click on the Copy or Modify button next to a row of port information to display a configuration panel for that port.

- Select Modify to change parameters for a port.
- Select Copy to apply a port's parameter settings to another port.

Here is an example of the Port configuration panel.

Port

Slot:4 Port:24 MAC:00-e0-52-f0-4f-00	
Name:	<input type="text"/>
Speed:	<input checked="" type="radio"/> 10/100 Auto <input type="radio"/> 10 Mbps <input type="radio"/> 100 Mbps
Mode:	<input checked="" type="radio"/> Full Duplex <input type="radio"/> Half Duplex
Status:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Flow Control:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Lock Address:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable MAC Address <input type="text"/>
Route Only:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
IEEE Tagging:	<input type="radio"/> Tag <input checked="" type="radio"/> Untag
QOS:	<input type="text" value="0"/>
Monitoring:	<input type="text" value="Disable"/>

[\[Show\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

NOTE: A slot option appears on the chassis port configuration sheet. Slot corresponds to a module slot number.

NOTE: The IEEE Tagging option appears only on the Port configuration sheet when tagging is enabled at the system level and a VLAN is defined on the system.

NOTE: The port speed option 1 Gbps is displayed only when a 1000BaseSX, 1000BaseLX, 1000BaseLH, or 1000BaseT Gigabit port or module is resident on the Layer 2 Switch or Layer 3 Switch. Additionally, only the full-duplex mode is visible. When a 10/100BaseTX Ethernet port or module is being configured, the options are 10/100 Auto, 10 Mbps, and 100 Mbps.

Assigning a Port Name

A port name can be assigned to help identify interfaces on the network. You can assign a port name to physical ports, virtual routing interfaces, and loopback interfaces.

USING THE CLI

To assign a name to a port:

```
BigIron(config)# interface e 2/8
BigIron(config-if-2/8)# port-name Marsha Markey
```

Syntax: port-name <text>

The <text> parameter is an alphanumeric string. The name can be up to 64 characters long on Stackable devices and up to 255 characters long on Chassis devices, including the TurboIron/8. The name can contain blanks. You do not need to use quotation marks around the string, even when it contains blanks.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.

3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, [Ethernet](#)) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Enter a name in the Name field.
7. Click Apply to save the changes to the device's running-config file.
8. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Modifying Port Speed

Each of the 10BaseT/100BaseTX ports is designed to auto-sense and auto-negotiate the speed and mode of the connected device. If the attached device does not support this operation, you can manually enter the port speed to operate at either 10 Mbps or 100 Mbps. The default value for 10BaseT/100BaseTX ports is 10/100 Auto-sense.

The 100BaseFX ports operate in the full-duplex mode at 100 Mbps only and cannot be modified.

The 1000BaseSX, 1000BaseLX, 1000BaseLH, and 1000BaseT ports operate in the full-duplex mode at one Gigabit only and cannot be modified.

NOTE: Modifying the port speed of a port that has a pre-configured rate limit policy may result in the inability to remove the port's rate limit policy.

USING THE CLI

To change the port speed of interface 1/8 from the default of 10/100 auto-sense to 10 Mbps operating at full-duplex, enter the following:

```
BigIron(config)# interface e 1/8
BigIron(config-if-1/8)# speed-duplex 10-full
```

Syntax: speed-duplex <value>

The <value> can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto

The default is auto.

USING THE WEB MANAGEMENT INTERFACE

To modify port speed:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, [Ethernet](#)) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Click next to Full Duplex if you want to change the mode to full-duplex only. (This applies only to 10/100 ports.)
7. Click Disable or Enable next to Auto Negotiate to enable or disable auto-negotiation.

- Click Apply to save the changes to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Modifying Port Mode

You can configure a port to accept either full-duplex (bi-directional) or half-duplex (uni-directional) traffic. This option is available only for 10/100 Mbps ports. The 100BaseFx, 1000BaseSx, and 1000BaseLx ports operate only at full-duplex.

USING THE CLI

Port duplex mode and port speed are modified by the same command.

To change the port speed of interface 1/8 from the default of 10/100 auto-sense to 10 Mbps operating at full-duplex, enter the following:

```
BigIron(config)# interface e 1/8
BigIron(config-if-1/8)# speed-duplex 10-full
```

Syntax: speed-duplex <value>

The <value> can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- auto

The default is auto.

USING THE WEB MANAGEMENT INTERFACE

To modify port mode:

- Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
- Click on the plus sign next to Configure in the tree view to display the configuration options.
- Click on the plus sign next to Port in the tree view to display the configuration options.
- Select the link to the port type you want (for example, [Ethernet](#)) to display the Port table.
- Click on the Modify button next to the row of information for the port you want to reconfigure.
- Click next to Full Duplex to select or de-select full duplex mode. Full-duplex mode is selected when the radio button (small circle) next to Full Duplex contains a black dot.
- Click Apply to save the changes to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Disabling or Re-Enabling a Port

The port can be made inactive (disable) or active (enable) by selecting the appropriate status option. The default value for a port is enabled.

USING THE CLI

To disable port 8 on module 1 of a Foundry Layer 3 Switch, enter the following:

```
BigIron(config)# interface e 1/8
BigIron(config-if-1/8)# disable
```

Syntax: disable

Syntax: enable

You also can disable or re-enable a virtual routing interface. To do so, enter commands such as the following:

```
BigIron(config)# interface ve v1
BigIron(config-vif-1)# disable
```

Syntax: disable

To re-enable a virtual routing interface, enter the **enable** command at the Interface configuration level. For example, to re-enable virtual routing interface v1, enter the following command:

```
BigIron(config-vif-1)# enable
```

Syntax: enable

USING THE WEB MANAGEMENT INTERFACE

To disable or enable a port:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, Ethernet) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select either Enable or Disable option next to the Status option.
7. Click Apply to save the changes to the device's running-config file.
8. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You cannot disable or re-enable a virtual routing interface using the Web management interface.

Disabling or Re-Enabling Flow Control

You can configure full-duplex ports on a system to operate with or without flow control (802.3x). Flow control is enabled by default.

USING THE CLI

To disable flow control on full-duplex ports on a system, enter the following:

```
BigIron(config)# no flow-control
```

To turn the feature back on:

```
BigIron(config)# flow-control
```

Syntax: [no] flow-control

USING THE WEB MANAGEMENT INTERFACE

To disable or enable flow control on full-duplex ports on a system:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, Ethernet) to display the Port table.

5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select either Enable or Disable next to Flow Control.
7. Click Apply to save the changes to the device's running-config file.
8. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Specifying Threshold Values for Flow Control

The 802.3x flow control specification provides a method for slowing traffic from a sender when a port is receiving more traffic than it can handle. Specifically, the receiving device can send out 802.3x PAUSE frames that request that the sender stop sending traffic for a period of time.

In software release 07.6.02 and higher, the Foundry device generates 802.3x PAUSE frames when the number of buffers available to a module's Buffer Manager (BM) drops below a threshold value. A module's BM can start running out of buffers when a port receives more traffic than it can handle. In addition, the device drops the lowest priority traffic when the number of available buffers drops below a second threshold. When the number of available buffers returns to a higher level, the device sends out another PAUSE frame that tells the sender to resume sending traffic normally. You can specify values for both thresholds, as well as the module where the thresholds are to take effect.

NOTE: To use this feature, 802.3x flow control must be enabled globally on the device. By default, 802.3x flow control is enabled on Foundry devices, but can be disabled with the **no flow-control** command.

To specify threshold values for flow control, enter the following command:

```
BigIron(config)# qd-flow sink 75 sunk 50 slot 1
```

Syntax: qd-flow sink <sinking-threshold> sunk <sunk-threshold> slot <slot>

The threshold values are percentages of the total number of buffers available to a module's Buffer Manager.

When the <sinking-threshold> is reached, the Foundry device sends out 802.3x PAUSE frames telling the sender to stop sending traffic for a period of time.

When the <sunk-threshold> is reached, the Foundry device drops traffic at the specified priority level.

The <slot> parameter specifies the location of the module where the thresholds are to take effect.

Changing the 802.3x Gigabit Negotiation Mode

On Chassis devices, the globally configured Gigabit negotiation mode for 802.3x flow control is the default mode for all Gigabit ports. You can override the globally configured default and set individual ports to the following:

- Negotiate-full-auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information). This is the default for Chassis devices (including the Turbolron/8).
- Auto-Gigabit – The port tries to perform a handshake with the other port to exchange capability information.
- Negotiation-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.

On Stackable devices, the default negotiation mode is negotiation-off but you can enable auto-Gigabit on individual Gigabit ports. Negotiate-full-auto is not supported on Stackable devices.

USING THE CLI

To change the mode for individual ports on a Chassis device, enter commands such as the following:

```
BigIron(config)# int ethernet 4/1 to 4/4
BigIron(config-mif-4/1-4/4)# gig-default auto-gig
```

This command overrides the global setting and sets the negotiation mode to auto-Gigabit for ports 4/1 – 4/4.

The following syntax applies to Chassis devices.

Syntax: gig-default neg-full-auto | auto-gig | neg-off

To change a Gigabit port on a Stackable device to auto-Gigabit, enter commands such as the following:

```
BigIron(config)# int ethernet 1/4
```

```
BigIron(config-if-1/4)# auto-gig
```

The following syntax applies to Stackable devices.

Syntax: [no] auto-gig

[USING THE WEB MANAGEMENT INTERFACE](#)

To override the global 802.3x negotiation mode for an Gigabit individual port on a Chassis device:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, [Ethernet](#)) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select one of the following values from the Gig Port Default field's pulldown menu:
 - Default – The port uses the negotiation mode that was set at the global level.
 - Neg-off – The port does not try to perform a handshake. Instead, the port uses configuration information manually configured by an administrator.
 - Auto-Gig – The port tries to perform a handshake with the other port to exchange capability information.
 - Neg-Full-Auto – The port first tries to perform a handshake with the other port to exchange capability information. If the other port does not respond to the handshake attempt, the port uses the manually configured configuration information (or the defaults if an administrator has not set the information).
7. Click Apply to save the changes to the device's running-config file.
8. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You also can access the dialog for saving configuration changes by clicking on Command in the tree view, then clicking on [Save to Flash](#).

To enable or disable auto-negotiate on a Gigabit port on a Stackable device:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, [Ethernet](#)) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select either Enable or Disable next to Auto Negotiate.
7. Click Apply to save the changes to the device's running-config file.
8. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You also can access the dialog for saving configuration changes by clicking on Command in the tree view, then clicking on [Save to Flash](#).

Modifying Port Priority (QoS)

You can give preference to the inbound traffic on specific ports by changing the Quality of Service (QoS) level on those ports. For information and procedures, see the “Configuring IronClad Quality of Service” chapter in the *Foundry Enterprise Configuration and Management Guide*.

Configuring Basic Layer 2 Parameters

The procedures in this section describe how to configure the following Layer 2 parameters. Note that some of these parameters apply only to Foundry Layer 2 Switches, not Layer 3 Switches.

- Spanning Tree Protocol (STP) – see “Enabling or Disabling the Spanning Tree Protocol (STP)” on page 9-30

NOTE: The procedures in this chapter describe how to configure standard STP. For information about Foundry’s IronClad STP, see “Configuring Spanning Tree Protocol (STP) and IronSpan Features” on page 10-1.

- Aging time for learned MAC address entries – see “Changing the MAC Age Time” on page 9-35
- Static, non-aging MAC address entries – see “Configuring Static MAC Entries” on page 9-35
- Port-based VLANs – see “Enabling Port-Based VLANs” on page 9-38
- MAC address filters – see “Defining MAC Address Filters” on page 9-39
- Broadcast and Multicast Filters – see “Defining Broadcast and Multicast Filters” on page 9-45
- Port locks – see “Locking a Port To Restrict Addresses” on page 9-47

Enabling or Disabling the Spanning Tree Protocol (STP)

The STP (IEEE 802.1d bridge protocol) is supported on all Foundry switches and routers. STP detects and eliminates logical loops in the network. STP also ensures that the least cost path is taken when multiple paths exist between ports or VLANs. If the selected path fails, STP searches for and then establishes an alternate path to prevent or limit retransmission of data.

STP must be enabled at the system level to allow assignment of this capability on the VLAN level. On Foundry Layer 2 Switches, STP is enabled by default. On Foundry Layer 3 Switches, STP is disabled by default.

NOTE: The procedures in this chapter describe how to configure basic STP parameters. For more information about Foundry’s IronClad STP, see “Configuring Spanning Tree Protocol (STP) and IronSpan Features” on page 10-1.

USING THE CLI

To enable STP for all ports on a Foundry Layer 2 Switch or Layer 3 Switch:

```
BigIron(config)# spanning tree
```

Syntax: [no] spanning-tree

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select Enable next to Spanning Tree.

NOTE: For information about the Single and Fast checkboxes, see “Single Spanning Tree (SSTP)” on page 10-62 and “Fast Uplink Span” on page 10-21.

3. Click Apply to save the changes to the device's running-config file.
4. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Modifying STP Bridge and Port Parameters

You can modify the following STP Parameters:

- Bridge parameters – forward delay, maximum age, hello time, and priority
- Port parameters – priority and path cost

STP Bridge Parameters

You can configure the following STP parameters:

- Forward Delay – The period of time a bridge will wait (the listen and learn period) before forwarding data packets. Possible values: 4 – 30 seconds. Default is 15.
- Maximum Age – The interval a bridge will wait for receipt of a hello packet before initiating a topology change. Possible values: 6 – 40 seconds. Default is 20.
- Hello Time – The interval of time between each configuration BPDU sent by the root bridge. Possible values: 1 – 10 seconds. Default is 2.
- Priority – A parameter used to identify the root bridge in a network. The bridge with the lowest value has the highest priority and is the root. Possible values: 0 – 65,535. Default is 32,768.

STP Port Parameters

Spanning Tree Protocol port parameters priority and path cost are preconfigured with default values. If the default parameters meet your network requirements, no other action is required.

You can configure the following STP port parameters:

- Port Priority – This parameter can be used to assign a higher (or lower) priority to a port. In the event that traffic is re-routed, this parameter gives the port forwarding preference over lower priority ports within a VLAN or on the Layer 2 Switch or Layer 3 Switch (when no VLANs are configured for the system). Ports are re-routed based on their priority. A higher numerical value means a lower priority; thus, the highest priority is 0. Possible values: 0 – 255. Default is 128.
- Path Cost – This parameter can be used to assign a higher or lower path cost to a port. This value can be used to bias traffic toward or away from a certain path during periods of rerouting. For example, if you wish to bias traffic away from a certain port, assign it a higher value than other ports within the VLAN or all other ports (when VLANs are not active on the Layer 2 Switch or Layer 3 Switch). Possible values are 1 – 65535. The default values are listed in Table 9.3.

Table 9.3: Default STP Port Path Costs

Port Type	Default Path Cost
10 Mbps	100
100 Mbps	19
Gigabit	4
OC-3c	200
OC-12c	80

Table 9.3: Default STP Port Path Costs (Continued)

Port Type	Default Path Cost
OC-48c	20

Notice that the path costs favor 100 Mbps and faster Ethernet ports over Packet over SONET (POS) ports. STP applies to POS ports only when they are configured for remote bridging to support Point-to-Point Protocol (PPP). (See “Configuring POS for Layer 2 Switching” on page 7-16.) The POS paths are remote (WAN) paths and the Ethernet paths are local paths. The default path costs therefore favor local paths over remote paths.

USING THE CLI

EXAMPLE:

Suppose you want to enable STP on a system in which no port-based VLANs are active and change the hello-time from the default value of 2 to 8 seconds. Additionally, suppose you want to change the path and priority costs for port 5 only. To do so, enter the following commands.

```
BigIron(config)# span hello-time 8
```

```
BigIron(config)# span ethernet 5 path-cost 15 priority 64
```

Here is the syntax for global STP parameters.

Syntax: span [forward-delay <value>] | [hello-time <value>] | [maximum-age <time>] | [priority <value>]

Here is the syntax for STP port parameters.

Syntax: span ethernet <portnum> path-cost <value> | priority <value>

USING THE WEB MANAGEMENT INTERFACE

To modify the STP parameters:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the STP link to display the STP bridge and port parameters.

- Click the Modify button in the STP bridge row to display the STP configuration panel, as shown in the following example.

STP	
VLAN ID:	<input type="text" value="1"/>
Bridge	
Forward Delay (Seconds):	<input type="text" value="15"/>
Maximum Age (Seconds):	<input type="text" value="20"/>
Hello Time (Seconds):	<input type="text" value="2"/>
Priority:	<input type="text" value="32768"/>
<input type="button" value="Apply"/>	
Port	
Priority:	<input type="text" value="128"/>
Path Cost:	<input type="text" value="0"/>
Slot:	<input type="text" value="1"/>
Port:	<input type="text" value="1"/>
<input type="button" value="Apply Port STP"/>	
<input type="button" value="Apply To All Ports"/>	

[\[Show\]](#)[\[Statistic\]](#)

[\[Home\]](#)[\[Site Map\]](#)[\[Logout\]](#)[\[Save\]](#)[\[Frame Enable\]](#)[\[Disable\]](#)[\[TELNET\]](#)

- Modify the bridge STP parameters to the values desired.
- Click Apply to save the changes to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

To modify the STP port parameters:

- Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
- Click on the plus sign next to Configure in the tree view to display the configuration options.
- Select the [STP](#) link to display the STP bridge and port parameters.
- If you are modifying the settings for a specific port, select the port (and slot if applicable) from the Port and Slot pulldown lists.
- Enter the desired changes to the priority and path cost fields.
- Click Apply STP Port to apply the changes to only the selected port or select Apply To All Ports to apply the changes to all the ports.

NOTE: If you want to save the priority and path costs of one port to all other ports on the Layer 2 Switch or Layer 3 Switch within a VLAN, you can click the Apply To All Ports button.

- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Enabling or Disabling Layer 2 Switching (Layer 3 Switches only)

By default, Foundry Layer 3 Switches support Layer 2 switching. These devices switch the routing protocols that are not supported on the devices. If IPX routing is not enabled, then IPX traffic also is switched. By default IPX routing is disabled. If you want to disable Layer 2 switching, you can do so globally or on individual ports.

NOTE: Layer 2 switching is disabled by default on the NetIron stackable and NetIron Internet Backbone router when shipped from the factory with software release 07.1.00 or higher.

NOTE: Make sure you really want to disable all Layer 2 switching operations before you use this option. Consult your reseller or Foundry Networks for information.

USING THE CLI

To globally disable Layer 2 switching on a Layer 3 Switch, enter commands such as the following:

```
BigIron(config)# route-only
BigIron(config)# exit
BigIron# write memory
BigIron# reload
```

To re-enable Layer 2 switching on a Layer 3 Switch, enter the following:

```
BigIron(config)# no route-only
BigIron(config)# exit
BigIron# write memory
BigIron# reload
```

Syntax: [no] route-only

To disable Layer 2 switching only on a specific interface, go to the Interface configuration level for that interface, then disable the feature. The following commands show how to disable Layer 2 switching on port 3/2:

```
BigIron(config)# interface ethernet 3/2
BigIron(config-if-3/2)# route-only
```

Syntax: [no] route-only

To re-enable Layer 2 switching, enter the command with “no”, as in the following example:

```
BigIron(config-if-3/2)# no route-only
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select Enable or Disable next to L2 Switching.
3. Click Apply to save the changes to the device's running-config file.
4. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

To disable or re-enable Layer 2 switching for an individual port:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the [Port](#) link to display the Port table.
4. Click on the Modify button next to the row of information for the port you want to reconfigure.
5. Select Disable or Enable next to Route Only.
6. Click Apply to save the changes to the device's running-config file.
7. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Changing the MAC Age Time

This parameter sets the aging period for ports on the device, defining how long a port address remains active in the address table. This parameter value can be 0 or a number from 67 – 65535 seconds. The zero value results in no address aging. The default value for this field is 300 (seconds).

USING THE CLI

To change the aging period for MAC addresses from the default value of 300 seconds to 600 seconds:

```
BigIron(config)# mac-age-time 600
```

Syntax: [no] mac-age-time <age-time>

The <age-time> can be 0 or a number from 67 – 65535.

USING THE WEB MANAGEMENT INTERFACE

To change the aging period for MAC addresses to 600 seconds:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the [Advance](#) link.
3. Enter the new value in the Switch Age Time field. You can enter a value from 0 – 65535.
4. Click Apply to save the changes to the device's running-config file.
5. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring Static MAC Entries

Static MAC addresses can be assigned to Foundry Layer 2 Switches and Layer 3 Switches.

NOTE: Foundry Layer 3 Switches also support the assignment of static IP Routes, static ARP, and static RARP entries. For details on configuring these types of static entries, see the “Configuring Static Routes” and “Creating Static ARP Entries” sections in the “Configuring IP” chapter of the *Foundry Enterprise Configuration and Management Guide*.

You can manually input the MAC address of a device to prevent it from being aged out of the system address table.

This option can be used to prevent traffic for a specific device, such as a server, from flooding the network with traffic when it is down. Additionally, the static MAC address entry is used to assign higher priorities to specific MAC addresses.

You can specify port priority (QoS) and VLAN membership (VLAN ID) for the MAC Address as well as specify device type of either router or host.

The default and maximum configurable MAC table sizes can differ depending on the device. To determine the default and maximum MAC table sizes for your device, display the system parameter values. See “Displaying and Modifying System Parameter Default Settings” on page 9-48.

For example, the MAC table on devices with JetCore modules running Service Provider IronWare software releases 09.1.00 and later can now have up to 1 million entries. Previously, the MAC table can hold only up to 64,000 entries. The MAC entries are stored in the CAM. The ability of the CAM to store up to a million MAC entries, depends on the following factors:

- The number of source MAC address being learned by the CAM.
- The number of destination MAC addresses being forwarded by the CAM
- The distribution of the MAC entries across ports. For example, if one port is learning all the source MAC addresses, the available of the CAM for that port will be depleted.

Also, a large number of MAC entries in the MAC table could increase CPU utilization. To alleviate the load on the CPU, use this feature with the Control Plane Security option.

EXAMPLE:

To add a static entry for a server with a MAC address of 1145.5563.67FF and a priority of 7 to port 2 of module 1 of a BigIron Layer 3 Switch:

USING THE CLI

```
BigIron(config)# static-mac-address 1145.5563.67FF e 1/2 priority 7
```

Here is the syntax for Chassis devices.

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [to <portnum> ethernet <portnum>]
[priority <number>] [host-type | router-type | fixed-host]

Here is the syntax for Stackable devices.

Syntax: static-mac-address <mac-addr> ethernet <port-num> [priority <number>]
[host-type | router-type]

NOTE: On devices running Service Provider IronWare software releases 09.1.00 and later, you can configure up to 255 static MAC addresses on a device

The priority can be 0 – 7 (0 is lowest priority and 7 is highest priority) for Chassis devices, and either normal-priority or high-priority for Stackable devices.

The default priority is 0 or normal-priority. The default type is host-type.

NOTE: The location of the **static-mac-address** command in the CLI depends on whether you configure port-based VLANs on the device. If the device does not have more than one port-based VLAN (VLAN 1, which is the default VLAN that contains all the ports), the **static-mac-address** command is at the global CONFIG level of the CLI. If the device has more than one port-based VLAN, then the **static-mac-address** command is not available at the global CONFIG level. In this case, the command is available at the configuration level for each port-based VLAN.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the list of configuration options.
3. Select the [Static Station](#) link.
 - If the system already contains static MAC addresses and you are adding a new static MAC address, click on the [Add Static Station](#) link to display the Static Station Table configuration panel, as shown in the following example.
 - If you are modifying an existing static MAC address, click on the Modify button to the right of the row describing the static MAC address to display the Static Station Table configuration panel, as shown in the

following example.

Static Station Table

MAC Address:	ab-cd-ab-cd-ab-cd	
VLAN ID:	1	
Slot:	1	Port: 1
QoS:	0	

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Disable Frame\]](#)
[\[TELNET\]](#)

4. Enter or edit the MAC address, if needed. Specify the address in the following format: xx-xx-xx-xx-xx-xx.
5. Change the VLAN number if needed by editing the value in the VLAN ID field.
6. Select the port number from the Slot (for Chassis devices) and Port pulldown lists.
7. Select a QoS level from 0 – 7 from the QoS field's pulldown menu. For information about QoS, see the "Configuring IronClad Quality of Service" chapter in the *Foundry Enterprise Configuration and Management Guide*.
8. Click the Add button (to add a new static MAC entry) or the Modify button (if you are modifying an existing entry) to save the change to the device's running-config file.
9. Click the Apply button to save the change to the device's running-config file.
10. Select the [Save](#) link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring Static ARP Entries

Foundry recommends that you configure a static ARP entry to match the static MAC entry. In fact, the software automatically creates a static MAC entry when you create a static ARP entry.

NOTE: When a static MAC entry has a corresponding static ARP entry, you cannot delete the static MAC entry unless you first delete the static ARP entry.

To create a static ARP entry for a static MAC entry, enter a command such as the following:

```
BigIron(config)# arp 1 192.53.4.2 aaaa.bbbb.cccc ethernet 1
```

The **arp** command allows you to specify only one port number. To create a static ARP entry for a static MAC entry that is associated with multiple ports, specify the first (lowest-numbered) port associated with the static MAC entry.

Syntax: [no] arp <num> <ip-addr> <mac-addr> ethernet <portnum>

The <num> parameter specifies the entry number.

Configuring a Local MAC Address for Layer 2 Management Traffic

NOTE: This feature is supported on the FES running releases 03.4.00 and later.

By default, Foundry Layer 2 devices use the MAC address of the first port as the MAC address for Layer 2 management traffic. For example, when the Foundry device receives an ARP request for its management IP address, it responds with the first port's MAC address. This may cause problems in some configurations where the Foundry device uses the same MAC address for management traffic as for switched traffic.

Starting with software release 03.4.00, you can configure the Foundry device to use a different MAC address for Layer 2 management traffic than for switched traffic. When you issue the **use-local-management-mac**, the Foundry device will change a local bit in the first port's MAC address and will use this MAC address for management traffic. Switched traffic will continue to use the first port's MAC address without the local bit setting.

EXAMPLE:

```
FES Switch(config)# use-local-management-mac
FES Switch(config)# write memory
FES Switch(config)# end
FES Switch# reload
```

Syntax: [no] use-local-management-mac

NOTE: You must save the configuration and reload the software to place the change into effect.

Enabling Port-Based VLANs

Port and protocol VLANs must first be enabled at the system (global) level before they can be configured at the VLAN level. For details on configuring VLANs, refer to “Configuring Virtual LANs (VLANs)” on page 22-1.

USING THE CLI

When using the CLI, port and protocol-based VLANs are created by entering one of the following commands at the global CONFIG level of the CLI.

To create a port-based VLAN, enter commands such as the following:

```
BigIron(config)# vlan 222 by port
BigIron(config)# vlan 222 name Mktg
```

Syntax: vlan <num> by port

Syntax: vlan <num> name <string>

The <num> parameter specifies the VLAN ID. The valid range for VLAN IDs starts at 1 on all systems but the upper limit of the range differs depending on the device. In addition, you can change the upper limit on some devices using the **vlan max-vlans...** command. See the *Foundry Switch and Router Command Line Interface Reference*.

The <string> parameter is the VLAN name and can be a string up to 16 characters. You can use blank spaces in the name if you enclose the name in double quotes (for example, “Product Marketing”).

NOTE: The second command is optional and also creates the VLAN if the VLAN does not already exist. You can enter the first command after you enter the second command if you first exit to the global CONFIG level of the CLI.

USING THE WEB MANAGEMENT INTERFACE

To enable port-based VLANs on the Layer 2 Switch or Layer 3 Switch:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click the box for Port, next to Policy Based VLANs to enable port-based VLANs.
3. Click Apply to save the changes to the device's running-config file.
4. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Assigning IEEE 802.1q Tagging to a Port

When a port is tagged, it allows communication among the different VLANs to which it is assigned. A common use for this might be to place an email server that multiple groups may need access to on a tagged port, which in turn, is resident in all VLANs that need access to the server.

NOTE: Tagging is disabled by default on individual ports.

NOTE: Tagging does not apply to the default VLAN.

For details on configuring port-based VLANs, refer to “Configuring Virtual LANs (VLANs)” on page 22-1.

USING THE CLI

When using the CLI, ports are defined as either tagged or untagged at the VLAN level.

EXAMPLE:

Suppose you want to make port 5 on module 1 a member of port-based VLAN 4, a tagged port. To do so, enter the following:

```
BigIron(config)# vlan 4
BigIron(config-vlan-4)# tagged e 1/5
```

Syntax: tagged ethernet <portnum> [to <portnum> [ethernet <portnum>]]

USING THE WEB MANAGEMENT INTERFACE

To apply 802.1q tagging to a port:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, Ethernet) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select Enable next to IEEE Tagging.

NOTE: This option appears only if you are modifying a port that is a member of a port-based VLAN other than the default VLAN. Tagging does not apply to ports that are not in a port-based VLAN and does not apply to the default VLAN.

7. Click Apply to save the changes to the device's running-config file.
8. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Defining MAC Address Filters

MAC layer filtering enables you to build access lists based on MAC layer headers in the Ethernet/IEEE 802.3 frame. You can filter on the source and destination MAC addresses as well as other information such as the EtherType, LLC1 DSAP or SSAP numbers, and a SNAP EtherType. The filters apply to incoming traffic only.

NOTE: On the FastIron Edge Switch, you can filter on source and destination MAC addresses only.

NOTE: MAC filters do not block management access to the Foundry device. For example, if you apply a filter to block a specific host, the filter blocks switch traffic from the host but does not prevent the host from establishing a management connection to the device through Telnet. To block management access, use an Access Control List (ACL). See the “IP Access Control Lists (ACLs)” chapter of the *Foundry Enterprise Configuration and Management Guide*.

NOTE: You cannot use Layer 2 filters to filter Layer 4 information. To filter Layer 4 information, use IP access policies. See the “Policies and Filters” appendix in the *Foundry Enterprise Configuration and Management Guide*.

You configure MAC filters globally, then apply them to individual interfaces. To apply MAC filters to an interface, you add the filters to that interface’s MAC filter group.

NOTE: In software release 07.6.03, you can apply MAC filters to virtual routing interfaces. For more information, see “Configuring MAC Address Filters for Virtual Routing Ports” on page 9-44.

The device takes the action associated with the first matching filter. If the packet does not match any of the filters in the access list, the default action is to drop the packet. If you want the system to permit traffic by default, you must specifically indicate this by making the last entry in the access list a permit filter. Here is an example:

mac filter <last-index-number> **permit any any**

For Layer 3 Switches, the MAC filter is applied only to those inbound packets that are to be switched. This includes those ports associated with a virtual routing interface. However, the filter is not applied to the virtual routing interface. It is applied to the physical port.

NOTE: Inbound traffic on a port to which a Layer 2 MAC filter is assigned is sent to the CPU for processing.

NOTE: Use MAC Layer 2 filters only for switched traffic. If a routing protocol (for example, IP or IPX) is configured on an interface, a MAC filter defined on that interface is not applied to inbound packets. If you want to filter inbound route traffic, configure a route filter.

When you create a MAC filter, it takes effect immediately. You do not need to reset the system. However, you do need to save the configuration to flash memory to retain the filters across system resets.

For complete MAC filter examples, see the *Foundry Switch and Router Command Line Interface Reference*.

Configuring MAC Address Filters for Physical Ports

NOTE: In software releases 07.6.03 and later, you can apply MAC filters to virtual routing interfaces. For more information, see “Configuring MAC Address Filters for Virtual Routing Ports” on page 9-44.

To define a MAC filter, use one of the following methods.

USING THE CLI

To configure and apply a MAC filter, enter commands such as the following:

```
BigIron(config)# mac filter 1 deny 3565.3475.3676 ffff.0000.0000 any etype eq 806
BigIron(config)# mac filter 1 permit any any
BigIron(config)# int e 1/1
BigIron(config-if-1/1)# mac filter-group 1
```

These commands configure a filter to deny ARP traffic with a source MAC address that begins with “3565” to any destination. The second filter permits all traffic that is not denied by another filter.

NOTE: Once you apply a MAC filter to a port, the device drops all Layer 2 traffic on the port that does not match a MAC permit filter on the port.

Syntax: `mac filter <filter-num> permit | deny any | <H.H.H> any | <H.H.H> etype | llc | snap <operator> <frame-type>`

The <filter-num> can be a number from 1 – 128.

The **permit | deny** argument determines the action the software takes when a match occurs.

The <src-mac> <mask> | **any** parameter specifies the source MAC address. You can enter a specific address value and a comparison mask or the keyword **any** to filter on all MAC addresses. Specify the mask using f’s (ones) and zeros. For example, to match on the first two bytes of the address aabb.ccdd.eeff, use the mask

ffff.0000.0000. In this case, the filter matches on all MAC addresses that contain "aabb" as the first two bytes. The filter accepts any value for the remaining bytes of the MAC address. If you specify **any**, do not specify a mask. In this case, the filter matches on all MAC addresses.

The <dest-mac> <mask> | **any** parameter specifies the destination MAC address. The syntax rules are the same as those for the <src-mac> <mask> | **any** parameter.

Use the **etype** | **llc** | **snap** argument if you want to filter on information beyond the source and destination address. The MAC filter allows for you to filter on the following encapsulation types:

- **etype** (Ethertype) – a two byte field indicating the protocol type of the frame. This can range from 0x0600 to 0xFFFF.
- **llc** (IEEE 802.3 LLC1 SSAP and DSAP) – a two byte sequence providing similar function as the EtherType but for an IEEE 802.3 frame.
- **snap** (IEEE 802.3 LLC1 SNAP) – a specific LLC1 type packet.

To determine which type of frame is used on your network, use a protocol analyzer. If byte 12 of an Ethernet packet is equal to or greater than 0600 (hex), it is an Ethernet framed packet. Any number below this indicates an IEEE 802.3 frame (byte 12 will now indicate the length of the data field). Some well-known Ethernet types are 0800 (TCP/IP), 0600 (XNS), and 8137 (Novell Netware). Refer to RFC 1042 for a complete listing of EtherTypes.

For IEEE 802.3 frame, you can further distinguish the SSAP and DSAP of LLC header. Some well-known SAPs include: FE (OSI), F0 (NetBIOS), 42 (Spanning Tree BPDU), and AA (SNAP). Usually the DSAP and SSAP are the same.

NOTE: You must type in both bytes, otherwise the software will fill the field, left justified with a 00. Refer to RFC 1042 for a complete listing of SAP numbers.

SNAP is defined as an IEEE 802.3 frame with the SSAP, DSAP, and control field set to AA, AA, and 03. Immediately following these is a five-byte SNAP header. The first three bytes in this header are not used by the MAC filters. However, the next two bytes usually are set to the EtherType, so you can define the EtherType inside the SNAP header that you want to filter on.

The **eq** | **gt** | **lt** | **neq** argument specifies the possible operator: eq (equal), gt (greater than), lt (less than) and neq (not equal).

The <frame-type> argument is a hexadecimal number for the frame type. For example, the hex number for ARP is 806.

To globally enable logging for filtered packets, enter the following command:

```
BigIron(config)# mac filter log-enable
```

Syntax: mac filter log-enable

To enable logging for filtered packets on a specific port, enter the following commands:

```
BigIron(config)# int e 1/1
BigIron(config-if-1/1)# mac filter-group log-enable
```

Syntax: mac filter-group log-enable

To assign MAC filter 1 to interface port 1 on slot 1, enter the following commands:

```
BigIron(config)# int e 1/1
BigIron(config-if-1/1)# mac filter-group 1
```

Syntax: mac filter-group <filter-list>

NOTE: The filters must be applied as a group. For example, if you want to apply four filters to an interface, they must all appear on the same command line.

NOTE: You cannot add or remove individual filters in the group. To add or remove a filter on an interface, apply the filter group again containing all the filters you want to apply to the port.

NOTE: If you apply a filter group to a port that already has a filter group applied, the older filter group is replaced by the new filter group.

USING THE WEB MANAGEMENT INTERFACE

To define a MAC filter:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to System in the tree view to display the system configuration options.
4. Select the MAC Filter link.
 - If the device does not have any MAC filters configured, the MAC Filter configuration panel is displayed, as shown in the following example.
 - If a MAC filter is already configured and you are adding a new one, click on the Add MAC Filter link to display the MAC Filter configuration panel, as shown in the following example.
 - If you are modifying an existing MAC filter, click on the Modify button to the right of the row describing the filter to display the MAC Filter configuration panel, as shown in the following example.

MAC Filter

ID:	<input type="text" value="1"/>
Action:	<input checked="" type="radio"/> Deny <input type="radio"/> Permit
Source Address:	<input type="text" value="12-34-56-78-9a-bc"/>
Source Mask:	<input type="text" value="ff-ff-ff-00-00-00"/>
Destination Address:	<input type="text" value="ab-cd-ab-cd-ab-cd"/>
Destination Mask:	<input type="text" value="ff-ff-ff-ff-ff-ff"/>
Frame Type:	<input type="text" value="llc"/> ▾
Operator:	<input type="text" value="Equal"/> ▾
Protocol:	<input type="text" value="0000"/> <input type="button" value="System Define"/>

[Show][Filter Group]

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

5. Edit the value in the ID field if you want to assign the filter a different ID. The software automatically increments this field each time you add a MAC filter.
6. Select the filter action by selecting Permit or Deny next to Action.
7. Enter the source MAC address in the Source Address field. Separate the bytes in the address with dashes.
8. Enter the comparison mask for the source address in the Source Mask field. The mask consists of “f”s and “0”s or the word “any”.
 - An “f” indicates a significant bit. The software checks the indicated bit in each packet’s source MAC address.

- A “0” indicates an insignificant bit. The software does not care what value is in the bit position.
 - “any” matches all bits and is equivalent to entering “ff-ff-ff-ff-ff”.
9. Enter the destination MAC address in the Destination Address field. Separate the bytes in the address with dashes.
 10. Enter the comparison mask for the destination address in the Destination Mask field.
 11. Select the frame type from the Frame Type field’s pulldown menu.
 12. Select an operator from the Operator field’s pulldown menu to filter by protocol type.
 13. Enter a protocol in the Protocol field.
 14. Click the Add button to save the filter to the device’s running-config file. The filter is now configured in the software but has not yet been applied to a port.
 15. Select the [Filter Group](#) link.
 - If the device does not have any MAC filter groups configured, the Filter Group configuration panel is displayed, as shown in the following example.
 - If a MAC filter group is already configured and you are adding a new one, click on the Show link to display the MAC Filter Group list. Then click on the [Add MAC Filter Group](#) link to display the Filter Group configuration panel, as shown in the following example.
 - If you are modifying an existing MAC filter group, click on the Modify button to the right of the row describing the filter group to display the Filter Group configuration panel, as shown in the following example.

Filter Group

Slot:	1 ▾	Port:	1 ▾
Filter ID List:	1 2 3 1024		

[\[Show\]\[MAC Filter\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

16. Select the port (and slot, if applicable) for which you are configuring the filter group. You can configure one MAC filter group on each port.
17. Enter the filter numbers in the Filter ID List field. Separate each filter number from the next one by a single space. The software applies the filters in the order you list them, from left to right. When a packet matches a filter, the software stops comparing the packet against the filter list and applies the action specified in the matching filter.

NOTE: The filters must be applied as a group. For example, if you want to apply four filters to an interface, they must all appear on the same command line.

NOTE: You cannot add or remove individual filters in the group. To add or remove a filter on an interface, apply the filter group again containing all the filters you want to apply to the port.

NOTE: If you apply a filter group to a port that already has a filter group applied, the older filter group is replaced by the new filter group.

18. Click the Add button to save the filter to the device’s running-config file.

19. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring MAC Address Filters for Virtual Routing Ports

Software release 07.6.03 allows you to apply MAC filters to virtual routing interfaces; however, MAC filters used on a virtual routing interface can only deny packets. Permit is not available. Packets are denied based on their source MAC address. The Layer 3 Switch will drop any Layer 2 or Layer 3 packet that originated from the specified source MAC address.

NOTE: No etype arguments will be checked.

To apply a MAC filter on a virtual routing interface using the CLI, first create a filter group that denies specific source MAC addresses using the **mac filter-group** command. (Refer to the *Foundry Switch and Router Installation and Basic Configuration Guide* for details.) Then use the **mac deny-src-mac-filter-grp...** command to apply them to virtual routing interfaces. Enter commands such as the following:

```
BigIron(config)# interface ve 2
BigIron(config-vif-2)# mac filter 1 deny 00a0.cc77.a18d ffff.ffff.ffff any
BigIron(config-vif-2)# mac filter 2 deny 0010.2222.3333 ffff.ffff.ffff any
BigIron(config-vif-2)# mac deny-src-mac-filter-grp 1 2
```

Syntax: [no] mac deny-src-mac-filter-group <number>

<number> is the number of the ID of the filter that you've defined. You can enter up to eight filter IDs.

Enabling Logging of Packets Denied by MAC Filters

You can configure the Foundry device to generate Syslog entries and SNMP traps for packets that are denied by Layer 2 MAC filters. You can enable logging of denied packets on a global basis or an individual port basis.

The first time an entry in a MAC filter denies a packet and logging is enabled for that entry, the software generates a Syslog message and an SNMP trap. Messages for packets denied by MAC filters are at the warning level of the Syslog.

When the first Syslog entry for a packet denied by a MAC filter is generated, the software starts a five-minute MAC filter timer. After this, the software sends Syslog messages every five minutes. The messages list the number of packets denied by each MAC filter during the previous five-minute interval. If a MAC filter does not deny any packets during the five-minute interval, the software does not generate a Syslog entry for that MAC filter.

NOTE: For a MAC filter to be eligible to generate a Syslog entry for denied packets, logging must be enabled for the filter. The Syslog contains entries only for the MAC filters that deny packets and have logging enabled.

When the software places the first entry in the log, the software also starts the five-minute timer for subsequent log entries. Thus, five minutes after the first log entry, the software generates another log entry and SNMP trap for denied packets.

USING THE CLI

To configure Layer 2 MAC filter logging globally, enter the following CLI commands at the global CONFIG level:

```
BigIron(config)# mac filter log-enable
BigIron(config)# write memory
```

Syntax: [no] mac filter log-enable

To configure Layer 2 MAC filter logging for MAC filters applied to ports 1/1 and 3/3, enter the following CLI commands:

```
BigIron(config)# int ethernet 1/1
BigIron(config-if-1/1)# mac filter-group log-enable
BigIron(config-if-1/1)# int ethernet 3/3
BigIron(config-if-3/3)# mac filter-group log-enable
BigIron(config-if-3/3)# write memory
```

Syntax: [no] mac filter-group log-enable

USING THE WEB MANAGEMENT INTERFACE

You cannot configure a Layer 2 MAC filter to generate Syslog entries and SNMP traps for denied packets using the Web management interface.

Defining Broadcast and Multicast Filters

You can filter Layer 2 broadcast and multicast packets on specific ports.

- Layer 2 broadcast packets have the value "FFFFFFFFFFFF" (all ones) in the destination MAC address field. You can configure broadcast filters for all types of IP packets or for UDP packets.
- Layer 2 multicast packets have a multicast address in the destination MAC address field. You can configure multicast filters to filter on all MAC addresses or a specific multicast address.

You can configure up to eight of each type of filter.

To configure a Layer 2 broadcast or multicast filter, you define the filter globally to either filter out all types of broadcasts or to filter out only IP UDP broadcasts. After configuring a broadcast or multicast filter, you apply it to specific ports. Broadcast and multicast filters apply only to outbound traffic.

When defining the filter, you can specify a port-based VLAN ID. If a port is a member of more than one VLAN and is a tagged port, specifying a VLAN ID causes the filter to be applied only to traffic for the specified VLAN on the tagged ports to which you apply the filter. Otherwise, the filter applies to all the VLANs of which the port is a member.

The filters are applied in numerical order, beginning with filter number 1. As soon as the software finds a matching filter for a given packet, the filtering process stops for that packet. For example, if you configure filter 1 to filter all broadcast traffic and filter 2 to filter only IP UDP traffic, filter 1 will always be true for any broadcast packet, and thus the software will never consult filter 2 for ports that you configure to use filter 1.

Configuring a Layer 2 Broadcast Filter

To configure a broadcast filter, you must have access to the CONFIG level of the CLI. You can configure up to eight broadcast filters on a device.

Syntax: [no] broadcast filter <filter-id> any | ip udp [vlan <vlan-id>]

Syntax: [no] exclude-ports ethernet <portnum> to <portnum>

Or

Syntax: [no] exclude-ports ethernet <portnum> ethernet <portnum>

The **exclude-ports** command specifies the ports to which the filter applies.

The <filter-id> specifies the filter number and can a number from 1 – 8. The software applies the filters in ascending numerical order. As soon as a match is found, the software takes the action specified by the filter (block the broadcast) does not compare the packet against additional broadcast filters.

You can specify **any** or **ip udp** as the type of broadcast traffic to filter. The **any** parameter prevents all broadcast traffic from being sent on the specified ports. The **ip udp** parameter prevents all IP UDP broadcasts from being sent on the specified ports but allows other types of broadcast traffic.

If you specify a port-based VLAN ID, the filter applies only to the broadcast domain of the specified VLAN, not to all broadcast domains (VLANs) on the device.

As soon as you press Enter after entering the command, the CLI changes to the configuration level for the filter you are configuring. You specify the ports to which the filter applies at the filter's configuration level.

NOTE: This is the same command syntax as that used for configuring port-based VLANs. Use the first command for adding a range of ports. Use the second command for adding separate ports (not in a range). You also can combine the syntax. For example, you can enter **exclude-ports ethernet 1/4 ethernet 2/6 to 2/9**.

Configuration Examples

To configure a Layer 2 broadcast filter to filter all types of broadcasts, then apply the filter to ports 1/1, 1/2, and 1/3, enter the following commands:

```
BigIron(config)# broadcast filter 1 any
BigIron(config-bcast-filter-id-1)# exclude-ports ethernet 1/1 to 1/3
BigIron(config-bcast-filter-id-1)# write memory
```

To configure two filters, one to filter IP UDP traffic on ports 1/1 – 1/4, and the other to filter all broadcast traffic on port 4/6, enter the following commands:

```
BigIron(config)# broadcast filter 2 ip udp
BigIron(config-bcast-filter-id-2)# exclude-ports ethernet 1/1 to 1/4
BigIron(config-bcast-filter-id-2)# exit
BigIron(config)# broadcast filter 3 any
BigIron(config-bcast-filter-id-3)# exclude-ports ethernet 4/6
BigIron(config-bcast-filter-id-3)# write memory
```

To configure an IP UDP broadcast filter and apply that applies only to port-based VLAN 10, then apply the filter to two ports within the VLAN, enter the following commands:

```
BigIron(config)# broadcast filter 4 ip udp vlan 10
BigIron(config-bcast-filter-id-4)# exclude-ports eth 1/1 eth 1/3
BigIron(config-bcast-filter-id-4)# write memory
```

Configuring a Layer 2 Multicast Filter

To configure a multicast filter, you must have access to the CONFIG level of the CLI. You can configure up to eight multicast filters on a device.

Syntax: [no] multicast filter <filter-id> any | ip udp mac <multicast-address> | any [mask <mask>]
[vlan <vlan-id>]

The parameter values are the same as the for the broadcast filter command. In addition, the multicast filter command requires the **mac <multicast-address> | any** parameter, which specifies the multicast address. Enter **mac any** to filter on all multicast addresses.

Enter **mac** followed by a specific multicast address to filter only on that multicast address. To filter on a range of multicast addresses, use the **mask <mask>** parameter. For example, to filter on multicast groups 0100.5e00.5200 – 0100.5e00.52ff, use **mask ffff.ffff.ff00**. The default mask matches all bits (is all Fs). You can leave the mask off if you want the filter to match on all bits in the multicast address.

Configuration Examples

To configure a Layer 2 multicast filter to filter all multicast groups, then apply the filter to ports 2/4, 2/5, and 2/8, enter the following commands:

```
BigIron(config)# multicast filter 1 any
BigIron(config-mcast-filter-id-1)# exclude-ports ethernet 2/4 to 2/5 ethernet 2/8
BigIron(config-mcast-filter-id-1)# write memory
```

To configure a multicast filter to block all multicast traffic destined for multicast addresses 0100.5e00.5200 – 0100.5e00.52ff on port 4/8, enter the following commands:

```
BigIron(config)# multicast filter 2 any 0100.5e00.5200 mask ffff.ffff.ff00
BigIron(config-mcast-filter-id-2)# exclude-ports ethernet 4/8
BigIron(config-mcast-filter-id-2)# write memory
```

The software calculates the range by combining the mask with the multicast address. In this example, all but the last eight bits in the mask are “significant bits” (ones). The last eight bits are zeros and thus match on any value. Each “f” or “0” is four bits.

Locking a Port To Restrict Addresses

Address-lock filters allow you to limit the number of devices that have access to a specific port. Access violations are reported as SNMP traps. By default this feature is disabled. A maximum of 2048 entries can be specified for access. The default address count is eight.

NOTE: In release 07.6.01, a more robust version of this feature was introduced. See “Using the MAC Port Security Feature” in the *Foundry Security Guide*.

USING THE CLI

EXAMPLE:

To enable address locking for port 2/1 and place a limit of 15 entries:

```
BigIron(config)# lock e 2/1 addr 15
```

Syntax: lock-address ethernet <portnum> [addr-count <num>]

USING THE WEB MANAGEMENT INTERFACE

To enable address locking on a port:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Click on the plus sign next to Port in the tree view to display the configuration options.
4. Select the link to the port type you want (for example, Ethernet) to display the Port table.
5. Click on the Modify button next to the row of information for the port you want to reconfigure.
6. Select Enable next to Lock Address.
7. Enter the maximum number of MAC addresses you want the device to learn on the port in the MAC Address field.
8. Click Apply to save the changes to the device’s running-config file.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Enabling or Disabling Routing Protocols

Foundry Layer 3 Switches support the following protocols:

- AppleTalk
- BGP4
- DVMRP
- FSRP
- IP
- IPX
- OSPF
- PIM
- RIP

- VRRP
- VRRPE

By default, IP routing is enabled on Layer 3 Switches. All other protocols are disabled, so you must enable them to configure and use them.

NOTE: The following protocols require a system reset before the protocol will be active on the system: PIM, DVMRP, RIP, FSRP, and IPX. To reset a system, select the [Reload](#) link (Web) or enter the **reload** command at the privileged level of the CLI.

USING THE CLI

To enable a protocol on a Foundry Layer 3 Switch, enter **router** at the global CONFIG level, followed by the protocol to be enabled. The following example shows how to enable OSPF:

```
BigIron(config)# router ospf
BigIron(config)# end
BigIron# write memory
BigIron# reload
```

Syntax: router appletalk | bgp | dvmrp | fsrp | ipx | ospf | pim | rip | vrrp | vrrpe

USING THE WEB MANAGEMENT INTERFACE

To enable protocols on a Layer 3 Switch:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the Enable option next to the protocol(s) to be enabled.

NOTE: If you are enabling BGP4, you must also specify the local AS number in the Local AS field.

NOTE: Do not enable both FSRP and VRRP. Foundry Networks recommends that you use only one of these router redundancy protocols on a Layer 3 Switch.

3. Click Apply to save the changes to the device's running-config file.
4. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

NOTE: You also can access the dialog for saving configuration changes by clicking on Command in the tree view, then clicking on [Save to Flash](#).

If you enable PIM, DVMRP, RIP, FSRP, or IPX, you must reload the software to place the change into effect.

1. Click on the plus sign next to Command in the tree view to list the command options.
2. Select the [Reload](#) link and select Yes when the Web management interface asks you whether you really want to reload the software.

Displaying and Modifying System Parameter Default Settings

Foundry devices have default table sizes for the following parameters. The table sizes determine the maximum number of entries the tables can hold. You can adjust individual table sizes to accommodate your configuration needs.

- MAC address entries
- Layer 2 Port VLANs supported on a system

- Layer 3 Protocol VLANs supported on a system
- Layer 4 sessions supported
- IP cache size
- ARP entries
- IP routes
- IP route filters
- IP subnets per port and per device
- Static routes
- IGMP
- DVMRP routes
- IPX/SAP entries
- IPX/RIP entries
- IPX/SAP filters
- IPX/RIP filters
- IPX forwarding filters
- AppleTalk routes
- AppleTalk zones

The tables you can configure as well the defaults and valid ranges for each table differ depending on the Foundry device you are configuring.

NOTE: If you increase the number of subnet addresses you can configure on each port to a higher amount, you might also need to increase the total number of subnets that you can configure on the device.

To display and configure the adjustable tables on a device, use one of the following methods.

NOTE: Changing the table size for a parameter reconfigures the device's memory. Whenever you reconfigure the memory on a Foundry device, you must save the change to the startup-config file, then reload the software to place the change into effect.

USING THE CLI

To display the configurable tables and their defaults and maximum values, enter the following command at any level of the CLI:

```
BigIron# show default values
```

```

sys log buffers:50          mac age time:300 sec      telnet sessions:5
ip arp age:10 min         bootp relay max hops:4    ip ttl:64 hops
ip addr per intf:24

when multicast enabled :
igmp group memb.:140 sec  igmp query:60 sec

when ospf enabled :
ospf dead:40 sec          ospf hello:10 sec        ospf retrans:5 sec
ospf transit delay:1 sec

when bgp enabled :
bgp local pref.:100       bgp keep alive:60 sec    bgp hold:180 sec
bgp metric:10             bgp local as:1           bgp cluster id:0
bgp ext. distance:20      bgp int. distance:200    bgp local distance:200

```

System Parameters	Default	Maximum	Current
ip-arp	8000	64000	8000
ip-static-arp	1024	2048	1024
atalk-route	512	1536	512
atalk-zone-port	64	255	64
atalk-zone-sys	255	1024	255
dvmrp	2048	32000	2048
igmp	256	1024	256
ip-cache	128000	256000	128000
ip-filter-port	512	4096	512
ip-filter-sys	1024	8192	1024
ipx-forward-filter	256	1024	256
ipx-rip-entry	3072	32728	3072
ipx-rip-filter	256	1024	256
ipx-sap-entry	6144	32768	6144
ipx-sap-filter	256	1024	256
l3-vlan	32	2048	32
ip-qos-session	2048	32000	2048
l4-real-server	1024	2048	1024
l4-virtual-server	256	512	256
l4-server-port	2048	4096	2048
mac	8000	64000	8000
ip-route	128000	200000	128000
ip-static-route	512	2048	512
vlan	16	2048	16
spanning-tree	32	128	32
mac-filter-port	32	512	32
mac-filter-sys	64	1024	64
ip-subnet-port	24	128	24
session-limit	131072	500000	131072
view	10	65535	10
virtual-interface	255	2048	255

Information for the configurable tables appears under the columns that are shown in bold type in this example. To simplify configuration, the command parameter you enter to configure the table is used for the table name. For example, to increase the capacity of the IP route table, enter the following commands:

```
BigIron(config)# system-max ip-route 120000
BigIron(config)# write memory
BigIron(config)# exit
BigIron# reload
```

NOTE: If you accidentally enter a value that is not within the valid range of values, the CLI will display the valid range for you.

To increase the number of IP subnet interfaces you can configure on each port on a NetIron Layer 3 Switch from 24 to 64, then increase the total number of IP interfaces you can configure on the device from 256 to 512, enter the following commands:

```
BigIron(config)# system-max subnet-per-interface 64
BigIron(config)# write memory
BigIron(config)# exit
BigIron# reload
```

Syntax: system-max subnet-per-interface <num>

The <num> parameter specifies the maximum number of subnet addresses per port and can be from 1 – 64. The default is 24.

Syntax: system-max subnet-per-system <num>

The <num> parameter specifies the maximum number of subnet addresses for the entire device and can be from 1 – 512. The default is 256.

```
BigIron(config)# system-max subnet-per-system 512
BigIron(config)# write memory
BigIron(config)# exit
BigIron# reload
```

You can increase the size of the IP route table for static routes by entering the following command:

```
NetIron(config)# system-max ip-static-route 8192
```

Syntax: system-max ip-static-route <num>

On most devices, the maximum number of static routes you can define is 4096. On devices running Service Provider software Release 09.1.02, you can have a maximum of 8192 static routes.

NOTE: You must reload the software for the change to take effect.

USING THE WEB MANAGEMENT INTERFACE

To modify a table size using the Web management interface:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the [Max-Parameter](#) link to display the Configure System Parameter Maximum Value table. This table lists the settings and valid ranges for all the configurable table sizes on the device.
3. Click the Modify button next to the row for the table you want to change.

4. Enter the new value for the table size. The value you enter specifies the maximum number of entries the table can hold.
5. Click Apply to save the changes to the device's running-config file.
6. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
7. Click on the plus sign next to Command in the tree view to list the command options.
8. Select the [Reload](#) link and select Yes when the Web management interface asks you whether you really want to reload the software. Changes to table sizes do not take effect until you reload the software.

Using the Temperature Sensor

The following products and modules have a temperature sensor:

- FastIron 4802
- Velocity Management Module
- Management Modules 2, and 3, and 4
- Network Processor Architecture (NPA) forwarding modules

The temperature sensor generates a Syslog message and SNMP trap if the temperature exceeds a specified warning level or shutdown level, and can shut the module down if the temperature exceeds the safe threshold. You can use the CLI or Web management interface to display the temperature and to change the warning and shutdown temperature levels. The software reads the temperature sensor according to the chassis poll time, which is 60 seconds by default.

If the temperature equals or exceeds the shutdown temperature for five consecutive polls of the temperature by the software, the software shuts down the module to prevent damage.

You can display the temperature of the module. You also can change the warning and shutdown temperatures and the chassis poll time.

Displaying the Temperature

By default, the software polls the temperature sensor on the module every 60 seconds to get the current temperature. This poll rate is controlled by the chassis poll time, which also controls how often the software polls other system components. You can display the temperature of the module using either of the following methods.

USING THE CLI

To display the temperature of a module, enter the following command at any level of the CLI:

```
BigIron> show chassis

power supply 1 not present
power supply 2 not present
power supply 3 ok
power supply 4 not present
power supply 1 to 4 from bottom to top
fan 1 ok
fan 2 bad
fan 3 ok
fan 4 ok
Current temperature : 34.5 C degrees
Warning level : 45 C degrees, shutdown level : 55 C degrees
```

Syntax: show chassis

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Monitor in the tree view to display the monitoring options.
3. Select the Device link to display the Device Information panel. The temperature is listed in the Temperature field. The temperature information is color coded to indicate the state.
 - Green indicates the temperature is within the normal operating range.
 - Orange indicates the temperature has reached the warning level.
 - Red indicates the temperature has reached the shutdown level.

NOTE: You also can display the Device Information panel by clicking on the graphic of the chassis panel, in the upper right frame. The graphic is shown only if the Web management interface frames are enabled.

Displaying Temperature Messages

The software sends a Syslog message and an SNMP trap if the temperature crosses the warning or shutdown thresholds. The following methods describe how to view the system log on the device. If you have configured the device to use a Syslog server or SNMP trap receiver, see the documentation for the server or receiver.

USING THE CLI

To display the system log, enter the following command at any CLI level:

```
BigIron# show log

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
Buffer logging: level ACDMEINW, 8 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
I=informational N=notification W=warning

Static Log Buffer:

Dynamic Log Buffer (50 entries):

at 0 days 0 hours 2 minutes 0 seconds, level alert
Temperature 48.0 C degrees, warning level 45.0 C degrees, shutdown level 55.0 C
degrees

at 0 days 0 hours 1 minutes 0 seconds, level alert
Temperature 50.0 C degrees, warning level 45.0 C degrees, shutdown level 55.0 C
degrees
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Monitor in the tree view to display the Monitor options.
3. Select the System Log link to display the system log.

Changing Temperature Warning and Shutdown Levels

The default warning temperature is 45.0 C degrees. The default shutdown temperature is 55.0 C degrees. You can change the warning and shutdown temperatures using the following commands. The valid range for each value is 0 – 125 C degrees.

NOTE: You cannot set the warning temperature to a value higher than the shutdown temperature.

USING THE CLI

To change the temperature at which the module sends a warning, enter a command such as the following at the Privileged EXEC level of the CLI:

```
BigIron# temperature warning 47
```

Syntax: temperature warning <value>

The <value> can be 0 – 125.

To change the shutdown temperature, enter a command such as the following at Privileged EXEC level of the CLI:

```
BigIron# temperature shutdown 57
```

Syntax: temperature shutdown <value>

The <value> can be 0 – 125.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Select the [Advance](#) link to display the following panel.

System

Tag Type:	<input type="text" value="8100"/>
Broadcast Limit:	<input type="text" value="0"/>
Switch Age Time:	<input type="text" value="300"/>
Default VLAN ID:	<input type="text" value="1"/>
Chassis Poll Interval (sec):	<input type="text" value="60"/>
Temperature Warning Threshold(C):	<input type="text" value="45"/>
Temperature Shutdown Threshold(C):	<input type="text" value="55"/>
Gig Port Default:	<input type="text" value="Neg-Full-Auto"/>
Mirror Slot:	<input type="text" value="None"/> <input type="text" value="Port: None"/>

[Home](#) | [Site Map](#) | [Logout](#) | [Save](#) | [Frame Enable](#) | [Disable](#) | [TELNET](#)

3. Edit the value in the Temperature Warning Threshold field to change the warning temperature.
4. Edit the value in the Temperature Shutdown Threshold field to change the shutdown temperature.
5. Click the Apply button to send the configuration change to the active module's running-config file.
6. If you want the change to remain in effect following the next system reload, select the [Save](#) link to save the configuration change to the startup-config file.

Changing the Chassis Polling Interval

The software reads the temperature sensor and polls other hardware sensors according to the value set for the chassis poll time, which is 60 seconds by default. You can change chassis poll time using the CLI.

USING THE CLI

To change the chassis poll time, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# chassis poll-time 200
```

Syntax: chassis poll-time <value>

The <value> can be 0 – 65535.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Select the [Advance](#) link to display the following panel

System	
Tag Type:	8100
Broadcast Limit:	0
Switch Age Time:	300
Default VLAN ID:	1
Chassis Poll Interval (sec):	60
Temperature Warning Threshold(C):	45
Temperature Shutdown Threshold(C):	55
Gig Port Default:	Neg-Full-Auto
Mirror Slot:	None Port: None

[Home](#) | [Site Map](#) | [Logout](#) | [Save](#) | [Frame Enable](#) | [Disable](#) | [TELNET](#)

3. Edit the value in the Chassis Poll Interval field to change polling interval. You can enter a value from 0 – 65535. The default is 60 seconds.
4. Click the Apply button to send the configuration change to the active module's running-config file.
5. If you want the change to remain in effect following the next system reload, select the [Save](#) link to save the configuration change to the startup-config file.

Assigning a Mirror Port and Monitor Ports

You can monitor traffic on Foundry ports by configuring another port to “mirror” the traffic on the ports you want to monitor. By attaching a protocol analyzer to the mirror port, you can observe the traffic on the monitored ports.

Monitoring traffic on a port is a two-step process:

- Enable a port to act as the mirror port. This is the port to which you connect your protocol analyzer.
- Enable monitoring on the ports you want to monitor.

You can monitor input traffic, output traffic, or both. Any port can operate as a mirror port and you can configure more than one mirror port. You can configure up to 64 mirror ports. You can configure the mirror ports on different modules and you can configure more than one mirror port on the same module.

Each mirror port can have its own set of monitored ports. For example, you can configure ports 1/1 and 5/1 as mirror ports, and monitor ports 1/2 – 1/8 on port 1/1 and ports 5/2 – 5/8 on port 5/1. The mirror port and monitored ports also can be on different slots.

Configuration Guidelines for Monitoring Inbound Traffic

Use the following considerations when configuring mirroring for inbound traffic on a Layer 3 Switch or a FastIron 4802 (Stackable device). The guidelines are applicable whether you configure multiple mirror ports or just one mirror port.

Guidelines for Chassis Devices

- Configure only one mirror port to monitor input traffic on a given module. If you configure multiple mirror ports on the same module, the inbound traffic for all the monitored ports on the module is sent to all the mirror ports on the same module. For example, if you configure ports 1/1 and 1/13 as mirror ports, then enable monitoring of inbound traffic on ports 1/2 and 1/14, the traffic from both ports is mirrored to both the mirror ports, 1/1 and 1/13. This occurs regardless of the mirror ports you assign to the monitor ports.
- When inbound traffic on a monitored port on one module is switched normally to another module, the switched traffic will be mirrored to the mirror ports on the other module. For example, if inbound traffic on a monitored port on the module in slot 1 is switched to the module in slot 2, mirror ports on the module in slot 2 will receive copies of the traffic. These guidelines do not apply to outbound traffic.
- If you are concurrently monitoring more than one set of ports on a non-Terathon device, there are additional restrictions on which ports can be mirror ports for monitoring inbound traffic:
 - On IronCore, do not use any of the ports on the management module as a mirror port for monitoring inbound traffic.
 - On JetCore, do not use any of the 4 lowest-numbered Gigabit Ethernet ports, or the 24 lowest-numbered 10/100 ports on the management module as a mirror port for monitoring inbound traffic.

NOTE: These restrictions do not apply to Terathon devices. On Terathon devices, any port can be mirrored and monitored except for the management port.

Guidelines for the FastIron 4802 Only

Use the following guidelines when configuring port monitoring for inbound traffic:

- In the current release, you cannot monitor inbound traffic on Gigabit ports 49 – 50 if they are configured as a trunk group. This restriction does not apply if the ports are not configured as a trunk group.
- If you configure more than one mirror port to monitor inbound traffic on the same IPC, each of the mirror ports on the IPC receives all the inbound traffic from all the monitored ports on the same IPC.
- When inbound traffic on a monitored port on one IPC is switched normally to another IPC, the switched traffic will be mirrored to the mirror ports on the other IPC. For example, if inbound traffic on a monitored port on IPC 1 is switched to IPC 2, mirror ports on IPC 2 will receive copies of the traffic.
- If you are concurrently monitoring more than one set of ports on the device, do not use 10/100 ports 1 – 24 or Gigabit Ethernet port 49 as a mirror port for monitoring inbound traffic.

These guidelines do not apply to monitoring outbound traffic. You can monitor traffic between IPCs and use multiple mirror ports.

Notes Regarding Monitoring of Router Traffic

- For inbound traffic that is routed (not switched), if the traffic is forwarded by the hardware and thus bypasses the CPU, the port that receives the traffic changes the source and destination MAC addresses of the packet before sending the packet to its outbound port and the mirror port.
- For outbound traffic that is routed (not switched), the source MAC address of the traffic that is copied to the mirror port has the MAC address of the mirror port rather than the monitored port's MAC address.

This happens because the routed traffic sent by the router interface must address itself as the sender of the packet, to the neighboring router. This behavior cannot be turned off for the monitored traffic, so the mirror port's MAC address is substituted for the mirror copy of the packet. In this case, the source MAC address of the mirror port is equivalent to that of the monitored port.

Configuring Port Mirroring and Monitoring on Non-Terathon Devices

Use one of the methods below to configure port mirroring and monitoring for non-Terathon devices.

USING THE CLI

Suppose you want to diagnose the in and out traffic on port 3 on a module in slot 4 of a BigIron, and use port 1 in slot 4 as the mirror port. To do so, enter the following commands:

```
BigIron(config)# mirror-port ethernet 4/1
BigIron(config)# interface ethernet 4/3
BigIron(config-if-4/3)# monitor ethernet 4/1 both
```

Syntax: [no] mirror-port ethernet <portnum>

The <portnum> parameter specifies the port. You can configure up to 64 mirror ports on a Layer 3 Switch and up to 50 mirror ports on a FastIron 4802.

Syntax: [no] monitor ethernet <portnum> [ethernet <portnum>...] both | in | out

The <portnum> parameter specifies the mirror port(s).

The **both | in | out** parameter specifies the traffic direction you want to monitor on the mirror port. There is no default.

NOTE: You can configure multiple mirror ports on the same module. However, if you mirror inbound traffic to any of the mirror ports on the module, the traffic is mirrored to all the mirror ports on the module. If you plan to mirror outbound traffic only, you can use multiple mirror ports on the same module without the traffic being duplicated on the other mirror ports on the module.

NOTE: The following behavior applies to the following releases:

- Enterprise software releases prior to release 07.8.00; it is not longer applicable with software releases 07.8.00 and later.
- ServerIron Chassis software releases
- Service Provider software releases

If you configure the device to monitor inbound traffic on multiple ports and use a single mirror port for the traffic, disabling monitoring on one of the ports also disables monitoring on the other ports. For example, if you configure the device to monitor inbound traffic on ports 1/1 and 1/2 and to mirror the traffic to port 2/1, if you then disable monitoring of inbound traffic on port 1/2, the software also disables monitoring of inbound traffic on port 1/1.

This guideline does not apply to monitoring outbound traffic. Disabling monitoring for outbound traffic does not affect other ports that use the same mirror port.

If your traffic is to be monitored in both directions, only the inbound traffic monitoring is disabled on the other ports.

To configure more than one mirror port, enter commands such as the following:

```
BigIron(config)# mirror-port ethernet 1/1
BigIron(config)# mirror-port ethernet 5/1
BigIron(config)# mirror-port ethernet 5/1
BigIron(config)# mirror-port ethernet 5/2
```

These commands configure four mirror ports.

The following commands configure ports on the module in slot 1 to be mirrored by port 1/1:

```
BigIron(config)# interface ethernet 1/2
BigIron(config-if-1/2)# monitor ethernet 1/1 in
BigIron(config-if-1/2)# interface ethernet 1/3
BigIron(config-if-1/3)# monitor ethernet 1/1 in
BigIron(config-if-1/3)# interface ethernet 1/4
BigIron(config-if-1/4)# monitor ethernet 1/1 in
```

These commands configure the inbound traffic on ports 1/2 – 1/4 to be mirrored to port 1/1.

USING THE WEB MANAGEMENT INTERFACE

Suppose you want to diagnose the in and out on traffic on port 3 on a module in slot 4 of a BigIron using port 1 in slot 4. To do so:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the [Advance](#) link to display the advanced system configuration panel.
3. Select the slot (if applicable) and port from the corresponding pulldown menus next to Mirror Slot. In this example, select slot 4 and port 1.
4. Click Apply to save the changes to the device's running-config file.
5. Click on the plus sign next to Configure in the tree view to display the configuration options.
6. Click on the plus sign next to Port in the tree view to display the configuration options.
7. Select the link to the port type you want (for example, [Ethernet](#)) to display the Port table.
8. Click the Modify button next to the port you want to monitor. In this example, select port 3 on the module in slot 4 (4/3).
9. Select the traffic direction you want to monitor. For this example, select the In & Out.
10. Click Apply to save the changes to the device's running-config file.
11. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring Port Mirroring and Monitoring on Terathon Devices

Port mirroring and monitoring was introduced in Terathon IronWare software Release 01.0.00.

You can configure multiple mirror ports on the same module. However, if you mirror inbound traffic to any of the mirror ports on the module, the traffic is mirrored to all the mirror ports on the module. If you plan to mirror outbound traffic only, you can use multiple mirror ports on the same module without the traffic being duplicated on the other mirror ports on the module.

NOTE: On Terathon devices, you cannot monitor outbound traffic from one armed router traffic.

The following example configures two mirror ports on the same module and one mirror port on another module. It will illustrate how inbound traffic is mirrored to the two mirror ports on the same module even if the traffic is configured to be mirrored to only one mirror port on the module.

```
BigIron MG8(config)# mirror-port ethernet 1/1
BigIron MG8(config)# mirror-port ethernet 1/2
BigIron MG8(config)# mirror-port ethernet 2/1

BigIron MG8(config)# interface ethernet 3/1
BigIron MG8(config-if-e10000-3/1)# monitor ethernet 1/1 both
BigIron MG8(config-if-e10000-3/1)# monitor ethernet 2/1 in

BigIron MG8(config-if-e10000-3/1)# interface ethernet 4/13
BigIron MG8(config-if-e10000-4/1)# monitor ethernet 1/2 both
```

This example configures two mirror ports 1/1 and 1/2 on the same module. It also configures input and output traffic from port 3/1 to be mirrored to mirror port 1/1 and input and output traffic from port 4/1 to be mirrored to mirror port 1/2. Because mirror ports 1/1 and 1/2 are configured on the same module, mirror port 1/1 will receive the input traffic from port 3/1 as well as port 4/1 and mirror port 1/2 will receive input traffic from port 4/1 as well as port 3/1 even if they are not explicitly configured to do so. The outbound traffic from port 3/1 is mirrored to port 1/1 only, as configured and the outbound traffic from port 4/1 is mirrored to port 1/2 only as configured.

This example also configures one mirror port 2/1 on another module, to which inbound traffic from port 3/1 is mirrored. Because only one mirror port is configured on this module, the traffic is mirrored as configured.

If input monitoring is enabled on two ports controlled by the same packet processor, then the input traffic on these two ports will be mirrored to all the ports configured as mirror ports for these two monitored ports. This restriction does not apply to outbound monitoring.

```
BigIron MG8(config)# mirror-port ethernet 1/1
BigIron MG8(config)# mirror-port ethernet 2/1
BigIron MG8(config)# interface ethernet 3/1
BigIron MG8(config-if-e1000-3/1)# monitor ethernet 1/1 both
BigIron MG8(config-if-e1000-3/1)# interface ethernet 3/2
BigIron MG8(config-if-e1000-3/2)# monitor ethernet 2/1 both
```

The above example configures two mirror ports 1/1 and 2/1 on different modules. Port 3/1 uses port 1/1 for inbound and outbound mirroring. Port 3/2 uses port 2/1 for inbound and outbound mirroring. If 3/1 and 3/2 are controlled by the same packet processor, inbound traffic from 3/1 will be mirrored to 1/1 as well as 2/1 and similarly, inbound traffic from 3/2 will be mirrored to 2/1 as well as 1/1. The outbound traffic on 3/1 and 3/2 are mirrored according to the configuration.

Monitoring an Individual Trunk Port

By default, when you monitor the primary port in a trunk group, aggregated traffic for all the ports in the trunk group is copied to the mirror port. You can configure the device to monitor individual ports in a trunk group. You can monitor the primary port or a secondary port individually.

NOTE: In the current release, you can use only one mirror port for each monitored trunk port.

To monitor traffic on an individual port in a trunk group, enter commands such as the following:

```
BigIron(config)# mirror ethernet 2/1
BigIron(config)# trunk switch ethernet 4/1 to 4/8
BigIron(config-trunk-4/1-4/8)# config-trunk-ind
BigIron(config-trunk-4/1-4/8)# monitor ethe-port-monitored 4/5 ethernet 2/1 in
```

Syntax: [no] config-trunk-ind

Syntax: [no] monitor ethe-port-monitored <portnum> | named-port-monitored <portname>
ethernet | pos <portnum> in | out | both

The **config-trunk-ind** command enables configuration of individual ports in the trunk group. You need to enter the **config-trunk-ind** command only once in a trunk group. After you enter the command, all applicable port configuration commands apply to individual ports only.

NOTE: If you enter **no config-trunk-ind**, all port configuration commands are removed from the individual ports and the configuration of the primary port is applied to all the ports. Also, once you enter the **no config-trunk-ind** command, the **enable**, **disable**, and **monitor** commands are valid only on the primary port and apply to the entire trunk group.

The **monitor ethe-port-monitored** command in this example enables monitoring of the inbound traffic on port 4/5.

- The **ethe-port-monitored <portnum> | named-port-monitored <portname>** parameter specifies the trunk port you want to monitor. Use **ethe-port-monitored <portnum>** to specify a port number. Use **named-port-monitored <portname>** to specify a trunk port name.
- The **ethernet | pos <portnum>** parameter specifies the port to which the traffic analyzer is attached.
- The **in | out | both** parameter specifies the traffic direction to be monitored.

Monitoring 802.3ad Aggregate Links

Starting in software release 07.8.00, you can monitor 802.3ad aggregate links, as well as individual ports within 802.3ad aggregate links.

NOTE: The terms *802.3ad aggregate link* and *dynamic trunk group* are used interchangeably in this section and mean the same thing.

Configuration Note

- This feature is supported in software releases 07.8.00 and later.
- This feature is supported on any port that can be configured with 802.3ad link aggregation.

Configuring Port Monitoring on 802.3ad Aggregate Links

By default, when you enable monitoring on the primary port of an 802.3ad aggregate link, the device copies the traffic for all the ports in the dynamic trunk group to the mirror port.

To monitor all of the ports in an 802.3ad aggregate link, enter commands such as the following on the primary port of the dynamic trunk group:

```
BigIron(config)# interface e1/1
BigIron(config-if-e100-1/1)# link-aggregate monitor ethernet-port-monitored e 1/1 e
1/10 both
```

These commands enable monitoring of the entire dynamic trunk group and copy both incoming and outgoing traffic to port 1/10, the assigned mirror port. Note that the mirror port (in this case, port 1/10) must already be configured as a mirror port.

Syntax: link-aggregate monitor ethernet-port-monitored ethernet <monitor slot/port> <mirror slot/port> both | in | out

The <monitor slot/port> parameter specifies the port to monitor.

The <mirror slot/port> parameter specifies the port that will receive copies of the monitored port's traffic.

The **both | in | out** parameter specifies the traffic direction to monitor. There is no default.

Configuring Port Monitoring on an Individual Port in an 802.3ad Aggregate Link

To monitor traffic on an individual port in a dynamic trunk group, enter commands such as the following:

```
BigIron(config)#interface e1/1
BigIron(config-if-e100-1/1)# link-aggregate config-ind-monitor
BigIron(config-if-e100-1/1)# link-aggregate monitor ethernet-port-monitored
ethernet 1/1 ethernet 1/10 in
```

Syntax: [no] link-aggregate config-ind-monitor

Syntax: link-aggregate monitor ethernet-port-monitored ethernet <monitor slot/port> <mirror slot/port> in | out | both

The **link-aggregate config-ind-monitor** command enables configuration of individual ports in the dynamic trunk group. Enter this command only once in a dynamic trunk group configuration. After you enter this command, all applicable port configuration commands apply to individual ports only.

NOTE: If you enter **no link-aggregate config-ind-monitor**, the device removes all monitor configuration commands from the individual ports and applies the primary port's configuration to all the ports. Also, once you enter the **no link-aggregate config-ind-monitor** command, any monitor configuration command you enter thereafter applies to the entire trunk group.

The **link-aggregate monitor ethernet-port-monitored ethernet** command in this example enables monitoring of inbound traffic on port 1/1.

- The <monitor slot/port> parameter specifies the port to monitor.

- The <mirror slot/port> parameter specifies the port that will receive copies of the monitored port's traffic.
- The **in** | **out** | **both** parameter specifies the traffic direction to monitor. There is no default.

Mirror Ports for Policy-Based Routing (PBR) Traffic

NOTE: This feature applies to hardware-based PBR, which is currently supported only on JetCore and FastIron 4802 (FWS 4802) premium devices, and on 10 Gigabit Ethernet modules. Also, this feature is not supported on Terathon devices.

Software release 07.6.03 and later allows you to mirror traffic on ports that have policy-based routing (PBR) enabled. This feature is useful for monitoring traffic, debugging, and enabling application-specific mirroring.

The PBR mirror interface feature allows continued hardware forwarding and, at the same time, enables you to determine exactly which traffic flows get routed using the policies defined by PBR.

The following section provides a general overview of hardware-based PBR. For more specific information about hardware based PBR, see the chapter “JetCore Hardware-Based IP Access Control Lists (ACLs)” in the *Foundry Enterprise Configuration and Management Guide*.

About Hardware-Based PBR

Hardware-based Policy-Based Routing (PBR) routes traffic in hardware based on policies you define. A PBR policy specifies the next hop for traffic that matches the policy. A PBR policy also can use an ACL to perform QoS mapping and marking for traffic that matches the policy.

To configure PBR, you define the policies using IP ACLs and route maps, then enable PBR globally or on individual interfaces. The device programs the ACLs into the Layer 4 CAM on the interfaces and routes traffic that matches the ACLs according to the instructions in the route maps. You also can map and mark the traffic's QoS information using the QoS options of the ACLs.

Configuring Mirror Ports for PBR Traffic

When you configure a physical or virtual port to act as a mirror port for PBR traffic, outgoing packets that match the permit Access Control List (ACL) clause in the route map are copied to the mirror port(s) that you specify. You can specify up to four mirror ports for each PBR route map instance.

For example, to capture all traffic forwarded to an SSL port and mirror it to port 5, enter commands such as the following:

```
BigIron(config)# route-map ssl-pbr-map permit 1
BigIron(config-routemap ssl-pbr-map)# match ip address 100
BigIron(config-routemap ssl-pbr-map)# set mirror-interface 5
BigIron(config-routemap ssl-pbr-map)# set next-hop 10.10.10.1
BigIron(config-routemap ssl-pbr-map)# exit
BigIron(config)# interface e 5
BigIron(config-if-5)# port-name mirror-port
BigIron(config-if-mirror-port)# interface e 10
BigIron(config-if-10)# ip policy route-map ssl-pbr-map
BigIron(config-if-10)# exit
BigIron(config)# access-list 100 permit tcp any any eq ssl
```

The above commands complete the following configuration tasks:

1. Configures an entry in the PBR route map named “ssl-pbr-map”. The **match** statement matches on IP information in ACL 100. The **set mirror-interface** statement specifies interface e 5 as the mirror port for matched ACL permit clauses. The **set next-hop** statement sets the IP address of the route's next hop router to 10.10.10.1.
2. Identifies interface e 5 as a mirror port by assigning the name “mirror-port”.
3. Enables PBR and applies the route map “ssl-pbr-map” on interface e 10.
4. Creates an extended ACL (100) that permits all TCP traffic destined for an for an SSL port.

NOTE: This section describes the syntax for the new CLI Route Map level command, **set mirror-interface**. For more information about the other existing commands and syntax shown in the above example, see the *Foundry Switch and Router Command Line Interface Reference* or the *Foundry Enterprise Configuration and Management Guide*.

Syntax: set mirror-interface <slot number>/<port number>

The <slot number> parameter specifies the port number on a Foundry Layer 3 Switch. This parameter is not applicable to stackable devices.

The <port number> parameter specifies the mirror port number.

You can specify up to 4 mirror ports for each PBR route map instance. To do so, enter the **set mirror interface** command for each mirror port.

Displaying the Current Mirror and Monitor Port Configuration on Non-Terathon Devices

You can display the current port mirroring and monitoring configuration using the following CLI method.

USING THE CLI

To display the current mirroring and monitoring configuration, enter the following command at any level of the CLI:

```
BigIron(config)# show monitor
Mirror Interface:      ethernet 4/1
Monitored Interfaces:
  Both                Input                Output
-----
  ethernet 4/3
```

Syntax: show monitor

This example shows the monitoring and mirroring configuration set up by the commands in the example in the previous section. Port 4/1 is the mirror interface, to which the software copies (“mirrors”) the traffic on port 4/3. In this case, both directions of traffic on the monitored port are mirrored to port 4/1.

If only the incoming traffic is mirrored, the monitored interface is listed under Input. If only the outbound traffic is mirrored, the monitored interface is listed under Output.

USING THE WEB MANAGEMENT INTERFACE

You cannot display this information using the Web management interface.

Displaying Mirror and Monitor Port Configuration on Terathon Devices

To display the inbound and outbound traffic mirrored to each mirror port as configured on a Terathon device, enter the following command at any level of the CLI:

```
BigIron MG8# show monitor config
Monitored Port 3/1
  Input traffic mirrored to: 1/1 2/1
  Output traffic mirrored to: 1/1
Monitored Port 4/1
  Input traffic mirrored to: 1/2
  Output traffic mirrored to: 1/2
```

Syntax: show monitor config

This output does not display the input traffic mirrored to mirror port 1/2 from port 3/1 and mirrored to mirror port 1/1 from port 4/1 because the mirroring of this traffic is not explicitly configured.

To display the actual traffic mirrored to each mirror port despite the configuration on a Terathon device, enter the following command at any level of the CLI:

```
BigIron MG8# show monitor actual
Monitored Port 3/1
  Input traffic mirrored to: 1/1(configured) 1/2 2/1(configured)
  Output traffic mirrored to: 1/1
Monitored Port 4/1
  Input traffic mirrored to: 1/2(configured) 1/1
  Output traffic mirrored to: 1/2
```

Syntax: show monitor actual

This output displays the input traffic mirrored to mirror port 1/2 from port 3/1 and mirrored to mirror port 1/1 from port 4/1, which are not explicitly configured.

Chapter 10

Configuring Spanning Tree Protocol (STP) and IronSpan Features

The Spanning Tree Protocol (STP) eliminates Layer 2 loops in networks, by selectively blocking some ports and allowing other ports to forward traffic, based on global (bridge) and local (port) parameters you can configure.

This chapter describes how to configure Spanning Tree Protocol (STP) parameters on Foundry Layer 3 Switches.

You can enable or disable STP on a global basis (for the entire device), a port-based VLAN basis (for the individual Layer 2 broadcast domain), or an individual port basis.

Configuration procedures are provided for the standard STP bridge and port parameters as well as Foundry IronSpan parameters.

IronSpan is a set of Layer 2 features that enable you to overcome limitations in the standard 802.1d Spanning Tree Protocol (STP). IronSpan includes the following features:

- Fast Port Span
- Fast Uplink Span
- Rapid Spanning Tree (both 802.1W Draft 3 and full 802.1W are supported)
- Single-instance STP
- SuperSpan™
- STP per VLAN group
- Per VLAN Spanning Tree (PVST) and PVST+ Compatibility

These enhancements extend the operation of standard STP. IronSpan enables you to fine tune standard STP and avoid some of its limitations.

- To configure standard STP parameters, see “Configuring Standard STP Parameters”.
- To configure IronSpan parameters, see “Configuring IronSpan Features” on page 10-19.

Configuring Standard STP Parameters

Foundry Layer 2 Switches and s support standard STP as described in the IEEE 802.1D specification. STP is enabled by default on Layer 2 Switches but disabled by default on Layer 3 Switches.

By default, each port-based VLAN on a Foundry device runs a separate spanning tree (a separate instance of STP). A Foundry device has one port-based VLAN (VLAN 1) by default that contains all the device's ports. Thus, by default each Foundry device has one spanning tree. However, if you configure additional port-based VLANs on a Foundry device, then each of those VLANs on which STP is enabled and VLAN 1 all run separate spanning trees.

If you configure a port-based VLAN on the device, the VLAN has the same STP state as the default STP state on the device. Thus, on Layer 2 Switches, new VLANs have STP enabled by default. On Layer 3 Switches, new VLANs have STP disabled by default. You can enable or disable STP in each VLAN separately. In addition, you can enable or disable STP on individual ports.

STP Parameters and Defaults

Table 10.1 lists the default STP states for Foundry devices.

Table 10.1: Default STP States

Device Type	Default STP Type	Default STP State	Default STP State of New VLANs ^a
Layer 2 Switch	MSTP ^b	Enabled	Enabled
Layer 3 Switch	MSTP	Disabled	Disabled
ServerIron	MSTP	Enabled	Enabled

a. When you create a port-based VLAN, the new VLAN's STP state is the same as the default STP state on the device. The new VLAN does not inherit the STP state of the default VLAN.

b. MSTP stands for "Multiple Spanning Tree Protocol". In this type of STP, each port-based VLAN, including the default VLAN, has its own spanning tree. References in this documentation to "STP" apply to MSTP. The Single Spanning Tree Protocol (SSTP) is another type of STP. SSTP includes all VLANs on which STP is enabled in a single spanning tree. See "Single Spanning Tree (SSTP)" on page 10-62.

Table 10.2 lists the default STP bridge parameters. The bridge parameters affect the entire spanning tree. If you are using MSTP, the parameters affect the VLAN. If you are using SSTP, the parameters affect all VLANs that are members of the single spanning tree.

Table 10.2: Default STP Bridge Parameters

Parameter	Description	Default and Valid Values
Forward Delay	The period of time a bridge will wait (the listen and learn period) before beginning to forward data packets.	15 seconds Possible values: 4 – 30 seconds
Maximum Age	The interval a bridge will wait for a hello packet from the root bridge before initiating a topology change.	20 seconds Possible values: 6 – 40 seconds
Hello Time	The interval of time between each configuration BPDU sent by the root bridge.	2 seconds Possible values: 1 – 10 seconds
Priority	A parameter used to identify the root bridge in a spanning tree (instance of STP). The bridge with the lowest value has the highest priority and is the root. A higher numerical value means a lower priority; thus, the highest priority is 0.	32768 Possible values: 0 – 65535

NOTE: If you plan to change STP bridge timers, Foundry recommends that you stay within the following ranges, from section 8.10.2 of the IEEE STP specification.

$$2 * (\text{forward_delay} - 1) \geq \text{max_age}$$

$$\text{max_age} \geq 2 * (\text{hello_time} + 1)$$

Table 10.3 lists the default STP port parameters. The port parameters affect individual ports and are separately configurable on each port.

Table 10.3: Default STP Port Parameters

Parameter	Description	Default and Valid Values
Priority	The preference that STP gives this port relative to other ports for forwarding traffic out of the spanning tree. A higher numerical value means a lower priority; thus, the highest priority is 8.	128 Possible values: 8 – 252 (configurable in increments of 4)
Path Cost	The cost of using the port to reach the root bridge. When selecting among multiple links to the root bridge, STP chooses the link with the lowest path cost and blocks the other paths. Each port type has its own default STP path cost.	10 Mbps – 100 100 Mbps – 19 Gigabit – 4 10 Gigabit – 2 OC-3c – 200 OC-12c – 80 OC-48c – 20 Possible values are 0 – 65535

Notice that the path costs favor 100 Mbps and faster Ethernet ports over Packet over SONET (POS) ports. STP applies to POS ports only when they are configured for Layer 2 switching. The POS paths are remote (WAN) paths and the Ethernet paths are local paths. The default path costs therefore favor local paths over remote paths. See “Configuring POS for Layer 2 Switching” on page 7-16.

Enabling or Disabling the Spanning Tree Protocol (STP)

You can enable or disable STP on the following levels:

- Globally – Affects all ports on the device.
- Port-based VLAN – Affects all ports within the specified port-based VLAN. When you enable or disable STP within a port-based VLAN, the setting overrides the global setting. Thus, you can enable STP for the ports within a port-based VLAN even when STP is globally disabled, or disable the ports within a port-based VLAN when STP is globally enabled.
- Individual port – Affects only the individual port. However, if you change the STP state of the primary port in a trunk group, the change affects all ports in the trunk group.

Enabling or Disabling STP Globally

Use the following methods to enable or disable STP on a device on which you have not configured port-based VLANs.

NOTE: When you configure a VLAN, the VLAN inherits the global STP settings. However, once you begin to define a VLAN, you can no longer configure standard STP parameters globally using the CLI. From that point on, you can configure STP only within individual VLANs.

USING THE CLI

To enable STP for all ports in all VLANs on a Foundry device, enter the following command:

```
BigIron(config)# spanning-tree
```

This command enables a separate spanning tree in each VLAN, including the default VLAN.

Syntax: [no] spanning-tree

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select Enable next to Spanning Tree.

NOTE: For information about the Single and Fast checkboxes, see “Single Spanning Tree (SSTP)” on page 10-62 and “Fast Uplink Span” on page 10-21.

3. Click Apply to save the changes to the device's running-config file.
4. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Enabling or Disabling STP in a Port-Based VLAN

Use the following procedure to disable or enable STP on a device on which you have configured a port-based VLAN. Changing the STP state in a VLAN affects only that VLAN.

USING THE CLI

To enable STP for all ports in a port-based VLAN, enter commands such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning-tree
```

Syntax: [no] spanning-tree

USING THE WEB MANAGEMENT INTERFACE

You cannot enable or disable STP on individual VLANs using the Web management interface.

Enabling or Disabling STP on an Individual Port

Use the following procedure to disable or enable STP on an individual port.

NOTE: If you change the STP state of the primary port in a trunk group, the change affects all ports in the trunk group.

USING THE CLI

To enable STP on an individual port, enter commands such as the following:

```
BigIron(config)# interface 1/1
BigIron(config-if-1/1)# spanning-tree
```

Syntax: [no] spanning-tree

USING THE WEB MANAGEMENT INTERFACE

You cannot enable or disable STP on individual ports using the Web management interface.

Changing STP Bridge and Port Parameters

Table 10.2 on page 10-2 and Table 10.3 on page 10-3 list the default STP parameters. If you need to change the default value for an STP parameter, use the following procedures.

Changing STP Bridge Parameters

To change STP bridge parameters, use either of the following methods.

NOTE: If you plan to change STP bridge timers, Foundry recommends that you stay within the following ranges, from section 8.10.2 of the IEEE STP specification.

$$2 * (\text{forward_delay} - 1) \geq \text{max_age}$$

$$\text{max_age} \geq 2 * (\text{hello_time} + 1)$$

USING THE CLI

To change a Foundry device's STP bridge priority to the highest value to make the device the root bridge, enter the following command:

```
BigIron(config)# spanning-tree priority 0
```

The command in this example changes the priority on a device on which you have not configured port-based VLANs. The change applies to the default VLAN. If you have configured a port-based VLAN on the device, you can configure the parameters only at the configuration level for individual VLANs. Enter commands such as the following:

```
BigIron(config)# vlan 20
BigIron(config-vlan-20)# spanning-tree priority 0
```

To make this change in the default VLAN, enter the following commands:

```
BigIron(config)# vlan 1
BigIron(config-vlan-1)# spanning-tree priority 0
```

Syntax: [no] spanning-tree [forward-delay <value>] | [hello-time <value>] | [maximum-age <value>] | [priority <value>]

The **forward-delay** <value> parameter specifies the forward delay and can be a value from 4 – 30 seconds. The default is 15 seconds.

NOTE: You can configure a Foundry device for faster convergence (including a shorter forward delay) using Fast Span or Fast Uplink Span. See “Configuring IronSpan Features” on page 10-19.

The **hello-time** <value> parameter specifies the hello time and can be a value from 1 – 10 seconds. The default is 2 seconds.

NOTE: This parameter applies only when this device or VLAN is the root bridge for its spanning tree.

The **maximum-age** <value> parameter specifies the amount of time the device waits for receipt of a hello packet before initiating a topology change. You can specify from 6 – 40 seconds. The default is 20 seconds.

The **priority** <value> parameter specifies the priority and can be a value from 0 – 65535. A higher numerical value means a lower priority. Thus, the highest priority is 0. The default is 32768.

You can specify some or all of these parameters on the same command line. If you specify more than one parameter, you must specify them in the order shown above, from left to right.

USING THE WEB MANAGEMENT INTERFACE

To modify the STP parameters:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.

- Click on the plus sign next to Configure in the tree view to display the configuration options.
- Select the [STP](#) link to display the STP bridge and port parameters.
- Click the Modify button in the STP bridge parameters table to display the STP configuration panel, as shown in the following example. If the device has multiple port-based VLANs, select the Modify button next to the VLAN on which you want to change the parameters. A dialog such as the following is displayed.

STP	
VLAN ID:	<input type="text" value="1"/>
Bridge	
Forward Delay (Seconds):	<input type="text" value="15"/>
Maximum Age (Seconds):	<input type="text" value="20"/>
Hello Time (Seconds):	<input type="text" value="2"/>
Priority:	<input type="text" value="32768"/>
<input type="button" value="Apply"/>	
Port	
Priority:	<input type="text" value="128"/>
Path Cost:	<input type="text" value="0"/>
Slot:	<input type="text" value="1"/> <input type="text" value="Port: 1"/>
<input type="button" value="Apply Port STP"/> <input type="button" value="Apply To All Ports"/>	

[\[Show\]](#)[\[Statistic\]](#)

[\[Home\]](#)[\[Site Map\]](#)[\[Logout\]](#)[\[Save\]](#)[\[Frame Enable\]](#)[\[Disable\]](#)[\[TELNET\]](#)

- Modify the bridge STP parameters to the values desired.
- Click Apply to save the changes to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Changing STP Port Parameters

To change STP port parameters, use either of the following methods.

USING THE CLI

To change the path and priority costs for a port, enter commands such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning-tree ethernet 1/5 path-cost 15 priority 64
```

Syntax: spanning-tree atm | ethernet | pos <portnum> path-cost <value> | priority <value> | disable | enable

The **atm | ethernet | pos** <portnum> parameter specifies the interface.

The **path-cost** <value> parameter specifies the port's cost as a path to the spanning tree's root bridge. STP prefers the path with the lowest cost. You can specify a value from 0 – 65535.

The default depends on the port type:

- 10 Mbps – 100
- 100 Mbps – 19
- Gigabit – 4
- 10 Gigabit – 2

- OC-3c – 200
- OC-12c – 80
- OC-48c – 20

The **priority** <value> parameter specifies the preference that STP gives this port relative to other ports for forwarding traffic out of the spanning tree. You can specify a value from 8 – 252, in increments of 4. If you enter a value that is not divisible by four the software rounds to the nearest value that is. The default is 128. A higher numerical value means a lower priority; thus, the highest priority is 8.

NOTE: The range in software releases earlier than 07.5.01 is 0 – 255. If you are upgrading a device that has a configuration saved under an earlier software release, and the configuration contains a value from 0 – 7 for a port's STP priority, the software changes the priority to the default when you save the configuration while running the new release.

The **disable | enable** parameter disables or re-enables STP on the port. The STP state change affects only this VLAN. The port's STP state in other VLANs is not changed.

USING THE WEB MANAGEMENT INTERFACE

To modify the STP port parameters:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the STP link to display the STP bridge and port parameters.
4. Click the Modify button in the STP port parameters table to display the STP configuration panel, as shown in the following example. If the device has multiple port-based VLANs, select the Modify button next to the VLAN on which you want to change the parameters. A dialog such as the following is displayed.

STP	
VLAN ID:	<input type="text" value="1"/>
Bridge	
Forward Delay (Seconds):	<input type="text" value="15"/>
Maximum Age (Seconds):	<input type="text" value="20"/>
Hello Time (Seconds):	<input type="text" value="2"/>
Priority:	<input type="text" value="32768"/>
<input type="button" value="Apply"/>	
Port	
Priority:	<input type="text" value="128"/>
Path Cost:	<input type="text" value="0"/>
Slot:	<input type="text" value="1"/> Port: <input type="text" value="1"/>
<input type="button" value="Apply Port STP"/> <input type="button" value="Apply To All Ports"/>	

[Show][Statistic]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

5. Select the port (and slot if applicable) from the Port and Slot pulldown lists.
6. Enter the desired changes to the priority and path cost fields.
7. Click Apply STP Port to apply the changes to only the selected port or select Apply To All Ports to apply the changes to all the ports.

NOTE: If you want to save the priority and path costs of one port to all other ports on the device or within the selected VLAN, you can click the Apply To All Ports button.

8. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Displaying STP Information

You can display the following STP information:

- All the global and interface STP settings
- CPU utilization statistics
- Detailed STP information for each interface
- STP state information for a port-based VLAN
- STP state information for an individual interface

Displaying STP Information for an Entire Device

To display STP information for an entire device, use either of the following methods.

USING THE CLI

To display STP information, enter the following command at any level of the CLI:

```
BigIron# show span

VLAN 1 BPDU cam_index is 3 and the Master DMA Are(HEX)
STP instance owned by VLAN 1

Global STP (IEEE 802.1D) Parameters:

VLAN Root          Root Root Prio Max He- Ho- Fwd Last      Chg  Bridge
ID   ID            Cost Port rity Age llo ld  dly Chang cnt  Address
                                Hex  sec sec sec  sec sec
    1 800000e0804d4a00 0   Root 8000 20  2  1  15  689    1   00e0804d4a00

Port STP Parameters:

Port  Prio Path  State      Fwd  Design  Designated  Designated
Num   rity Cost  State      Trans Cost  Root         Bridge
    Hex
  1    80  19   FORWARDING 1     0      800000e0804d4a00 800000e0804d4a00
  2    80  0    DISABLED   0     0      0000000000000000 0000000000000000
  3    80  0    DISABLED   0     0      0000000000000000 0000000000000000
  4    80  0    DISABLED   0     0      0000000000000000 0000000000000000
  5    80  19   FORWARDING 1     0      800000e0804d4a00 800000e0804d4a00
  6    80  19   BLOCKING   0     0      800000e0804d4a00 800000e0804d4a00
  7    80  0    DISABLED   0     0      0000000000000000 0000000000000000
<lines for remaining ports excluded for brevity>
```

Syntax: show span [vlan <vlan-id>] | [pvst-mode] | [<num>] | [detail [vlan <vlan-id>] [atm <portnum> | ethernet <portnum> | pos <portnum>] | <num>]]

The **vlan** <vlan-id> parameter displays STP information for the specified port-based VLAN.

The **pvst-mode** parameter displays STP information for the device's Per VLAN Spanning Tree (PVST+) compatibility configuration. See "PVST/PVST+ Compatibility" on page 10-75.

The <num> parameter displays only the entries after the number you specify. For example, on a device with three port-based VLANs, if you enter 1, then information for the second and third VLANs is displayed, but information for the first VLAN is not displayed. Information is displayed according to VLAN number, in ascending order. The entry number is not the same as the VLAN number. For example, if you have port-based VLANs 1, 10, and 2024, then the command output has three STP entries. To display information for VLANs 10 and 2024 only, enter **show span 1**.

The **detail** parameter and its additional optional parameters display detailed information for individual ports. See “Displaying Detailed STP Information for Each Interface” on page 10-14.

The **show span** command shows the following information.

Table 10.4: CLI Display of STP Information

This Field...	Displays...
Global STP Parameters	
VLAN ID	The port-based VLAN that contains this spanning tree (instance of STP). VLAN 1 is the default VLAN. If you have not configured port-based VLANs on this device, all STP information is for VLAN 1.
Root ID	The ID assigned by STP to the root bridge for this spanning tree.
Root Cost	The cumulative cost from this bridge to the root bridge. If this device is the root bridge, then the root cost is 0.
Root Port	The port on this device that connects to the root bridge. If this device is the root bridge, then the value is “Root” instead of a port number.
Priority Hex	This device or VLAN's STP priority. The value is shown in hexadecimal format. Note: If you configure this value, specify it in decimal format. See “Changing STP Bridge Parameters” on page 10-5.
Max age sec	The number of seconds this device or VLAN waits for a hello message from the root bridge before deciding the root has become unavailable and performing a reconvergence.
Hello sec	The interval between each configuration BPDU sent by the root bridge.
Hold sec	The minimum number of seconds that must elapse between transmissions of consecutive Configuration BPDUs on a port.
Fwd dly sec	The number of seconds this device or VLAN waits following a topology change and consequent reconvergence.
Last Chang sec	The number of seconds since the last time a topology change occurred.
Chg cnt	The number of times the topology has changed since this device was reloaded.
Bridge Address	The STP address of this device or VLAN. Note: If this address is the same as the Root ID, then this device or VLAN is the root bridge for its spanning tree.
Port STP Parameters	
Port Num	The port number.

Table 10.4: CLI Display of STP Information (Continued)

This Field...	Displays...
Priority Hex	<p>The port's STP priority, in hexadecimal format.</p> <p>Note: If you configure this value, specify it in decimal format. See "Changing STP Port Parameters" on page 10-6.</p>
Path Cost	<p>The port's STP path cost.</p>
State	<p>The port's STP state. The state can be one of the following:</p> <ul style="list-style-type: none"> • BLOCKING – STP has blocked Layer 2 traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is FORWARDING. When a port is in this state, the port does not transmit or receive user frames, but the port does continue to receive STP BPDUs. • DISABLED – The port is not participating in STP. This can occur when the port is disconnected or STP is disabled on the port. • FORWARDING – STP is allowing the port to send and receive frames. • LISTENING – STP is responding to a topology change and this port is listening for a BPDU from neighboring bridge(s) in order to determine the new topology. No user frames are transmitted or received during this state. • LEARNING – The port has passed through the LISTENING state and will change to the BLOCKING or FORWARDING state, depending on the results of STP's reconvergence. The port does not transmit or receive user frames during this state. However, the device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table.
Fwd Trans	<p>The number of times STP has changed the state of this port between BLOCKING and FORWARDING.</p>
Design Cost	<p>The cost to the root bridge as advertised by the designated bridge that is connected to this port. If the designated bridge is the root bridge itself, then the cost is 0. The identity of the designated bridge is shown in the Design Bridge field.</p>
Design Root	<p>The root bridge as recognized on this port. The value is the same as the root bridge ID listed in the Root ID field.</p>
Design Bridge	<p>The designated bridge to which this port is connected. The designated bridge is the device that connects the network segment on the port to the root bridge.</p>

USING THE WEB MANAGEMENT INTERFACE

To display STP information:

1. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to display the monitoring options.
3. Select the STP link to display the STP bridge and port parameters.

Table 10.5: Web Management Display of STP Information

This Field...	Displays...
STP Bridge Parameters (global parameters)	
VLAN ID	The port-based VLAN that contains this spanning tree (instance of STP). VLAN 1 is the default VLAN. If you have not configured port-based VLANs on this device, all STP information is for VLAN 1.
Root ID	The ID assigned by STP to the root bridge for this spanning tree.
Root Cost	The cumulative cost from this bridge to the root bridge. If this device is the root bridge, then the root cost is 0.
Root Port	The port on this device that connects to the root bridge. If this device is the root bridge, then the value is "Root" instead of a port number.
Priority	This device or VLAN's STP priority. The value is shown in hexadecimal format. Note: If you configure this value, specify it in decimal format. See "Changing STP Bridge Parameters" on page 10-5.
Max Age	The number of seconds this device or VLAN waits for a hello message from the root bridge before deciding the root has become unavailable and performing a reconvergence.
Hello Time	The interval between each configuration BPDU sent by the root bridge.
Hold Time	The minimum number of seconds that must elapse between transmissions of consecutive Configuration BPDUs on a port.
Forward Delay	The number of seconds this device or VLAN waits following a topology change and consequent reconvergence.
Topology Last Change	The number of seconds since the last time a topology change occurred.
Topology Change Counter	The number of times the topology has changed since this device was reloaded.
Bridge Address	The STP address of this device or VLAN. Note: If this address is the same as the Root ID, then this device or VLAN is the root bridge for its spanning tree.
STP Port Parameters	
VLAN	The VLAN that the port is in.
Port	The port number.
Priority	The port's STP priority, in hexadecimal format. Note: If you configure this value, specify it in decimal format. See "Changing STP Port Parameters" on page 10-6.
Path Cost	The port's STP path cost.

Table 10.5: Web Management Display of STP Information (Continued)

This Field...	Displays...
State	<p>The port's STP state. The state can be one of the following:</p> <ul style="list-style-type: none"> • BLOCKING – STP has blocked Layer 2 traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is FORWARDING. When a port is in this state, the port does not transmit or receive user frames, but the port does continue to receive STP BPDUs. • DISABLED – The port is not participating in STP. This can occur when the port is disconnected or STP is disabled on the port. • FORWARDING – STP is allowing the port to send and receive frames. • LISTENING – STP is responding to a topology change and this port is listening for a BPDU from neighboring bridge(s) in order to determine the new topology. No user frames are transmitted or received during this state. • LEARNING – The port has passed through the LISTENING state and will change to the BLOCKING or FORWARDING state, depending on the results of STP's reconvergence. The port does not transmit or receive user frames during this state. However, the device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table.
Transition	<p>The number of times STP has changed the state of this port between BLOCKING and FORWARDING.</p>
Cost	<p>The cost to the root bridge as advertised by the designated bridge that is connected to this port. If the designated bridge is the root bridge itself, then the cost is 0. The identity of the designated bridge is shown in the Design Bridge field.</p>
Root	<p>The root bridge as recognized on this port. The value is the same as the root bridge ID listed in the Root ID field.</p>
Bridge	<p>The designated bridge to which this port is connected. The designated bridge is the device that connects the network segment on the port to the root bridge.</p>

Displaying CPU Utilization Statistics

You can display CPU utilization statistics for STP and the IP protocols.

USING THE CLI

To display CPU utilization statistics for STP for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI:

```
BigIron# show process cpu
Process Name   5Sec(%)   1Min(%)   5Min(%)   15Min(%)   Runtime(ms)
ARP            0.01      0.03      0.09      0.22       9
BGP            0.04      0.06      0.08      0.14      13
GVRP           0.00      0.00      0.00      0.00       0
ICMP           0.00      0.00      0.00      0.00       0
IP             0.00      0.00      0.00      0.00       0
OSPF           0.00      0.00      0.00      0.00       0
RIP            0.00      0.00      0.00      0.00       0
STP          0.00    0.03    0.04    0.07     4
VRRP           0.00      0.00      0.00      0.00       0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example:

```
BigIron# show process cpu
The system has only been up for 6 seconds.
Process Name   5Sec(%)   1Min(%)   5Min(%)   15Min(%)   Runtime(ms)
ARP            0.01      0.00      0.00      0.00       0
BGP            0.00      0.00      0.00      0.00       0
GVRP           0.00      0.00      0.00      0.00       0
ICMP           0.01      0.00      0.00      0.00       1
IP             0.00      0.00      0.00      0.00       0
OSPF           0.00      0.00      0.00      0.00       0
RIP            0.00      0.00      0.00      0.00       0
STP            0.00      0.00      0.00      0.00       0
VRRP           0.00      0.00      0.00      0.00       0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following:

```
BigIron# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name   Sec(%)   Time(ms)
ARP            0.00     0
BGP            0.00     0
GVRP           0.00     0
ICMP           0.01     1
IP             0.00     0
OSPF           0.00     0
RIP            0.00     0
STP            0.01     0
VRRP           0.00     0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: show process cpu [<num>]

The <num> parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

USING THE WEB MANAGEMENT INTERFACE

You cannot display this information using the Web management interface.

Displaying the STP State of a Port-Based VLAN

When you display information for a port-based VLAN, that information includes the STP state of the VLAN. Use either of the following methods to display port-based VLAN information.

USING THE CLI

To display information for a port-based VLAN, enter a command such as the following at any level of the CLI. The STP state is shown in bold type in this example.

```
BigIron(config)# show vlans

Total PORT-VLAN entries: 2
Maximum PORT-VLAN entries: 16

legend: [S=Slot]

PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree On
Untagged Ports: (S3) 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
Untagged Ports: (S3) 17 18 19 20 21 22 23 24
Untagged Ports: (S4) 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17
Untagged Ports: (S4) 18 19 20 21 22 23 24
    Tagged Ports: None
    Uplink Ports: None

PORT-VLAN 2, Name greenwell, Priority level0, Spanning tree Off
Untagged Ports: (S1) 1 2 3 4 5 6 7 8
Untagged Ports: (S4) 1
    Tagged Ports: None
    Uplink Ports: None
```

Syntax: show vlans [*<vlan-id>* | ethernet *<portnum>* | pos *<portnum>*]

The *<vlan-id>* parameter specifies a VLAN for which you want to display the configuration information.

The **ethernet** *<portnum>* | **pos** *<portnum>* parameter specifies a port. If you use this parameter, the command lists all the VLAN memberships for the port.

USING THE WEB MANAGEMENT INTERFACE

To display STP information for a specific VLAN:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view.
3. Click on the plus sign next to VLAN in the tree view
4. Select the Port link to display configuration information for the device's port-based VLANs. The STP state is shown in the STP column.

Displaying Detailed STP Information for Each Interface

To display detailed STP information for individual ports, use the following CLI method.

USING THE CLI

To display the detailed STP information, enter the following command at any level of the CLI:

```
BigIron# show span detail
=====
VLAN 1 - MULTIPLE SPANNING TREE (MSTP) ACTIVE
=====
Bridge identifier      - 0x800000e0804d4a00
Active global timers - Hello: 0

Port 1/1 is FORWARDING
  Port - Path cost: 19, Priority: 128, Root: 0x800000e052a9bb00
  Designated - Bridge: 0x800000e052a9bb00, Interface: 1, Path cost: 0
  Active Timers - None
  BPDUs - Sent: 11, Received: 0
Port 1/2 is DISABLED
Port 1/3 is DISABLED
Port 1/4 is DISABLED
<lines for remaining ports excluded for brevity>
```

If a port is disabled, the only information shown by this command is “DISABLED”. If a port is enabled, this display shows the following information.

Syntax: show span detail [vlan <vlan-id> [atm <portnum> | ethernet <portnum> | pos <portnum>] | <num>]

The **vlan** <vlan-id> parameter specifies a VLAN.

The **atm** <portnum> | **ethernet** <portnum> | **pos** <portnum> parameter specifies an individual port within the VLAN (if specified).

The <num> parameter specifies the number of VLANs you want the CLI to skip before displaying detailed STP information. For example, if the device has six VLANs configured (VLAN IDs 1, 2, 3, 99, 128, and 256) and you enter the command **show span detail 4**, detailed STP information is displayed for VLANs 128 and 256 only.

NOTE: If the configuration includes VLAN groups, the **show span detail** command displays the master VLANs of each group but not the member VLANs within the groups. However, the command does indicate that the VLAN is a master VLAN. The **show span detail vlan** <vlan-id> command displays the information for the VLAN even if it is a member VLAN. To list all the member VLANs within a VLAN group, enter the **show vlan-group** [<group-id>] command.

The **show span detail** command shows the following information.

Table 10.6: CLI Display of Detailed STP Information for Ports

This Field...	Displays...
Active Spanning Tree protocol	<p>The VLAN that contains the listed ports and the active Spanning Tree protocol.</p> <p>The STP type can be one of the following:</p> <ul style="list-style-type: none"> MULTIPLE SPANNNG TREE (MSTP) GLOBAL SINGLE SPANNING TREE (SSTP) <p>Note: If STP is disabled on a VLAN, the command displays the following message instead: “Spanning-tree of port-vlan <vlan-id> is disabled.”</p>

Table 10.6: CLI Display of Detailed STP Information for Ports (Continued)

This Field...	Displays...
Bridge identifier	The STP identity of this device.
Active global timers	<p>The global STP timers that are currently active, and their current values. The following timers can be listed:</p> <ul style="list-style-type: none"> • Hello – The interval between Hello packets. This timer applies only to the root bridge. • Topology Change (TC) – The amount of time during which the topology change flag in Hello packets will be marked, indicating a topology change. This timer applies only to the root bridge. • Topology Change Notification (TCN) – The interval between Topology Change Notification packets sent by a non-root bridge toward the root bridge. This timer applies only to non-root bridges.
Port number and STP state	<p>The internal port number and the port's STP state.</p> <p>The internal port number is one of the following:</p> <ul style="list-style-type: none"> • The port's interface number, if the port is the designated port for the LAN. • The interface number of the designated port from the received BPDU, if the interface is not the designated port for the LAN. <p>The state can be one of the following:</p> <ul style="list-style-type: none"> • BLOCKING – STP has blocked Layer 2 traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is FORWARDING. When a port is in this state, the port does not transmit or receive user frames, but the port does continue to receive STP BPDUs. • DISABLED – The port is not participating in STP. This can occur when the port is disconnected or STP is administratively disabled on the port. • FORWARDING – STP is allowing the port to send and receive frames. • LISTENING – STP is responding to a topology change and this port is listening for a BPDU from neighboring bridge(s) in order to determine the new topology. No user frames are transmitted or received during this state. • LEARNING – The port has passed through the LISTENING state and will change to the BLOCKING or FORWARDING state, depending on the results of STP's reconvergence. The port does not transmit or receive user frames during this state. However, the device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table. <p>Note: If the state is DISABLED, no further STP information is displayed for the port.</p>
Port Path cost	The port's STP path cost.

Table 10.6: CLI Display of Detailed STP Information for Ports (Continued)

This Field...	Displays...
Port Priority	This port's STP priority. The value is shown as a hexadecimal number.
Root	The ID assigned by STP to the root bridge for this spanning tree.
Designated Bridge	The MAC address of the designated bridge to which this port is connected. The designated bridge is the device that connects the network segment on the port to the root bridge.
Designated Port	The port number sent from the designated bridge.
Designated Path Cost	The cost to the root bridge as advertised by the designated bridge that is connected to this port. If the bridge is the root bridge itself, then the cost is 0. The identity of the designated bridge is shown in the Designated Bridge field.
Active Timers	The current values for the following timers, if active: <ul style="list-style-type: none"> • Message age – The number of seconds this port has been waiting for a hello message from the root bridge. • Forward delay – The number of seconds that have passed since the last topology change and consequent reconvergence. • Hold time – The number of seconds that have elapsed since transmission of the last Configuration BPDU.
BPDU's Sent and Received	The number of BPDU's sent and received on this port since the software was reloaded.

Displaying Detailed STP Information for a Single Port in a Specific VLAN

Enter a command such as the following to display STP information for an individual port in a specific VLAN.

```
BigIron(config)# show span detail vlan 1 ethernet 7/1
Port 7/1 is FORWARDING
  Port - Path cost: 19, Priority: 128, Root: 0x800000e052a9bb00
  Designated - Bridge: 0x800000e052a9bb00, Interface: 7, Path cost: 0
  Active Timers - None
  BPDU's - Sent: 29, Received: 0
```

Syntax: show span detail [vlan <vlan-id> [atm <portnum> | ethernet <portnum> | pos <portnum>] | <num>]

USING THE WEB MANAGEMENT INTERFACE

The detailed display is not supported in the Web management interface.

Displaying STP State Information for an Individual Interface

To display STP state information for an individual port, you can use the methods in “Displaying STP Information for an Entire Device” on page 10-8 or “Displaying Detailed STP Information for Each Interface”. You also can display STP state information for a specific port using either of the following methods.

USING THE CLI

To display information for a specific port, enter a command such as the following at any level of the CLI:

```
BigIron(config)# show interface ethernet 3/11

FastEthernet3/11 is up, line protocol is up
  Hardware is FastEthernet, address is 00e0.52a9.bb49 (bia 00e0.52a9.bb49)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Member of L2 VLAN ID 1, port is untagged, port state is FORWARDING
  STP configured to ON, priority is level0, flow control enabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  MTU 1518 bytes, encapsulation ethernet
  5 minute input rate: 352 bits/sec, 0 packets/sec, 0.00% utilization
  5 minute output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  1238 packets input, 79232 bytes, 0 no buffer
  Received 686 broadcasts, 0 runts, 0 giants
  0 input errors, 0 CRC, 0 frame, 0 ignored
  529 multicast
  918 packets output, 63766 bytes, 0 underruns
  0 output errors, 0 collisions
```

The STP information is shown in bold type in this example.

Syntax: show interfaces [ethernet | pos <portnum>] | [loopback <num>] | [slot <slot-num>] | [ve <num>] | [brief]

You also can display the STP states of all ports by entering a command such as the following, which uses the **brief** parameter:

```
BigIron(config)# show interface brief

Port  Link State      Dupl Speed Trunk Tag Priori MAC           Name
1/1   Down None         None None  None No  level0 00e0.52a9.bb00
1/2   Down None         None None  None No  level0 00e0.52a9.bb01
1/3   Down None         None None  None No  level0 00e0.52a9.bb02
1/4   Down None         None None  None No  level0 00e0.52a9.bb03
1/5   Down None         None None  None No  level0 00e0.52a9.bb04
1/6   Down None         None None  None No  level0 00e0.52a9.bb05
1/7   Down None         None None  None No  level0 00e0.52a9.bb06
1/8   Down None         None None  None No  level0 00e0.52a9.bb07

.
.  some rows omitted for brevity
.
3/10  Down None         None None  None No  level0 00e0.52a9.bb4a
3/11  Up   Forward      Full 100M  None No  level0 00e0.52a9.bb49
```

In this example, only one port, 3/11, is forwarding traffic toward the root bridge.

USING THE WEB MANAGEMENT INTERFACE

To display STP information for a specific port, use the same method as the one described in “Displaying STP Information for an Entire Device” on page 10-8:

1. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to display the monitoring options.
3. Select the STP link to display the STP bridge and port parameters.

Configuring IronSpan Features

This section describes how to configure the following features:

- Fast Port Span – see “Fast Port Span”
- Fast Uplink Span – see “Fast Uplink Span” on page 10-21
- 802.1W Rapid Spanning Tree (RSTP) – see “802.1W Rapid Spanning Tree (RSTP)” on page 10-22
- 802.1W Draft 3 RSTP – see “802.1W Draft 3” on page 10-58
- Single-instance STP – see “Single Spanning Tree (SSTP)” on page 10-62
- SuperSpan™ – see “SuperSpan™” on page 10-65
- STP per VLAN group – see “STP per VLAN Group” on page 10-71
- Per VLAN Spanning Tree+ (PVST+) Compatibility – see “PVST/PVST+ Compatibility” on page 10-75

Fast Port Span

When STP is running on a device, message forwarding is delayed during the spanning tree recalculation period following a topology change. The STP forward delay parameter specifies the period of time a bridge waits before forwarding data packets. The forward delay controls the listening and learning periods of STP convergence. You can configure the forward delay to a value from 4 – 30 seconds. The default is 15 seconds. Thus, using the standard forward delay, convergence requires 30 seconds (15 seconds for listening and an additional 15 seconds for learning) when the default value is used.

This slow convergence is undesirable and unnecessary in some circumstances. The Fast Port Span feature allows certain ports to enter the forwarding state in four seconds. Specifically, Fast Port Span allows faster convergence on ports that are attached to end stations and thus do not present the potential to cause Layer 2 forwarding loops. Because the end stations cannot cause forwarding loops, they can safely go through the STP state changes (blocking to listening to learning to forwarding) more quickly than is allowed by the standard STP convergence time. Fast Port Span performs the convergence on these ports in four seconds (two seconds for listening and two seconds for learning).

In addition, Fast Port Span enhances overall network performance in the following ways:

- Fast Port Span reduces the number of STP topology change notifications on the network. When an end station attached to a Fast Span port comes up or down, the Foundry device does not generate a topology change notification for the port. In this situation, the notification is unnecessary since a change in the state of the host does not affect the network's topology.
- Fast Port Span eliminates unnecessary MAC cache aging that can be caused by topology change notifications. Bridging devices age out the learned MAC addresses in their MAC caches if the addresses are unrefreshed for a given period of time, sometimes called the MAC aging interval. When STP sends a topology change notification, devices that receive the notification use the value of the STP forward delay to quickly age out their MAC caches. For example, if a device's normal MAC aging interval is 5 minutes, the aging interval changes temporarily to the value of the forward delay (for example, 15 seconds) in response to an STP topology change.

In normal STP, the accelerated cache aging occurs even when a single host goes up or down. Because Fast Port Span does not send a topology change notification when a host on a Fast Port Span port goes up or down, the unnecessary cache aging that can occur in these circumstances under normal STP is eliminated.

Fast Port Span is a system-wide parameter and is enabled by default. Thus, when you boot a device with software release 06.0.00 or later, all the ports that are attached only to end stations run Fast Port Span. For ports that are

not eligible for Fast Port Span, such as ports connected to other networking devices, the device automatically uses the normal STP settings. If a port matches any of the following criteria, the port is ineligible for Fast Port Span and uses normal STP instead:

- The port is 802.1q tagged
- The port is a member of a trunk group
- The port has learned more than one active MAC address
- An STP Configuration BPDU has been received on the port, thus indicating the presence of another bridge on the port.

You also can explicitly exclude individual ports from Fast Port Span if needed. For example, if the only uplink ports for a wiring closet switch are Gigabit ports, you can exclude the ports from Fast Port Span.

Disabling and Re-enabling Fast Port Span

Fast Port Span is a system-wide parameter and is enabled by default. Thus all ports that are eligible for Fast Port Span use it.

To disable or re-enable Fast Port Span, use one of the following methods.

USING THE CLI

To disable Fast Port Span, enter the following commands:

```
BigIron(config)# no fast port-span
BigIron(config)# write memory
```

Syntax: [no] fast port-span

NOTE: The **fast port-span** command has additional parameters that let you exclude specific ports. These parameters are shown in the following section.

To re-enable Fast Port Span, enter the following commands:

```
BigIron(config)# fast port-span
BigIron(config)# write memory
```

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access.
2. Click the Fast checkbox next to Spanning Tree to remove the checkmark from the box.
3. Click Apply to apply the change to the device's running-config.
4. Select the Save link at the bottom of the panel. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Excluding Specific Ports from Fast Port Span

You can exclude individual ports from Fast Port Span while leaving Fast Port Span enabled globally. To do so, use one of the following methods.

USING THE CLI

To exclude a port from Fast Port Span, enter commands such as the following:

```
BigIron(config)# fast port-span exclude ethernet 1/1
BigIron(config)# write memory
```

To exclude a set of ports from Fast Port Span, enter commands such as the following:

```
BigIron(config)# fast port-span exclude ethernet 1/1 ethernet 2/1 ethernet 3/2
BigIron(config)# write memory
```

To exclude a contiguous (unbroken) range of ports from Fast Span, enter commands such as the following:

```
BigIron(config)# fast port-span exclude ethernet 1/1 to 1/24
```

```
BigIron(config)# write memory
```

Syntax: [no] fast port-span [exclude ethernet <portnum> [ethernet <portnum>... | to <portnum>]]

To re-enable Fast Port Span on a port, enter a command such as the following:

```
BigIron(config)# no fast port-span exclude ethernet 1/1
BigIron(config)# write memory
```

This command re-enables Fast Port Span on port 1/1 only and does not re-enable Fast Port Span on other excluded ports. You also can re-enable Fast Port Span on a list or range of ports using the syntax shown above this example.

To re-enable Fast Port Span on all excluded ports, disable and then re-enable Fast Port Span by entering the following commands:

```
BigIron(config)# no fast port-span
BigIron(config)# fast port-span
BigIron(config)# write memory
```

Disabling and then re-enabling Fast Port Span clears the exclude settings and thus enables Fast Port Span on all eligible ports. To make sure Fast Port Span remains enabled on the ports following a system reset, save the configuration changes to the startup-config file after you re-enable Fast Port Span. Otherwise, when the system resets, those ports will again be excluded from Fast Port Span.

[USING THE WEB MANAGEMENT INTERFACE](#)

You cannot exclude individual ports from Fast Span using the Web management interface.

Fast Uplink Span

The Fast Port Span feature described in the previous section enhances STP performance for end stations. The Fast Uplink feature enhances STP performance for wiring closet switches with redundant uplinks. Using the default value for the standard STP forward delay, convergence following a transition from an active link to a redundant link can take 30 seconds (15 seconds for listening and an additional 15 seconds for learning).

You can use the Fast Uplink feature on a Foundry device deployed as a wiring closet switch to decrease the convergence time for the uplink ports to another device to just four seconds (two seconds for listening and two seconds for learning). The wiring closet switch must be a Foundry device but the device at the other end of the link can be a Foundry device or another vendor's switch. Configuration of the Fast Uplink Span feature takes place entirely on the Foundry device.

To configure the Fast Uplink Span feature, specify a group of ports that have redundant uplinks on the wiring closet switch (Foundry device) as members of a Fast Uplink Group. If the active link becomes unavailable, the Fast Uplink Span feature transitions the forwarding to one of the other ports in four seconds. You can configure one Fast Uplink Span group on the device. All Fast Uplink Span ports are members of the same Fast Uplink Span group.

NOTE: To avoid the potential for temporary bridging loops, Foundry Networks recommends that you use the Fast Uplink feature only for wiring closet switches (switches at the edge of the network cloud). In addition, enable the feature only on a group of ports intended for redundancy, so that at any given time only one of the ports is expected to be in the forwarding state.

NOTE: When the wiring closet switch (Foundry device) first comes up or when STP is first enabled, the uplink ports still must go through the standard STP state transition without any acceleration. This behavior guards against temporary routing loops as the switch tries to determine the states for all the ports. Fast Uplink Span acceleration applies only when a working uplink becomes unavailable.

Fast Uplink Span Rules for Trunk Groups

If you add a port to a Fast Uplink Span group that is a member of a trunk group, the following rules apply:

- If you add the primary port of a trunk group to the Fast Uplink Span group, all other ports in the trunk group are automatically included in the group. Similarly, if you remove the primary port in a trunk group from the

Fast Uplink Span group, the other ports in the trunk group are automatically removed from the Fast Uplink Span group.

- You cannot add a subset of the ports in a trunk group to the Fast Uplink Span group. All ports in a trunk group have the same Fast Uplink Span property, as they do for other port properties.
- If the working trunk group is partially down but not completely down, no switch-over to the backup occurs. This behavior is the same as in the standard STP feature.
- If the working trunk group is completely down, a backup trunk group can go through an accelerated transition only if the following are true:
 - The trunk group is included in the fast uplink group.
 - All other ports except those in this trunk group are either disabled or blocked. The accelerated transition applies to all ports in this trunk group.
- When the original working trunk group comes back (partially or fully), the transition back to the original topology is accelerated if the conditions listed above are met.

Configuring a Fast Uplink Port Group

To enable Fast Uplink, use one of the following methods.

USING THE CLI

To configure a group of ports for Fast Uplink Span, enter the following commands:

```
BigIron(config)# fast uplink-span ethernet 4/1 to 4/4
BigIron(config)# write memory
```

Syntax: [no] fast uplink-span [ethernet <portnum> [ethernet <portnum>... | to <portnum>]]

This example configures four ports, 4/1 – 4/4, as a Fast Uplink Span group. In this example, all four ports are connected to a wiring closet switch. Only one of the links is expected to be active at any time. The other links are redundant. For example, if the link on port 4/1 is the active link on the wiring closet switch but becomes unavailable, one of the other links takes over. Because the ports are configured in a Fast Uplink Span group, the STP convergence takes about four seconds instead of taking 30 seconds or longer using the standard STP forward delay.

If you add a port that is the primary port of a trunk group, all ports in the trunk group become members of the Fast Uplink Span group.

You can add ports to a Fast Uplink Span group by entering the **fast uplink-span** command additional times with additional ports. The device can have only one Fast Uplink Span group, so all the ports you identify as Fast Uplink Span ports are members of the same group.

To remove a Fast Uplink Span group or to remove individual ports from a group, use “no” in front of the appropriate **fast uplink-span** command. For example, to remove ports 4/3 and 4/4 from the Fast Uplink Span group configured above, enter the following commands:

```
BigIron(config)# no fast uplink-span ethernet 4/3 to 4/4
BigIron(config)# write memory
```

If you delete a port that is the primary port of a trunk group, all ports in the trunk group are removed from the Fast Uplink Span group.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure the Fast Uplink Span feature using the Web management interface.

802.1W Rapid Spanning Tree (RSTP)

Foundry’s earlier implementation of Rapid Spanning Tree Protocol (RSTP), which was 802.1W Draft 3, provided only a subset of the IEEE 802.1W standard; whereas the 802.1W RSTP feature provides the full standard. The implementation of the 802.1W Draft 3 is referred to as RSTP Draft 3.

RSTP Draft3 will continue to be supported on Foundry devices for backward compatibility. However, customers who are currently using RSTP Draft 3 should migrate to 802.1W.

The 802.1W feature is supported on all s, as well as on the FastIron 4802. It provides rapid traffic reconvergence for point-to-point links within a few milliseconds (0 – 500 milliseconds), following the failure of a bridge or bridge port. This reconvergence occurs more rapidly than the reconvergence provided by the 802.1D (Spanning Tree Protocol (STP)) or by RSTP Draft 3.

NOTE: This rapid convergence will not occur on ports connected to shared media devices, such as hubs. To take advantage of the rapid convergence provided by 802.1W, make sure to explicitly configure all point-to-point links in a topology.

The convergence provided by the standard 802.1W protocol occurs more rapidly than the convergence provided by previous spanning tree protocols because:

- Classic or legacy 802.1D STP protocol requires a newly selected Root port to go through listening and learning stages before traffic convergence can be achieved. The 802.1D traffic convergence time is calculated using the following formula:

$$2 \times FORWARD_DELAY + BRIDGE_MAX_AGE.$$

If default values are used in the parameter configuration, convergence can take up to 50 seconds. (In this document STP will be referred to as 802.1D.)

- RSTP Draft 3 works only on bridges that have Alternate ports, which are the precalculated “next best root port”. (Alternate ports provide back up paths to the root bridge.) Although convergence occurs from 0 – 500 milliseconds in RSTP Draft 3, the spanning tree topology reverts to the 802.1D convergence if an Alternate port is not found.
- Convergence in 802.1w bridge is not based on any timer values. Rather, it is based on the explicit handshakes between Designated ports and their connected Root ports to achieve convergence in less than 500 milliseconds.

Bridges and Bridge Port Roles

A bridge in an 802.1W rapid spanning tree topology is assigned as the root bridge if it has the highest priority (lowest bridge identifier) in the topology. Other bridges are referred to as non-root bridges.

Unique roles are assigned to ports on the root and non-root bridges. Role assignments are based on the following information contained in the Rapid Spanning Tree Bridge Packet Data Unit (RST BPDU):

- Root bridge ID
- Path cost value
- Transmitting bridge ID
- Designated port ID

802.1W algorithm uses this information to determine if the RST BPDU received by a port is superior to the RST BPDU that the port transmits. The two values are compared in the order as given above, starting with the Root bridge ID. The RST BPDU with a lower value is considered superior. The superiority and inferiority of the RST BPDU is used to assign a role to a port.

If the value of the received RST BPDU is the same as that of the transmitted RST BPDU, then the port ID in the RST BPDUs are compared. The RST BPDU with the lower port ID is superior. Port roles are then calculated appropriately.

The port's role is included in the BPDU that it transmits. The BPDU transmitted by an 802.1W port is referred to as an RST BPDU, while it is operating in 802.1W mode.

Ports can have one of the following roles:

- Root – Provides the lowest cost path to the root bridge from a specific bridge
- Designated – Provides the lowest cost path to the root bridge from a LAN to which it is connected
- Alternate – Provides an alternate path to the root bridge when the root port goes down
- Backup – Provides a backup to the LAN when the Designated port goes down

- Disabled – Has no role in the topology

Assignment of Port Roles

At system start-up, all 802.1W-enabled bridge ports assume a Designated role. Once start-up is complete, 802.1W algorithm calculates the superiority or inferiority of the RST BPDU that is received and transmitted on a port.

On a root bridge, each port is assigned a **Designated port** role, except for ports on the same bridge that are physically connected together. In these type of ports, the port that receives the superior RST BPDU becomes the **Backup port**, while the other port becomes the **Designated port**.

On non-root bridges, ports are assigned as follows:


- The port that receives the RST BPDU with the lowest path cost from the root bridge becomes the **Root port**.
- If two ports on the same bridge are physically connected, the port that receives the superior RST BPDU becomes the **Backup port**, while the other port becomes the **Designated port**.
- If a non-root bridge already has a Root port, then the port that receives an RST BPDU that is superior to those it can transmit becomes the **Alternate port**.
- If the RST BPDU that a port receives is inferior to the RST BPDUs it transmits, then the port becomes a **Designated port**.
- If the port is down or if 802.1W is disabled on the port, that port is given the role of **Disabled port**. Disabled ports have no role in the topology. However, if 802.1W is enabled on a port with a link down and the link of that port comes up, then that port assumes one of the following port roles: Root, Designated, Alternate, or Backup.

The following example (Figure 10.1) explains role assignments in a simple RSTP topology.

NOTE: All examples in this document assume that all ports in the illustrated topologies are point-to-point links and are homogeneous (they have the same path cost value) unless otherwise specified.

The topology in Figure 10.1 contains four bridges. Switch 1 is the root bridge since it has the lowest bridge priority. Switch 2 through Switch 4 are non-root bridges.

Figure 10.1 Simple 802.1W Topology



Ports on Switch 1

All ports on Switch 1, the root bridge, are assigned Designated port roles.

Ports on Switch 2

Port2 on Switch 2 directly connects to the root bridge; therefore, Port2 is the Root port.

Switch 2's bridge priority value is superior to that of Switch 3 and Switch 4; therefore, the ports on Switch 2 that connect to Switch 3 and Switch 4 are given the Designated port role.

Furthermore, Port7 and Port8 on Switch 2 are physically connected. The RST BPDUs transmitted by Port7 are superior to those Port8 transmits. Therefore, Switch 2 is the Backup port and Port7 is the Designated port.

Ports on Switch 3

Port2 on Switch 3 directly connects to the Designated port on the root bridge; therefore, it assumes the Root port role.

The root path cost of the RST BPDUs received on Port4/Switch 3 is inferior to the RST BPDUs transmitted by the port; therefore, Port4/Switch 3 becomes the Designated port.

Similarly Switch 3 has a bridge priority value inferior to Switch 2. Port3 on Switch 3 connects to Port 3 on Switch 2. This port will be given the Alternate port role, since a Root port is already established on this bridge.

Ports Switch 4


Switch 4 is not directly connected to the root bridge. It has two ports with superior incoming RST BPDUs from two separate LANs: Port3 and Port4. The RST BPDUs received on Port3 are superior to the RST BPDUs received on port 4; therefore, Port3 becomes the Root port and Port4 becomes the Alternate port.

Edge Ports and Edge Port Roles

Foundry's implementation of 802.1W allows ports that are configured as Edge ports to be present in an 802.1W topology. (Figure 10.2). Edge ports are ports of a bridge that connect to workstations or computers. Edge ports do not register any incoming BPDU activities.

Edge ports assume Designated port roles. Port flapping does not cause any topology change events on Edge ports since 802.1W does not consider Edge ports in the spanning tree calculations.

Figure 10.2 Topology with Edge Ports



However, if any incoming RST BPDU is received from a previously configured Edge port, 802.1W automatically makes the port as a non-edge port. This is extremely important to ensure a loop free Layer 2 operation since a non-edge port is part of the active RSTP topology.

The 802.1W protocol can auto-detect an Edge port and a non-edge port. An administrator can also configure a port to be an Edge port using the CLI. It is recommended that Edge ports are configured explicitly to take advantage of the Edge port feature, instead of allowing the protocol to auto-detect them.


Point-to-Point Ports

To take advantage of the 802.1W features, ports on an 802.1W topology should be explicitly configured as point-to-point links using the CLI. Shared media should not be configured as point-to-point links.

NOTE: Configuring shared media or non-point-to-point links as point-to-point links could lead to Layer 2 loops.

The topology in Figure 10.3 is an example of shared media that should not be configured as point-to-point links. In Figure 10.3, a port on a bridge communicates or is connected to at least two ports.

Figure 10.3 Example of Shared Media



Bridge Port States

Ports roles can have one of the following states:

- Forwarding – 802.1W is allowing the port to send and receive all packets.
- Discarding – 802.1W has blocked data traffic on this port to prevent a loop. The device or VLAN can reach the root bridge using another port, whose state is forwarding. When a port is in this state, the port does not transmit or receive data frames, but the port does continue to receive RST BPDUs. This state corresponds to the listening and blocking states of 802.1D.
- Learning – 802.1W is allowing MAC entries to be added to the filtering database but does not permit forwarding of data frames. The device can learn the MAC addresses of frames that the port receives during this state and make corresponding entries in the MAC table.
- Disabled – The port is not participating in 802.1W. This can occur when the port is disconnected or 802.1W is administratively disabled on the port.

A port on a non-root bridge with the role of Root port is always in a forwarding state. If another port on that bridge assumes the Root port role, then the old Root port moves into a discarding state as it assumes another port role.

A port on a non-root bridge with a Designated role starts in the discarding state. When that port becomes elected to the Root port role, 802.1W quickly places it into a forwarding state. However, if the Designated port is an Edge port, then the port starts and stays in a forwarding state and it cannot be elected as a Root port.

A port with an Alternate or Backup role is always in a discarding state. If the port's role changes to Designated, then the port changes into a forwarding state.

If a port on one bridge has a Designated role and that port is connected to a port on another bridge that has an Alternate or Backup role, the port with a Designated role cannot be given a Root port role until two instances of the forward delay timer expires on that port.

Edge Port and Non-Edge Port States

As soon as a port is configured as an Edge port using the CLI, it goes into a forwarding state instantly (in less than 100 msec):

When the link to a port comes up and 802.1W detects that the port is an Edge port, that port instantly goes into a forwarding state.

If 802.1W detects that port as a non-edge port, the port goes into a forwarding state within four seconds of link up or after two hello timer expires on the port.

Changes to Port Roles and States

To achieve convergence in a topology, a port's role and state changes as it receives and transmits new RST BPDUs. Changes in a port's role and state constitute a topology change. Besides the superiority and inferiority of the RST BPDU, bridge-wide and per-port state machines are used to determine a port's role as well as a port's state. Port state machines also determine when port role and state changes occur.

State Machines

The bridge uses the Port Role Selection state machine to determine if port role changes are required on the bridge. This state machine performs a computation when one of the following events occur:

- New information is received on any port on the bridge
- The timer expires for the current information on a port on the bridge

Each port uses the following state machines:

- Port Information – This state machine keeps track of spanning-tree information currently used by the port. It records the origin of the information and ages out any information that was derived from an incoming BPDU.
- Port Role Transition – This state machine keeps track of the current port role and transitions the port to the appropriate role when required. It moves the Root port and the Designated port into forwarding states and moves the Alternate and Backup ports into discarding states.
- Port Transmit – This state machine is responsible for BPDU transmission. It checks to ensure only the

maximum number of BPDUs per hello interval are sent every second. Based on what mode it is operating in, it sends out either legacy BPDUs or RST BPDUs. In this document legacy BPDUs are also referred to as STP BPDUs.

- Port Protocol Migration – This state machine deals with compatibility with 802.1D bridges. When a legacy BPDU is detected on a port, this state machine configures the port to transmit and receive legacy BPDUs and operate in the legacy mode.
- Topology Change – This state machine detects, generates, and propagates topology change notifications. It acknowledges Topology Change Notice (TCN) messages when operating in 802.1D mode. It also flushes the MAC table when a topology change event takes place.
- Port State Transition – This state machine transitions the port to a discarding, learning, or forwarding state and performs any necessary processing associated with the state changes.
- Port Timers – This state machine is responsible for triggering any of the state machines described above, based on expiration of specific port timers.

In contrast to the 802.1D standard, the 802.1W standard does not have any bridge specific timers. All timers in the CLI are applied on a per-port basis, even though they are configured under bridge parameters.

802.1W state machines attempt to quickly place the ports into either a forwarding or discarding state. Root ports are quickly placed in forwarding state when both of the following events occur:

- It is assigned to be the Root port.
- It receives an RST BPDU with a proposal flag from a Designated port. The proposal flag is sent by ports with a Designated role when they are ready to move into a forwarding state.

When a the role of Root port is given to another port, the old Root port is instructed to reroot. The old Root port goes into a discarding state and negotiates with its peer port for a new role and a new state. A peer port is the port on the other bridge to which the port is connected. For example, in Figure 10.4, Port1 of Switch 200 is the peer port of Port2 of Switch 100.

A port with a Designated role is quickly placed into a forwarding state if one of the following occurs:

- The Designated port receives an RST BPDU that contains an agreement flag from a Root port
- The Designated port is an Edge port

However, a Designated port that is attached to an Alternate port or a Backup port must wait until the forward delay timer expires twice on that port while it is still in a Designated role, before it can proceed to the forwarding state.

Backup ports are quickly placed into discarding states.

Alternate ports are quickly placed into discarding states.

A port operating in 802.1W mode may enter a learning state to allow MAC entries to be added to the filtering database; however, this state is transient and lasts only a few milliseconds, if the port is operating in 802.1W mode and if the port meets the conditions for rapid transition.

Handshake Mechanisms

To rapidly transition a Designated or Root port into a forwarding state, the Port Role Transition state machine uses handshake mechanisms to ensure loop free operations. It uses one type of handshake if no Root port has been assigned on a bridge, and another type if a Root port has already been assigned.

Handshake When No Root Port is Elected

If a Root port has not been assigned on a bridge, 802.1W uses the Proposing -> Proposed -> Sync -> Synced -> Agreed handshake:

- Proposing – The Designated port on the root bridge sends an RST BPDU packet to its peer port that contains a proposal flag. The proposal flag is a signal that indicates that the Designated port is ready to put itself in a forwarding state (Figure 10.4). The Designated port continues to send this flag in its RST BPDU until it is placed in a forwarding state (Figure 10.7) or is forced to operate in 802.1D mode. (See “Compatibility of 802.1W with 802.1D” on page 48.)
- Proposed – When a port receives an RST BPDU with a proposal flag from the Designated port on its point-to-


point link, it asserts the Proposed signal and one of the following occurs (Figure 10.4):

- If the RST BPDU that the port receives is superior to what it can transmit, the port assumes the role of a Root port. (See the section on “Bridges and Bridge Port Roles” on page 10-23.)
- If the RST BPDU that the port receives is inferior to what it can transmit, then the port is given the role of Designated port.

NOTE: Proposed will never be asserted if the port is connected on a shared media link.


In Figure 10.4, Port3/Switch 200 is elected as the Root port

Figure 10.4 Proposing and Proposed Stage




- Sync – Once the Root port is elected, it sets a sync signal on all the ports on the bridge. The signal tells the ports to synchronize their roles and states (Figure 10.5). Ports that are non-edge ports with a role of Designated port change into a discarding state. These ports have to negotiate with their peer ports to establish their new roles and states.

Figure 10.5 Sync Stage




- Synced – Once the Designated port changes into a discarding state, it asserts a synced signal. Immediately, Alternate ports and Backup ports are synced. The Root port monitors the synced signals from all the bridge ports. Once all bridge ports asserts a synced signal, the Root port asserts its own synced signal (Figure 10.6).

Figure 10.6 Synced Stage



- **Agreed** – The Root port sends back an RST BPDU containing an agreed flag to its peer Designated port and moves into the forwarding state. When the peer Designated port receives the RST BPDU, it rapidly transitions into a forwarding state.

Figure 10.7 Agree Stage



At this point, the handshake mechanism is complete between Switch 100, the root bridge, and Switch 200.

Switch 200 updates the information on the Switch 200's Designated ports (Port2 and Port3) and identifies the new root bridge. The Designated ports send RST BPDUs, containing proposal flags, to their downstream bridges, without waiting for the hello timers to expire on them. This process starts the handshake with the downstream bridges.


For example, Port2/Switch 200 sends an RST BPDU to Port2/Switch 300 that contains a proposal flag. Port2/Switch 300 asserts a proposed signal. Ports in Switch 300 then set sync signals on the ports to synchronize and negotiate their roles and states. Then the ports assert a synced signal and when the Root port in Switch 300 asserts its synced signal, it sends an RST BPDU to Switch 200 with an agreed flag.

This handshake is repeated between Switch 200 and Switch 400 until all Designated and Root ports are in forwarding states.

Handshake When a Root Port Has Been Elected

If a non-root bridge already has a Root port, 802.1W uses a different type of handshake. For example, in Figure 10.8, a new root bridge is added to the topology.

Figure 10.8 Addition of a New Root Bridge




The handshake that occurs between Switch 60 and Switch 100 follows the one described in the previous section (“Handshake When No Root Port is Elected” on page 10-28). The former root bridge becomes a non-root bridge and establishes a Root port (Figure 10.9).

However, since Switch 200 already had a Root port in a forwarding state, 802.1W uses the Proposing -> Proposed -> Sync and Reroot -> Sync and Rerooted -> Rerooted and Synced -> Agreed handshake:

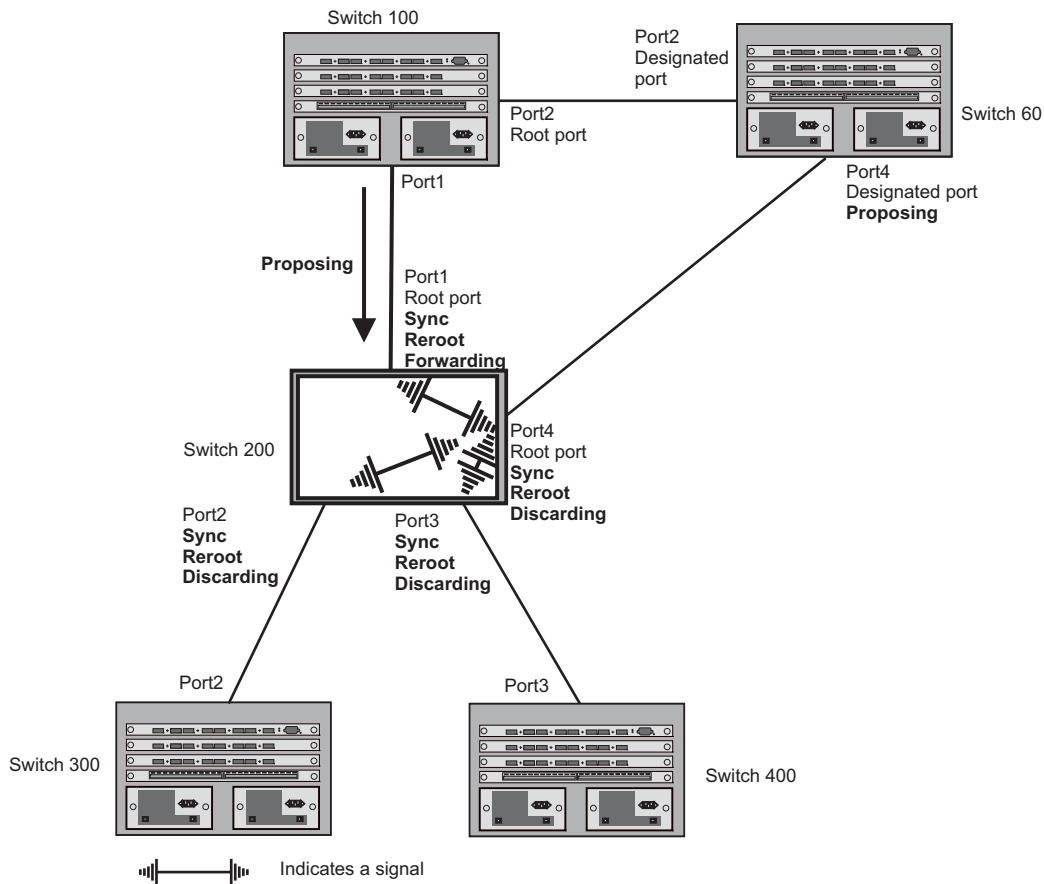
- Proposing and Proposed – The Designated port on the new root bridge (Port4/Switch 60) sends an RST BPDUs that contains a proposing signal to Port4/Switch 200 to inform the port that it is ready to put itself in a forwarding state (Figure 10.9). 802.1W algorithm determines that the RST BPDUs that Port4/Switch 200 received is superior to what it can generate, so Port4/Switch 200 assumes a Root port role.

Figure 10.9 New Root Bridge Sending a Proposal Flag




- Sync and Reroot – The Root port then asserts a sync and a reroot signal on all the ports on the bridge. The signal tells the ports that a new Root port has been assigned and they are to renegotiate their new roles and states. The other ports on the bridge assert their sync and reroot signals. Information about the old Root port is discarded from all ports. Designated ports change into discarding states (Figure 10.10).

Figure 10.10 Sync and Reroot



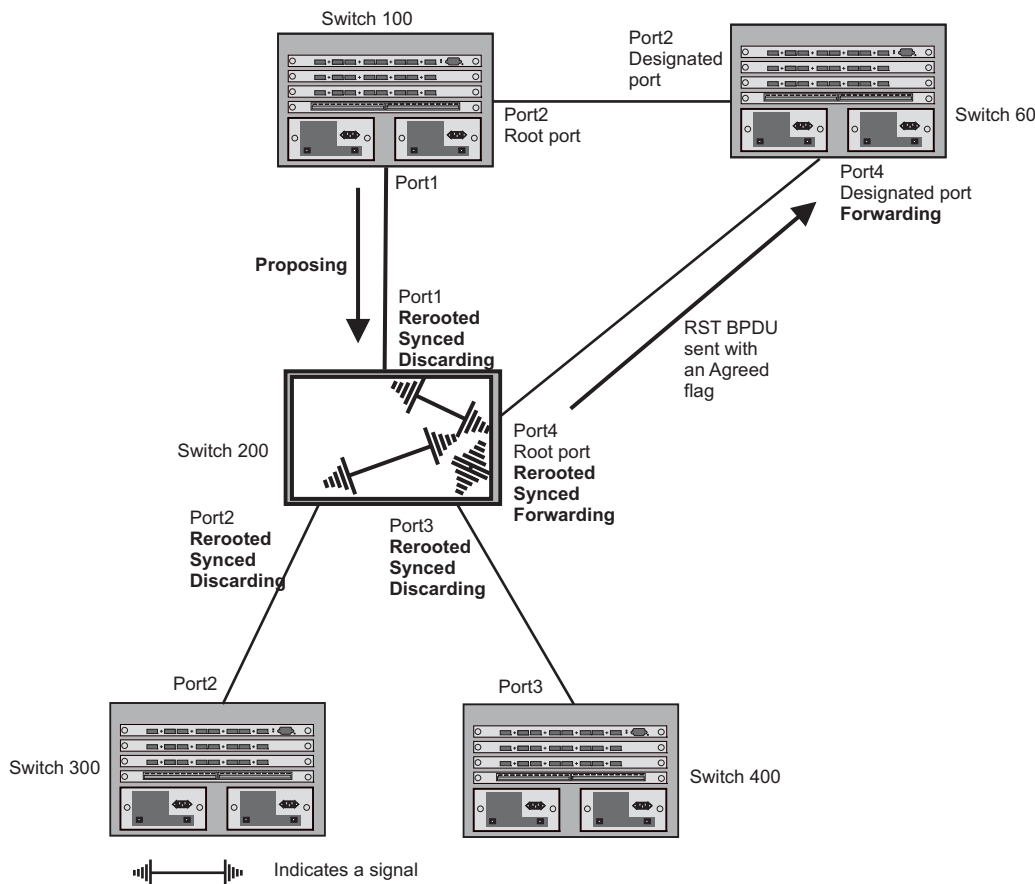
- Sync and Rerooted – When the ports on Switch 200 have completed the reroot phase, they assert their rerooted signals and continue to assert their sync signals as they continue in their discarding states. They also continue to negotiate their roles and states with their peer ports (Figure 10.11).

Figure 10.11 Sync and Rerooted



- Synced and Agree – When all the ports on the bridge assert their synced signals, the new Root port asserts its own synced signal and sends an RST BPDU to Port4/Switch 60 that contains an agreed flag (Figure 10.11). The Root port also moves into a forwarding state.


Figure 10.12 Rerouted, Synced, and Agreed



The old Root port on Switch 200 becomes an Alternate Port (Figure 10.13). Other ports on that bridge are elected to appropriate roles.

The Designated port on Switch 60 goes into a forwarding state once it receives the RST BPDU with the agreed flag.

Figure 10.13 Handshake Completed After Election of New Root Port



Recall that Switch 200 sent the agreed flag to Port4/Switch 60 and not to Port1/Switch 100 (the port that connects Switch 100 to Switch 200). Therefore, Port1/Switch 100 does not go into forwarding state instantly. It waits until two instances of the forward delay timer expires on the port before it goes into forwarding state.

At this point the handshake between the Switch 60 and Switch 200 is complete.

The remaining bridges (Switch 300 and Switch 400) may have to go through the reroot handshake if a new Root port needs to be assigned.

Convergence in a Simple Topology


The examples in this section illustrate how 802.1W convergence occurs in a simple Layer 2 topology at start-up.

NOTE: The remaining examples assume that the appropriate handshake mechanisms occur as port roles and states change.

Convergence at Start Up

In Figure 10.14, two bridges Switch 2 and Switch 3 are powered up. There are point-to-point connections between Port3/Switch 2 and Port3/Switch 3.

Figure 10.14 Convergence Between Two Bridges



At power up, all ports on Switch 2 and Switch 3 assume Designated port roles and are at discarding states before they receive any RST BPDU.

Port3/Switch 2, with a Designated role, transmits an RST BPDU with a proposal flag to Port3/Switch 3. A ports with a Designated role sends the proposal flag in its RST BPDU when they are ready to move to a forwarding state.


Port3/Switch 3, which starts with a role of Designated port, receives the RST BPDU and finds that it is superior to what it can transmit; therefore, Port3/Switch 3 assumes a new port role, that of a Root port. Port3/Switch 3 transmits an RST BPDU with an agreed flag back to Switch 2 and immediately goes into a forwarding state.

Port3/Switch 2 receives the RST BPDU from Port3/Switch 3 and immediately goes into a forwarding state.

Now 802.1W has fully converged between the two bridges, with Port3/Switch 3 as an operational root port in forwarding state and Port3/Switch 2 as an operational Designated port in forwarding state.

Next, Switch 1 is powered up (Figure 10.15).

Figure 10.15 Simple Layer 2 Topology



The point-to-point connections between the three bridges are as follows:

- Port2/Switch 1 and Port2/Switch 2
- Port4/Switch 1 and Port4/Switch 3
- Port3/Switch 2 and Port3/Switch 3

Ports 3 and 5 on Switch 1 are physically connected together.

At start up, the ports on Switch 1 assume Designated port roles, which are in discarding state. They begin sending RST BPDUs with proposal flags to move into a forwarding state.

When Port4/Switch 3 receives these RST BPDUs 802.1W algorithm determines that they are better than the RST BPDUs that were previously received on Port3/Switch 3. Port4/Switch 3 is now selected as Root port. This new assignment signals Port3/Switch 3 to begin entering the discarding state and to assume an Alternate port role. As it goes through the transition, Port3/Switch 3 negotiates a new role and state with its peer port, Port3/Switch 2.

Port4/Switch 3 sends an RST BPDU with an agreed flag to Port4/Switch 1. Both ports go into forwarding states.

Port2/Switch 2 receives an RST BPDU. The 802.1W algorithm determines that these RST BPDUs that are superior to any that any port on Switch 2 can transmit; therefore, Port2/Switch 2 assumes the role of a Root port.

The new Root port then signals all ports on the bridge to start synchronization. Since none of the ports are Edge ports, they all enter the discarding state and assume the role of Designated ports. Port3/Switch 2, which previously had a Designated role with a forwarding state, starts the discarding state. They also negotiate port roles and states with their peer ports. Port3/Switch 2 also sends an RST BPU to Port3/Switch 3 with a proposal flag to request permission go into a forwarding state.

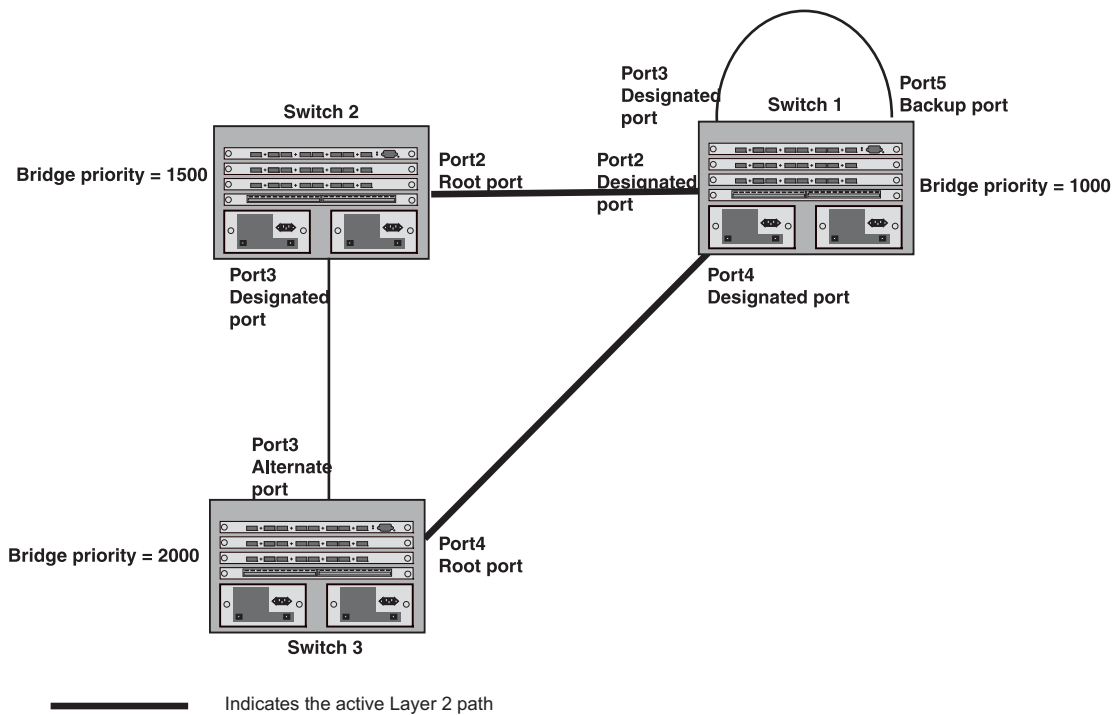
The Port2/Switch 2 bridge also sends an RST BPDU with an agreed flag Port2/Switch 1 that Port2 is the new Root port. Both ports go into forwarding states.

Now, Port3/Switch 3 is currently in a discarding state and is negotiating a port role. It received RST BPDUs from Port3/Switch 2. The 802.1W algorithm determines that the RST BPDUs Port3/Switch 3 received are superior to those it can transmit; however, they are not superior to those that are currently being received by the current Root port (Port4). Therefore, Port3 retains the role of Alternate port.

Ports 3/Switch 1 and Port5/Switch 1 are physically connected. Port5/Switch 1 received RST BPDUs that are superior to those received on Port3/Switch 1; therefore, Port5/Switch 1 is given the Backup port role while Port3 is given the Designated port role. Port3/Switch 1, does not go directly into a forwarding state. It waits until the forward delay time expires twice on that port before it can proceed to the forwarding state.

Once convergence is achieved, the active Layer 2 forwarding path converges as shown in Figure 10.16.

Figure 10.16 Active Layer 2 Path




Convergence After a Link Failure

What happens if a link in the 802.1W topology fails?

For example, Port2/Switch 2, which is the port that connects Switch 2 to the root bridge (Switch 1), fails. Both Switch 2 and Switch 1 notice the topology change (Figure 10.17).

Figure 10.17 Link Failure in the Topology



Switch 1 sets its Port2 into a discarding state.

At the same time, Switch 2 assumes the role of a root bridge since its root port failed and it has no operational Alternate port. Port3/Switch 2, which currently has a Designated port role, sends an RST BPDU to Switch 3. The RST BPDU contains a proposal flag and a bridge ID of Switch 2 as its root bridge ID.

When Port3/Switch 3 receives the RST BPDUs, 802.1W algorithm determines that they are inferior to those that the port can transmit. Therefore, Port3/Switch 3 is given a new role, that of a Designated port. Port3/Switch 3 then sends an RST BPDU with a proposal flag to Switch 2, along with the new role information. However, the root bridge ID transmitted in the RST BPDU is still Switch 1.

When Port3/Switch 2 receives the RST BPDU, 802.1W algorithm determines that it is superior to the RST BPDU that it can transmit; therefore, Port3/Switch 2 receives a new role; that of a Root port. Port3/Switch 2 then sends an RST BPDU with an agreed flag to Port3/Switch 3. Port3/Switch 2 goes into a forwarding state.

When Port3/Switch 3 receives the RST BPDU that Port3/Switch 2 sent, Port3/Switch 3 changes into a forwarding state, which then completes the full convergence of the topology.

Convergence at Link Restoration

When Port2/Switch 2 is restored, both Switch 2 and Switch 1 recognize the change. Port2/Switch 1 starts assuming the role of a Designated port and sends an RST BPDU containing a proposal flag to Port2/Switch 2.

When Port2/Switch 2 receives the RST BPDUs, 802.1W algorithm determines that the RST BPDUs the port received are better than those received on Port3/Switch 3; therefore, Port2/Switch 2 is given the role of a Root port. All the ports on Switch 2 are informed that a new Root port has been assigned which then signals all the ports to synchronize their roles and states. Port3/Switch 2, which was the previous Root port, enters a discarding state and negotiates with other ports on the bridge to establish its new role and state, until it finally assumes the role of a Designated port.

Next, the following happens:

- Port3/Switch 2, the Designated port, sends an RST BPDU, with a proposal flag to Port3/Switch 3.
- Port2/Switch 2 also sends an RST BPDU with an agreed flag to Port2/Switch 1 and then places itself into a forwarding state.

When Port2/Switch 1 receives the RST BPDU with an agreed flag sent by Port2/Switch 2, it puts that port into a forwarding state. The topology is now fully converged.

When Port3/Switch 3 receives the RST BPDU that Port3/Switch 2 sent, 802.1W algorithm determines that these RST BPDUs are superior to those that Port3/Switch 3 can transmit. Therefore, Port3/Switch 3 is given a new role, that of an Alternate port. Port3/Switch 3 immediately enters a discarding state.


Now Port3/Switch 2 does not go into a forwarding state instantly like the Root port. It waits until the forward delay timer expires twice on that port while it is still in a Designated role, before it can proceed to the forwarding state. The wait, however, does not cause a denial of service, since the essential connectivity in the topology has already been established.

When fully restored, the topology is the same as that shown on Figure 10.15.

Convergence in a Complex 802.1W Topology

The following is an example of a complex 802.1W topology.

Figure 10.18 Complex 802.1W Topology



In Figure 10.18, Switch 5 is selected as the root bridge since it is the bridge with the highest priority. Lines in the figure show the point-to-point connection to the bridges in the topology.

Switch 5 sends an RST BPDU that contains a proposal flag to Port5/Switch 2. When handshakes are completed in Switch 5, Port5/Switch 2 is selected as the Root port on Switch 2. All other ports on Switch 2 are given Designated port role with discarding states.

Port5/Switch 2 then sends an RST BPDU with an agreed flag to Switch 5 to confirm that it is the new Root port and the port enters a forwarding state. Port7 and Port8 are informed of the identity of the new Root port. 802.1W algorithm selects Port7 as the Designated port while Port8 becomes the Backup port.

Port3/Switch 5 sends an RST BPDU to Port3/Switch 6 with a proposal flag. When Port3/Switch 5 receives the RST BPDU, handshake mechanisms select Port3 as the Root port of Switch 6. All other ports are given a Designated port role with discarding states. Port3/Switch 6 then sends an RST BPDU with an agreed flag to Port3/Switch 5 to confirm that it is the Root port. The Root port then goes into a forwarding state.

Now, Port4/Switch 6 receives RST BPDUs that are superior to what it can transmit; therefore, it is given the Alternate port role. The port remains in discarding state.

Port5/Switch 6 receives RST BPDUs that are inferior to what it can transmit. The port is then given a Designated port role.

Next Switch 2 sends RST BPDUs with a proposal flag to Port3/Switch 4. Port3 becomes the Root port for the bridge; all other ports are given a Designated port role with discarding states. Port3/Switch 4 sends an RST BPDU with an agreed flag to Switch 2 to confirm that it is the new Root port. The port then goes into a forwarding state.

Now Port4/Switch 4 receives an RST BPDU that is superior to what it can transmit. The port is then given an Alternate port role, and remains in discarding state.

Likewise, Port5/Switch 4 receives an RST BPDU that is superior to what it can transmit. The port is also given an Alternate port role, and remains in discarding state.

Port2/Switch 2 transmits an RST BPDU with a proposal flag to Port2/Switch 1. Port2/Switch 1 becomes the Root port. All other ports on Switch 1 are given Designated port roles with discarding states.

Port2/Switch 1 sends an RST BPDU with an agreed flag to Port2/Switch 2 and Port2/Switch 1 goes into a forwarding state.

Port3/Switch 1 receives an RST BPDUs that is inferior to what it can transmit; therefore, the port retains its Designated port role and goes into forwarding state only after the forward delay timer expires twice on that port while it is still in a Designated role.

Port3/Switch 2 sends an RST BPDU to Port3/Switch 3 that contains a proposal flag. Port3/Switch 3 becomes the Root port, while all other ports on Switch 3 are given Designated port roles and go into discarding states. Port3/Switch 3 sends an RST BPDU with an agreed flag to Port3/Switch 2 and Port3/Switch 3 goes into a forwarding state.

Now, Port2/Switch 3 receives an RST BPDUs that is superior to what it can transmit so that port is given an Alternate port state.


Port4/Switch 3 receives an RST BPDU that is inferior to what it can transmit; therefore, the port retains its Designated port role.

Ports on all the bridges in the topology with Designated port roles that received RST BPDUs with agreed flags go into forwarding states instantly. However, Designated ports that did not receive RST BPDUs with agreed flags must wait until the forward delay timer expires twice on those port. Only then will these port move into forwarding states.

The entire 802.1W topology converges in less than 300 msec and the essential connectivity is established between the designated ports and their connected root ports.

After convergence is complete, Figure 10.19 shows the active Layer 2 path of the topology in Figure 10.18.

Figure 10.19 Active Layer 2 Path in Complex Topology



Propagation of Topology Change

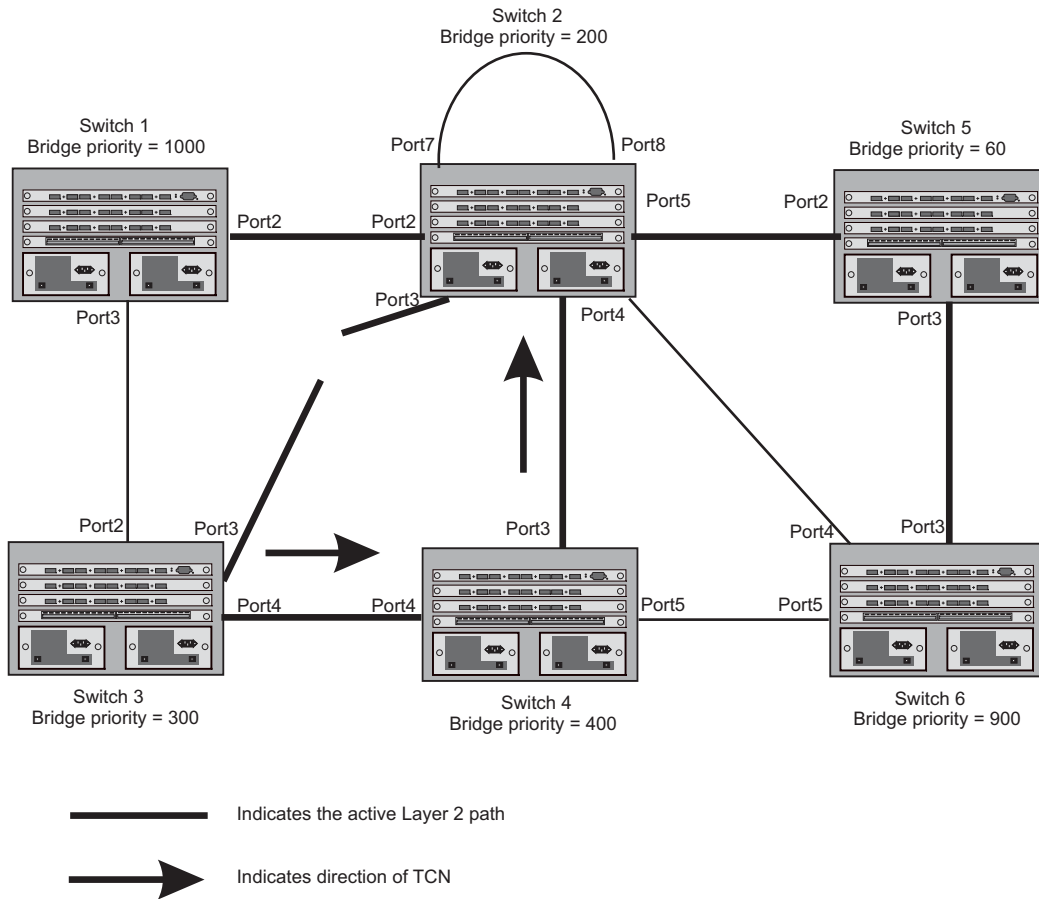
The Topology Change state machine generates and propagates the topology change notification messages on each port. When a Root port or a Designated port goes into a forwarding state, the Topology Change state machine on those ports send a topology change notice (TCN) to all the bridges in the topology to propagate the topology change.

NOTE: Edge ports, Alternate ports, or Backup ports do not need to propagate a topology change.

The TCN is sent in the RST BPDU that a port sends. Ports on other bridges in the topology then acknowledge the topology change once they receive the RST BPDU, and send the TCN to other bridges until all the bridges are informed of the topology change.

For example, Port3/Switch 2 in Figure 10.20, fails. Port4/Switch 3 becomes the new Root port. Port4/Switch 3 sends an RST BPDU with a TCN to Port4/Switch 4. To propagate the topology change, Port4/Switch 4 then starts a TCN timer on itself, on the bridge's Root port, and on other ports on that bridge with a Designated role. Then Port3/Switch 4 sends RST BPDU with the TCN to Port4/Switch 2. (Note the new active Layer 2 path in Figure 10.20.)


Figure 10.20 Beginning of Topology Change Notice



Switch 2 then starts the TCN timer on the Designated ports and sends RST BPDUs that contain the TCN as follows (Figure 10.21):

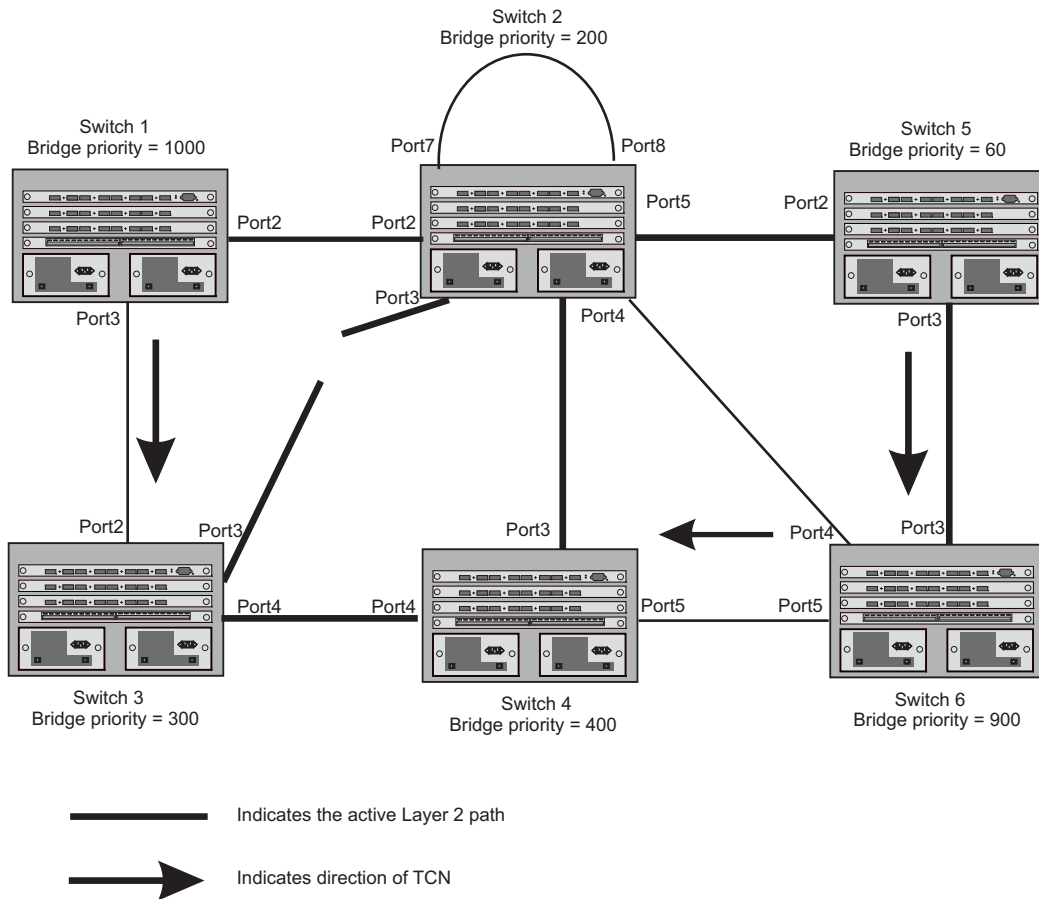
- Port5/Switch 2 sends the TCN to Port2/Switch 5
- Port4/Switch 2 sends the TCN to Port4/Switch 6
- Port2/Switch 2 sends the TCN to Port2/Switch 1

Figure 10.21 Sending TCN to Bridges Connected to Switch 2



Then FRY1, Switch 5, and Switch 6 send RST BPDUs that contain the TCN to Switch 3 and Switch 4 to complete the TCN propagation (Figure 10.22).

Figure 10.22 Completing the TCN Propagation



Compatibility of 802.1W with 802.1D

802.1W-enabled bridges are backward compatible with IEEE 802.1D bridges. This compatibility is managed on a per-port basis by the Port Migration state machine. **However, intermixing the two types of bridges in the network topology is not advisable if you want to take advantage of the rapid convergence feature.**


Compatibility with 802.1D means that an 802.1W-enabled port can send BPDUs in the STP or 802.1D format when one of the following events occur:

- The port receives a legacy BPDU. A legacy BPDU is an STP BPDU or a BPDU in an 802.1D format. The port that receives the legacy BPDU automatically configures itself to behave like a legacy port. It sends and receives legacy BPDUs only.
- The entire bridge is configured to operate in an 802.1D mode when an administrator sets the
- bridge parameter to zero at the CLI, forcing all ports on the bridge to send legacy BPDUs only.

Once a port operates in the 802.1D mode, 802.1D convergence times are used and rapid convergence is not realized.

For example, in Figure 10.23, Switch 10 and Switch 30 receive legacy BPDUs from Switch 20. Ports on Switch 10 and Switch 30 begin sending BPDUs in STP format to allow them to operate transparently with Switch 20.

Figure 10.23 802.1W Bridges with an 802.1D Bridge



Once Switch 20 is removed from the LAN, Switch 10 and Switch 30 receive and transmit BPDUs in the STP format to and from each other. This state will continue until the administrator enables the **force-migration-check** command to force the bridge to send RSTP BPDU during a migrate time period. If ports on the bridges continue to hear only STP BPDUs after this migrate time period, those ports will return to sending STP BPDUs. However, when the ports receive RST BPDUs during the migrate time period, the ports begin sending RST BPDUs. The migrate time period is non-configurable. It has a value of three seconds.

NOTE: The IEEE standards state that 802.1W bridges need to interoperate with 802.1D bridges. IEEE standards set the path cost of 802.1W bridges to be between 1 and 200,000,000; whereas path cost of 802.1D bridges are set between 1 and 65,535. In order for the two bridge types to be able to interoperate in the same topology, the administrator needs to configure the bridge path cost appropriately. Path costs for either 802.1W bridges or 802.1D bridges need to be changed; in most cases, path costs for 802.1W bridges need to be changed.

Configuring 802.1W Parameters on a Foundry Device

The remaining 802.1W sections explain how to configure the 802.1W protocol in a Foundry Layer 3 Switch.

s are shipped from the factory with 802.1W disabled. Use the following methods to enable or disable 802.1W. You can enable or disable 802.1W at the following levels:

- Port-based VLAN – Affects all ports within the specified port-based VLAN. When you enable or disable 802.1W within a port-based VLAN, the setting overrides the global setting. Thus, you can enable 802.1W for the ports within a port-based VLAN even when 802.1W is globally disabled, or disable the ports within a port-based VLAN when 802.1W is globally enabled.
- Individual port – Affects only the individual port. However, if you change the 802.1W state of the primary port in a trunk group, the change affects all ports in the trunk group.

Enabling or Disabling 802.1W in a Port-Based VLAN

Use the following procedure to disable or enable 802.1W on a device on which you have configured a port-based VLAN. Changing the 802.1W state in a VLAN affects only that VLAN.

USING THE CLI

To enable 802.1W for all ports in a port-based VLAN, enter commands such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning-tree 802-1w
```

Syntax: [no] spanning-tree 802-1w

USING THE WEB MANAGEMENT INTERFACE

You cannot enable or disable 802.1W on port-based VLAN using the Web management interface.

Enabling or Disabling 802.1W on a Single Spanning Tree

To enable 802.1W for all ports of a single spanning tree, use the procedure in this section.

USING THE CLI

Enter a command such as the following:

```
BigIron(config-vlan-10)# spanning-tree single 802-1w
```

Syntax: [no] spanning-tree single 802-1w

USING THE WEB MANAGEMENT INTERFACE

You cannot enable or disable 802.1W on a single spanning tree using the Web management interface.

Disabling or Enabling 802.1W on an Individual Port

The **spanning-tree 802-1w** or **spanning-tree single 802-1w** command must be used to initially enable 802.1W on ports. Both commands enable 802.1W on all ports that belong to the VLAN or to the single spanning tree.

Once 802.1W is enabled on a port, it can be disabled on individual ports. 802.1W that have been disabled on individual ports can then be enabled as required.

NOTE: If you change the 802.1W state of the primary port in a trunk group, the change affects all ports in that trunk group.

USING THE CLI

To disable or enable 802.1W on an individual port, enter commands such as the following:

```
BigIron(config)# interface 1/1
BigIron(config-if-1/1)# no spanning-tree
```

Syntax: [no] spanning-tree

USING THE WEB MANAGEMENT INTERFACE

You cannot enable or disable 802.1W on individual ports using the Web management interface.

Changing 802.1W Bridge Parameters

When you make changes to 802.1W bridge parameters, the changes are applied to individual ports on the bridge. To change 802.1W bridge parameters, use the following methods.

USING THE CLI

To designate a priority for a bridge, enter a command such as the following:

```
BigIron(config)# spanning-tree 802-1w priority 10
```

The command in this example changes the priority on a device on which you have not configured port-based VLANs. The change applies to the default VLAN. If you have configured a port-based VLAN on the device, you can configure the parameters only at the configuration level for individual VLANs. Enter commands such as the following:

```
BigIron(config)# vlan 20
BigIron(config-vlan-20)# spanning-tree 802-1w priority 0
```

To make this change in the default VLAN, enter the following commands:

```
BigIron(config)# vlan 1
BigIron(config-vlan-1)# spanning-tree 802-1w priority 0
```

Syntax: spanning-tree 802-1w [forward-delay <value>] | [hello-time <value>] | [max-age <time>] | [force-version <value>] | [priority <value>]

The **forward-delay** <value> parameter specifies how long a port waits before it forwards an RST BPDU after a topology change. This can be a value from 4 – 30 seconds. The default is 15 seconds.

The **hello-time** <value> parameter specifies the interval between two hello packets. This parameter can have a value from 1 – 10 seconds. The default is 2 seconds; however, set this value to at least 4 seconds to provide enough time for BPDUs to reach the root bridge before the timeout period expires on a non-root bridge port.

The **max-age** <value> parameter specifies the amount of time the device waits to receive a hello packet before it initiates a topology change. You can specify a value from 6 – 40 seconds. The default is 20 seconds.

Beginning with software release 07.6.03, the value of **max-age** must be greater than the value of **forward-delay** to ensure that the downstream bridges do not age out faster than the upstream bridges (those bridges that are closer to the root bridge).

The **force-version** <value> parameter forces the bridge to send BPDUs in a specific format. You can specify one of the following values:

- 0 – The STP compatibility mode. Only STP (or legacy) BPDUs will be sent.
- 2 – The default. RST BPDUs will be sent unless a legacy bridge is detected. If a legacy bridge is detected, STP BPDUs will be sent instead.

The default is 2.

The **priority** <value> parameter specifies the priority of the bridge. You can enter a value from 0 – 65535. A lower numerical value means a the bridge has a higher priority. Thus, the highest priority is 0. The default is 32768.

You can specify some or all of these parameters on the same command line. If you specify more than one parameter, you must specify them in the order shown above, from left to right.

USING THE WEB MANAGEMENT INTERFACE

You cannot modify 802.1W bridge parameters using the Web management interface.

Changing Port Parameters

The 802.1W port commands can be enabled on individual ports or on multiple ports, such as all ports that belong to a VLAN.

The 802.1W port parameters are preconfigured with default values. If the default parameters meet your network requirements, no other action is required.

You can change the following 802.1W port parameters using the following methods.

USING CLI

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning-tree 802-1w ethernet 1/5 path-cost 15 priority 64
```

Syntax: spanning-tree 802-1w ethernet <portnum> path-cost <value> | priority <value> | [admin-edge-port] | [admin-pt2pt-mac] | [force-migration-check]

The **ethernet | pos** <portnum> parameter specifies the interface used.

The **path-cost** <value> parameter specifies the cost of the port's path to the root bridge. 802.1W prefers the path with the lowest cost. You can specify a value from 1 – 20,000,000. Table 1 shows the recommended path cost values from the IEEE standards.

Table 1: Recommended Path Cost Values of 802.1W

Link Speed	Recommended (Default) 802.1W Path Cost Values	Recommended 802.1W Path Cost Range
Less than 100 kilobits per second	200,000,000	20,000,000 – 200,000,000
1 Megabit per second	20,000,000	2,000,000 – 200,000,000

Table 1: Recommended Path Cost Values of 802.1W

Link Speed	Recommended (Default) 802.1W Path Cost Values	Recommended 802.1W Path Cost Range
10 Megabits per second	2,000,000	200,000 – 200,000,000
100 Megabits per second	200,000	20,000 – 200,000,000
1 Gigabit per second	20,000	2,000 – 200,000,000
10 Gigabits per second	2,000	200 – 20,000
100 Gigabits per second	200	20 – 2,000
1 Terabits per second	20	2 – 200
10 Terabits per second	2	1 – 20
OC-3c	128,000	12,800 – 1,280,000
OC-12c	32,000	3,200 – 320,000
OC-48c	8,000	800 – 80,000
OC-192c	2,000	200 – 20,000

The **priority** <value> parameter specifies the preference that 802.1W gives to this port relative to other ports for forwarding traffic out of the topology. You can specify a value from 8 – 252, in increments of 4. If you enter a value that is not divisible by four the software rounds to the nearest value that is. The default is 128. A higher numerical value means a lower priority; thus, the highest priority is 8

Set the **admin-edge-port** to enabled or disabled. If set to enabled, then the port becomes an edge port in the domain.

Set the **admin-pt2pt-mac** to enabled or disabled. If set to enabled, then a port is connected to another port through a point-to-point link. The point-to-point link increases the speed of convergence. This parameter, however, does not auto-detect whether or not the link is a physical point-to-point link.

The **force-migration-check** parameter forces the specified port to send one RST BPDU. If only STP BPDUs are received in response to the sent RST BPDU, then the port will go return to sending STP BPDUs.

USING THE WEB MANAGEMENT INTERFACE

You cannot modify 802.1W port parameters using the Web management interface.

EXAMPLE:

Suppose you want to enable 802.1W on a system with no active port-based VLANs and change the hello-time from the default value of 2 to 8 seconds. Additionally, suppose you want to change the path and priority costs for port 5 only. To do so, enter the following commands.

```
BigIron(config)# spanning-tree 802-1w hello-time 8
```

```
BigIron(config)# spanning-tree 802-1w ethernet 5 path-cost 15 priority 64
```

Displaying Information About 802-1W

You can display a summary or details of the 802.1W information.

USING THE CLI

To display a summary of 802-1W, use the following command:

```
BigIron(config)#show 802-1w
--- VLAN 1 [ STP Instance owned by VLAN 1 ] -----
VLAN 1 BPDU cam_index is 2 and the IGC and DMA master Are(HEX) 0 1 2 3
Bridge IEEE 802.1W Parameters:
Bridge          Bridge  Bridge  Bridge  Force  tx
Identifier      MaxAge  Hello   FwdDly  Version Hold
hex            sec     sec     sec     cnt
800000e080541700 20      2       15      Default 3

RootBridge      RootPath  DesignatedBri-  Root  Max  Fwd  Hel
Identifier      Cost      dge Identifier  Port  Age  Dly  lo
hex             hex
800000e0804c9c00 200000    800000e0804c9c00 1     20  15  2

Port IEEE 802.1W Parameters:
      <--- Config Params -->|<----- Current state ----->
Port  Pri PortPath P2P Edge Role      State      Designa- Designated
Num   Cost  Mac Port  Port      State      ted cost  bridge
1     128 200000  F  F  ROOT      FORWARDING 0      800000e0804c9c00
2     128 200000  F  F  DESIGNATED FORWARDING 200000 800000e080541700
3     128 200000  F  F  DESIGNATED FORWARDING 200000 800000e080541700
4     128 200000  F  F  BACKUP    DISCARDING 200000 800000e080541700
```

Syntax: show 802-1w [vlan <vlan-id>]

The **vlan** <vlan-id> parameter displays 802.1W information for the specified port-based VLAN.

The **show 802.1w display** command shows the information listed in Table 2.

Table 2: CLI Display of 802.1W Summary

This Field...	Displays...
VLAN ID	The port-based VLAN that owns the STP instance. VLAN 1 is the default VLAN. If you have not configured port-based VLANs on this device, all 802.1W information is for VLAN 1.
Bridge IEEE 802.1W Parameters	
Bridge Identifier	The ID of the bridge.
Bridge Max Age	The configured max age for this bridge. The default is 20.
Bridge Hello	The configured hello time for this bridge. The default is 2.
Bridge FwdDly	The configured forward delay time for this bridge. The default is 15.

Table 2: CLI Display of 802.1W Summary (Continued)

This Field...	Displays...
Force-Version	<p>The configured force version value. One of the following value is displayed:</p> <ul style="list-style-type: none"> • 0 – The bridge has been forced to operate in an STP compatibility mode. • 2 – The bridge has been forced to operate in an 802.1W mode. (This is the default.)
txHoldCnt	<p>The number of BPDUs that can be transmitted per Hello Interval. The default is 3.</p>
Root Bridge Identifier	<p>ID of the Root bridge that is associated with this bridge</p>
Root Path Cost	<p>The cost to reach the root bridge from this bridge. If the bridge is the root bridge, then this parameter shows a value of zero.</p>
Designated Bridge Identifier	<p>The bridge from where the root information was received. It can be from the root bridge itself, but it could also be from another bridge.</p>
Root Port	<p>The port on which the root information was received. This is the port that is connected to the Designated Bridge.</p>
Max Age	<p>The max age is derived from the Root port. An 802.1W-enabled bridge uses this value, along with the hello and message age parameters to compute the effective age of an RST BPDU.</p> <p>The message age parameter is generated by the Designated port and transmitted in the RST BPDU. RST BPDUs transmitted by a Designated port of the root bridge contains a message value of zero.</p> <p>Effective age is the amount of time the Root port, Alternate port, or Backup port retains the information it received from its peer Designated port. Effective age is reset every time a port receives an RST BPDU from its Designated port. If a Root port does not receive an RST BPDU from its peer Designated port for a duration more than the effective age, the Root port ages out the existing information and recomputes the topology.</p> <p>If the port is operating in 802.1D compatible mode, then max age functionality is the same as in 802.1D (STP).</p>
Fwd Dly	<p>The number of seconds a non-edge Designated port waits until it can apply any of the following transitions, if the RST BPDU it receives does not have an agreed flag:</p> <ul style="list-style-type: none"> • Discarding state to learning state • Learning state to forwarding state <p>When a non-edge port receives the RST BPDU it goes into forwarding state within 4 seconds or after two hello timers expire on the port.</p> <p>Fwd Dly is also the number of seconds that a Root port waits for an RST BPDU with a proposal flag before it applies the state transitions listed above.</p> <p>If the port is operating in 802.1D compatible mode, then forward delay functionality is the same as in 802.1D (STP).</p>

Table 2: CLI Display of 802.1W Summary (Continued)

This Field...	Displays...
Hello	The hello value derived from the Root port. It is the number of seconds between two Hello packets.
Port IEEE 802.1W Parameters	
Port Num	The port number shown in a slot#/port# format.
Pri	The configured priority of the port. The default is 128 or 0x80.
Port Path Cost	The configured path cost on a link connected to this port.
P2P Mac	Indicates if the point-to-point-mac parameter is configured to be a point-to-point link: <ul style="list-style-type: none"> • T – The link is configured as a point-to-point link. • F – The link is not configured as a point-to-point link. This is the default.
Edge port	Indicates if the port is configured as an operational Edge port: <ul style="list-style-type: none"> • T – The port is configured as an Edge port. • F – The port is not configured as an Edge port. This is the default.
Role	The current role of the port: <ul style="list-style-type: none"> • Root • Designated • Alternate • Backup • Disabled Refer to “Bridges and Bridge Port Roles” on page 10-23 for definitions of the roles.
State	The port’s current 802.1W state. A port can have one of the following states: <ul style="list-style-type: none"> • Forwarding • Discarding • Learning • Disabled Refer to “Bridge Port States” on page 10-27 and “Edge Port and Non-Edge Port States” on page 10-27.
Designated Cost	The best root path cost that this port received, including the best root path cost that it can transmit.
Designated Bridge	The ID of the bridge that sent the best RST BPDU that was received on this port.

To display detailed information about 802-1W, using the following command:

```
BigIron(config)#show 802-1w detail
=====
VLAN 1 - MULTIPLE SPANNING TREE (MSTP - IEEE 802.1W) ACTIVE
=====
BridgeId 800000e080541700, forceVersion 2, txHoldCount 3

Port 1 - Role: ROOT - State: FORWARDING
  PathCost 200000, Priority 128, AdminOperEdge F, AdminPt2PtMac F
  DesignatedPriority - Root: 0x800000e0804c9c00, Bridge: 0x800000e080541700
  ActiveTimers - rrWhile 4 rcvdInfoWhile 4
  MachineStates - PIM: CURRENT, PRT: ROOT_PORT, PST: FORWARDING
  TCM: ACTIVE, PPM: SENDING_STP, PTX: TRANSMIT_IDLE
  Received - RST BPDUs 0, Config BPDUs 1017, TCN BPDUs 0

Port 2 - Role: DESIGNATED - State: FORWARDING
  PathCost 200000, Priority 128, AdminOperEdge F, AdminPt2PtMac F
  DesignatedPriority - Root: 0x800000e0804c9c00, Bridge: 0x800000e080541700
  ActiveTimers - helloWhen 0
  MachineStates - PIM: CURRENT, PRT: DESIGNATED_PORT, PST: FORWARDING
  TCM: ACTIVE, PPM: SENDING_RSTP, PTX: TRANSMIT_IDLE
  Received - RST BPDUs 0, Config BPDUs 0, TCN BPDUs 0
```

Syntax: show 802-1w detail [vlan <vlan-id>]

The **vlan <vlan-id>** parameter displays 802.1W information for the specified port-based VLAN.

The **show spanning-tree 802.1W** command shows the following information.

This Field...	Displays...
VLAN ID	ID of the VLAN that owns the instance of 802.1W and whether or not it is active.
Bridge ID	ID of the bridge.
forceVersion	the configured version of the bridge: <ul style="list-style-type: none"> • 0 – The bridge has been forced to operate in an STP compatible mode. • 2 – The bridge has been forced to operate in an 802.1W mode.
txHoldCount	The number of BPDUs that can be transmitted per Hello Interval. The default is 3.
Port	ID of the port in slot#/port# format.

This Field...	Displays...
Role	<p>The current role of the port:</p> <ul style="list-style-type: none"> • Root • Designated • Alternate • Backup • Disabled <p>Refer to “Bridges and Bridge Port Roles” on page 10-23 for definitions of the roles.</p>
State	<p>The port’s current 802.1W state. A port can have one of the following states:</p> <ul style="list-style-type: none"> • Forwarding • Discarding • Learning • Disabled <p>Refer to “Bridge Port States” on page 10-27 and “Edge Port and Non-Edge Port States” on page 10-27.</p>
Path Cost	<p>The configured path cost on a link connected to this port.</p>
Priority	<p>The configured priority of the port. The default is 128 or 0x80.</p>
AdminOperEdge	<p>Indicates if the port is an operational Edge port. Edge ports may either be auto-detected or configured (forced) to be Edge ports using the CLI:</p> <ul style="list-style-type: none"> • T – The port is and Edge port. • F – The port is not an Edge port. This is the default.
AdminP2PMac	<p>Indicates if the point-to-point-mac parameter is configured to be a point-to-point link:</p> <ul style="list-style-type: none"> • T – The link is a point-to-point link • F – The link is not a point-to-point link. This is the default.
DesignatedPriority	<p>Shows the following:</p> <ul style="list-style-type: none"> • Root – Shows the ID of the root bridge for this bridge. • Bridge – Shows the ID of the Designated bridge that is associated with this port.

This Field...	Displays...
ActiveTimers	<p>Shows what timers are currently active on this port and the number of seconds they have before they expire:</p> <ul style="list-style-type: none"> • rrWhile – Recent root timer. A non-zero value means that the port has recently been a Root port. • rcvdInfoWhile – Received information timer. Shows the time remaining before the information held by this port expires (ages out). This timer is initialized with the effective age parameter. (See “Max Age” on page 10-54.) • rbWhile – Recent backup timer. A non-zero value means that the port has recently been a Backup port. • helloWhen – Hello period timer. The value shown is the amount of time between hello messages. • tcWhile – Topology change timer. The value shown is the interval when topology change notices can be propagated on this port. • fdWhile – Forward delay timer. (See the explanation for Fwd Dly on page 54.) • mdelayWhile – Migration delay timer. The amount of time that a bridge on the same LAN has to synchronize its migration state with this port before another BPDU type can cause this port to change the BPDU that it transmits.
Machine States	<p>The current states of the various state machines on the port:</p> <ul style="list-style-type: none"> • PIM – State of the Port Information state machine. • PRT – State of the Port Role Transition state machine. • PST – State of the Port State Transition state machine. • TCM – State of the Topology Change state machine. • PPM – State of the Port Protocol Migration. • PTX – State of the Port Transmit state machine. <p>Refer to the section “State Machines” on page 10-27 for details on state machines.</p>
Received	<p>Shows the number of BPDU types the port has received:</p> <ul style="list-style-type: none"> • RST BPDU – BPDU in 802.1W format. • Config BPDU – Legacy configuration BPDU (802.1D format). • TCN BPDU – Legacy topology change BPDU (802.1D format).

802.1W Draft 3

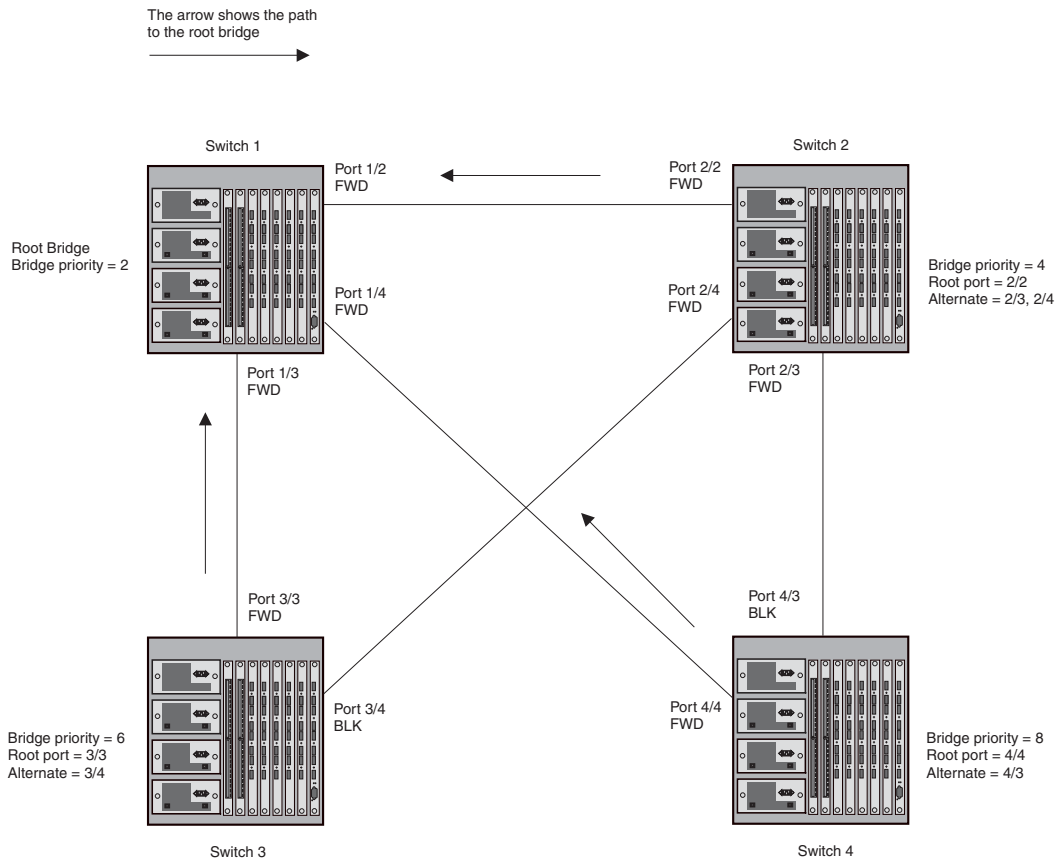
As an alternative to full 802.1W, you can configure 802.1W Draft 3. 802.1W Draft 3 provides a subset of the RSTP capabilities described in the 802.1W STP specification.

802.1W Draft 3 support is disabled by default. When the feature is enabled, if a root port on a Foundry device that is not the root bridge becomes unavailable, the device can automatically Switch over to an alternate root port, without reconvergence delays. 802.1W Draft 3 does not apply to the root bridge, since all the root bridge’s ports are always in the forwarding state.

Figure 10.24 shows an example of an optimal STP topology. In this topology, all the non-root bridges have at least two paths to the root bridge (Switch 1 in this example). One of the paths is through the root port. The other path


is a backup and is through the alternate port. While the root port is in the forwarding state, the alternate port is in the blocking state.

Figure 10.24 802.1W Draft 3 RSTP ready for failover



If the root port on a Switch becomes unavailable, 802.1W Draft 3 immediately fails over to the alternate port, as shown in Figure 10.25.

Figure 10.25 802.1W Draft 3 RSTP failover to alternate root port



In this example, port 3/3 on Switch 3 has become unavailable. In standard STP (802.1D), if the root port becomes unavailable, the Switch must go through the listening and learning stages on the alternate port to reconverge with the spanning tree. Thus, port 3/4 must go through the listening and learning states before entering the forwarding state and thus reconverging with the spanning tree.

8021.W Draft 3 avoids the reconvergence delay by calculating an alternate root port, and immediately failing over to the alternate port if the root port becomes unavailable. The alternate port is in the blocking state as long as the root port is in the forwarding state, but moves immediately to the active state if the root port becomes unavailable. Thus, using 8021.W Draft 3, Switch 3 immediately fails over to port 3/4, without the delays caused by the listening and learning states.

8021.W Draft 3 selects the port with the next-best cost to the root bridge. For example, on Switch 3, port 3/3 has the best cost to the root bridge and thus is selected by STP as the root port. Port 3/4 has the next-best cost to the root bridge, and thus is selected by 8021.W Draft 3 as the alternate path to the root bridge.

Once a failover occurs, the Switch no longer has an alternate root port. If the port that was an alternate port but became the root port fails, standard STP is used to reconverge with the network. You can minimize the reconvergence delay in this case by setting the forwarding delay on the root bridge to a lower value. For example, if the forwarding delay is set to 15 seconds (the default), change the forwarding delay to a value from 3 – 10 seconds.

During failover, 8021.W Draft 3 flushes the MAC addresses learned on the unavailable root port, selects the alternate port as the new root port, and places that port in the forwarding state. If traffic is flowing in both directions on the new root port, addresses are flushed (moved) in the rest of the spanning tree automatically.

Reconvergence Time

Spanning tree reconvergence using 8021.W Draft 3 can occur within one second.

After the spanning tree reconverges following the topology change, traffic also must reconverge on all the bridges attached to the spanning tree. This is true regardless of whether 8021.W Draft 3 or standard STP is used to reconverge the spanning tree.

Traffic reconvergence happens after the spanning tree reconvergence, and is achieved by flushing the Layer 2 information on the bridges.

- Following 8021.W Draft 3 reconvergence of the spanning tree, traffic reconvergence occurs in the time it takes for the bridge to detect the link changes plus the STP maximum age set on the bridge.
- If standard STP reconvergence occurs instead, traffic reconvergence takes two times the forward delay plus the maximum age.

NOTE: 8021.W Draft 3 does not apply when a failed root port comes back up. In this case, standard STP is used.

Configuration Considerations

8021.W Draft 3 is disabled by default. To ensure optimal performance of the feature before you enable it:

- Configure the bridge priorities so that the root bridge is one that supports 8021.W Draft 3. (Use a Foundry device or third-party device that supports 8021.W Draft 3.)
- Change the forwarding delay on the root bridge to a value lower than the default 15 seconds. Foundry recommends a value from 3 – 10 seconds. The lower forwarding delay helps reduce reconvergence delays in cases where 8021.W Draft 3 is not applicable, such as when a failed root port comes back up.
- Configure the bridge priorities and root port costs so that each device has an active path to the root bridge if its root port becomes unavailable. For example, port 3/4 is connected to port 2/4 on Switch 2, which has the second most favorable bridge priority in the spanning tree.

NOTE: If reconvergence involves changing the state of a root port on a bridge that supports 802.1D STP but not 8021.W Draft 3, then reconvergence still requires the amount of time it takes for the ports on the 802.1D bridge to change state to forwarding (as needed), and receive BPDUs from the root bridge for the new topology.

Enabling 8021.W Draft 3

8021.W Draft 3 is disabled by default. The procedure for enabling the feature differs depending on whether single STP is enabled on the device.

NOTE: STP must be enabled before you can enable 8021.W Draft 3.

Enabling 8021.W Draft 3 When Single STP Is Not Enabled

To enable 8021.W Draft 3 on a device that is not running single STP, use the following CLI method.

USING THE CLI

By default, each port-based VLAN on the device has its own spanning tree. To enable 8021.W Draft 3 in a port-based VLAN, enter commands such as the following:

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# spanning-tree rstp
```

Syntax: [no] spanning-tree rstp

This command enables 8021.W Draft 3. You must enter the command separately in each port-based VLAN in which you want to run 8021.W Draft 3.

NOTE: This command does not also enable STP. To enable STP, first enter the **spanning-tree** command without the **rstp** parameter. After you enable STP, enter the **spanning-tree rstp** command to enable 8021.W Draft 3.

To disable 8021.W Draft 3, enter the following command:

```
BigIron(config-vlan-10)# no spanning-tree rstp
```

Enabling 8021.W Draft 3 When Single STP Is Enabled

To enable 8021.W Draft 3 on a device that is running single STP, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# spanning-tree single rstp
```

Syntax: [no] spanning-tree single rstp

This command enables 8021.W Draft 3 on the whole device.

NOTE: This command does not also enable single STP. To enable single STP, first enter the **spanning-tree single** command without the **rstp** parameter. After you enable single STP, enter the **spanning-tree single rstp** command to enable 8021.W Draft 3.

To disable 8021.W Draft 3 on a device that is running single STP, enter the following command:

```
BigIron(config)# no spanning-tree single rstp
```

Single Spanning Tree (SSTP)

By default, each port-based VLAN on a Foundry device runs a separate spanning tree, which you can enable or disable on an individual VLAN basis.

Alternatively, you can configure a Foundry device to run a single spanning tree across all ports and VLANs on the device. The Single STP feature (SSTP) is especially useful for connecting a Foundry device to third-party devices that run a single spanning tree in accordance with the 802.1q specification.

SSTP uses the same parameters, with the same value ranges and defaults, as the default STP support on Foundry devices. See “STP Parameters and Defaults” on page 10-2.

SSTP Defaults

SSTP is disabled by default. When you enable the feature, all VLANs on which STP is enabled become members of a single spanning tree. All VLANs on which STP is disabled are excluded from the single spanning tree.

- To add a VLAN to the single spanning tree, enable STP on that VLAN.
- To remove a VLAN from the single spanning tree, disable STP on that VLAN.

When you enable SSTP, all the ports that are in port-based VLANs with STP enabled become members of a single spanning tree domain. Thus, the ports share a single BPDU broadcast domain. The Foundry device places all the ports in a non-configurable VLAN, 4094, to implement the SSTP domain. However, this VLAN does not affect port membership in the port-based VLANs you have configured. Other broadcast traffic is still contained within the individual port-based VLANs. Therefore, you can use SSTP while still using your existing VLAN configurations without changing your network. In addition, SSTP does not affect 802.1q tagging. Tagged and untagged ports alike can be members of the single spanning tree domain.

NOTE: When SSTP is enabled, the BPDUs on tagged ports go out untagged.

If you disable SSTP, all VLANs that were members of the single spanning tree run MSTP instead. In MSTP, each VLAN has its own spanning tree. VLANs that were not members of the single spanning tree were not enabled for STP. Therefore, STP remains disabled on those VLANs.

Enabling SSTP

To enable SSTP, use one of the following methods.

NOTE: If the device has only one port-based VLAN (the default VLAN), then the device is already running a single instance of STP. In this case, you do not need to enable SSTP. You need to enable SSTP only if the device contains more than one port-based VLAN and you want all the ports to be in the same STP broadcast domain.

USING THE CLI

To configure the Foundry device to run a single spanning tree, enter the following command at the global CONFIG level.

```
BigIron(config)# spanning-tree single
```

NOTE: If the device has only one port-based VLAN, the CLI command for enabling SSTP is not listed in the CLI. The command is listed only if you have configured a port-based VLAN.

To change a global STP parameter, enter a command such as the following at the global CONFIG level:

```
BigIron(config) spanning-tree single priority 2
```

This command changes the STP priority for all ports to 2.

To change an STP parameter for a specific port, enter commands such as the following:

```
BigIron(config) spanning-tree single ethernet 1/1 priority 10
```

The commands shown above override the global setting for STP priority and set the priority to 10 for port 1/1.

Here is the syntax for the global STP parameters.

Syntax: [no] spanning-tree single [forward-delay <value>]
[hello-time <value>] | [maximum-age <time>] | [priority <value>]

Here is the syntax for the STP port parameters.

Syntax: [no] spanning-tree single [ethernet <portnum> path-cost <value> | priority <value>]

NOTE: Both commands listed above are entered at the global CONFIG level.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access.
2. Click the Single checkbox next to Spanning Tree to place a checkmark in the box.
3. Make sure Enable, not Disable, is selected next to Spanning Tree.
4. Click Apply to apply the change to the device's running-config.
5. Select the [Save](#) link at the bottom of the panel. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Displaying SSTP information

To verify that SSTP is in effect, enter the following commands at any level of the CLI:

```
BigIron(config)# show span
```

Syntax: show span [vlan <vlan-id>] | [pvst-mode] | [<num>] |
[detail [vlan <vlan-id> [atm <portnum> | ethernet <portnum> | pos <portnum>] | <num>]]

The **vlan** <vlan-id> parameter displays STP information for the specified port-based VLAN.

The **pvst-mode** parameter displays STP information for the device's Per VLAN Spanning Tree (PVST+) compatibility configuration. See "PVST/PVST+ Compatibility" on page 10-75.

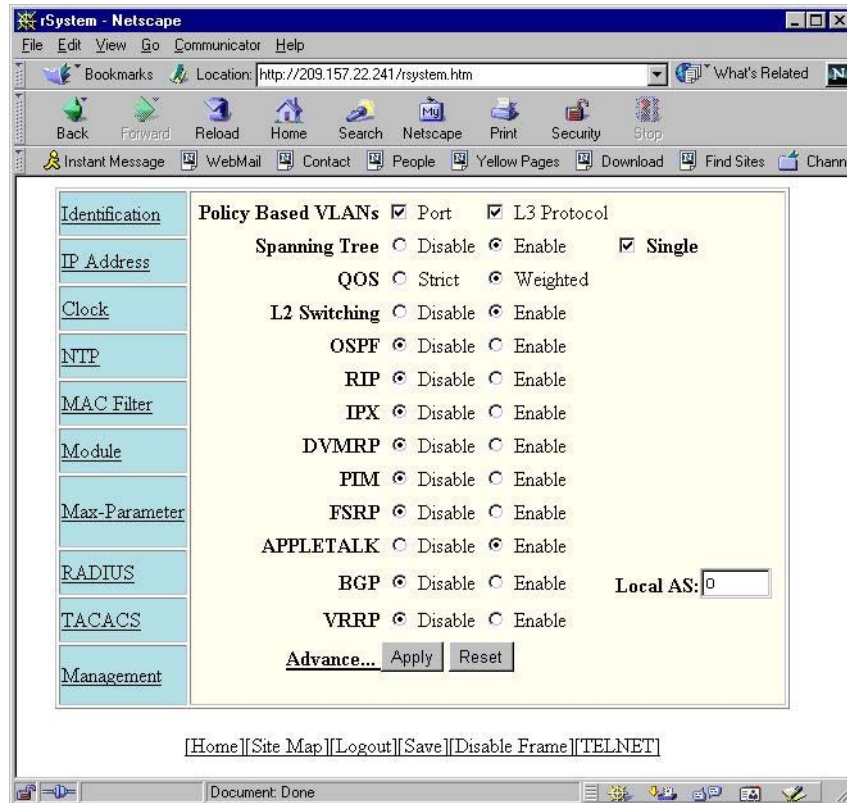
The <num> parameter displays only the entries after the number you specify. For example, on a device with three port-based VLANs, if you enter 1, then information for the second and third VLANs is displayed, but information for the first VLAN is not displayed. Information is displayed according to VLAN number, in ascending order. The entry number is not the same as the VLAN number. For example, if you have port-based VLANs 1, 10, and 2024,

then the command output has three STP entries. To display information for VLANs 10 and 2024 only, enter **show span 1**.

The **detail** parameter and its additional optional parameters display detailed information for individual ports. See “Displaying Detailed STP Information for Each Interface” on page 10-14.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration dialog is displayed.
2. Click on the Single checkbox next to Spanning Tree to place a checkmark in the box, as shown in the following example.



3. Click Apply to apply the change to the device's running-config.
4. Select the Save link at the bottom of the panel. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.


SuperSpan™

SuperSpan is a Foundry STP enhancement that allows Service Providers (SPs) to use STP in both SP networks and customer networks. The SP devices are Foundry devices and are configured to tunnel each customer's STP BPDUs through the SP. From the customer's perspective, the SP network is a loop-free non-blocking device or network. The SP network behaves like a hub in the sense that the necessary blocking occurs in the customer network, not in the SP.

The Foundry interfaces that connect the SP to a customer's network are configured as SuperSpan boundary interfaces. Each SuperSpan boundary interface is configured with a customer ID, to uniquely identify the customer's network within SuperSpan.

Figure 10.26 shows an example SuperSpan implementation. In this example, an SP's network is connected to multiple customers. Each customer network is running its own instance of standard STP. The Foundry devices in the SP are running SuperSpan.

Figure 10.26 SuperSpan example



In this example, the SP network contains two devices that are running SuperSpan. The SP is connected to two customer networks. Each customer network is running its own instance of STP. SuperSpan prevents Layer 2 loops in the traffic flow with each customer while at the same time isolating each customer's traffic and spanning tree from the traffic and spanning trees of other customers. For example, the SP devices provide loop prevention for Customer 1 while ensuring that Customer 1's traffic is never forwarded to Customer 2. In this example, customer 1 has two interfaces to the SP network, ports 1/1 and 1/2 connected to SP 1. The SP network behaves like a non-blocking hub. BPDUs are tunneled through the network. To prevent a Layer 2 loop, customer 1's port 1/2 enters the blocking state.

Customer ID

SuperSpan uses a SuperSpan customer ID to uniquely identify and forward traffic for each customer. You assign the customer ID as part of the SuperSpan configuration of the Foundry devices in the SP. In Figure 10.26, the spanning trees of customer 1 and customer 2 do not interfere with one another because the SP network isolates each customer's spanning tree based on the SuperSpan customer IDs in the traffic.

BPDU Forwarding

When a Foundry device receives a customer's BPDUs on a boundary interface, the device changes the destination MAC address of the BPDUs from the bridge group address (01-80-c2-00-00-00) as follows:

- The first byte (locally administered bit) is changed from 01 to 03, to indicate that the BPDUs need to be tunneled.

- The fourth and fifth bytes are changed to the customer STP ID specified on the boundary interface.

For example, if the customer's STP ID is 1, the destination MAC address of the customer's BPDUs is changed to the following: 03-80-c2-00-01-00.

Each Foundry device that is configured for SuperSpan forwards the BPDU using the changed destination MAC address. At the other end of the tunnel, the Foundry device connected to the customer's network changes the destination MAC address back to the bridge group address (01-80-c2-00-00-00).

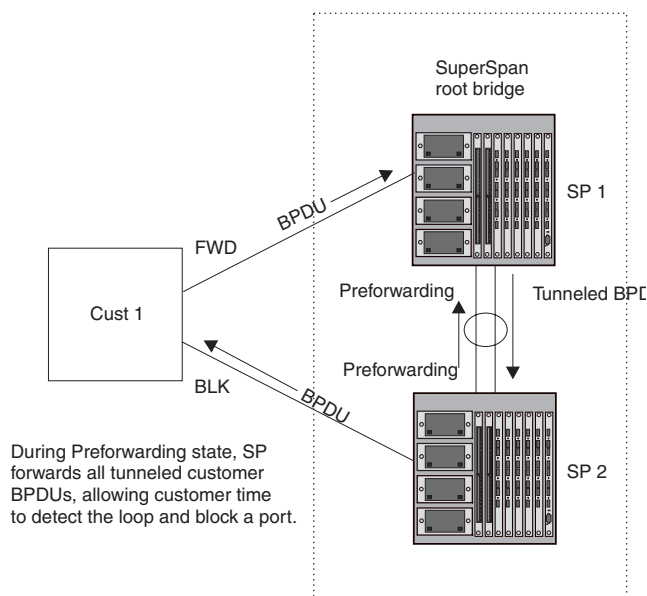
Preforwarding State

To ensure that the customer's network has time to converge at Layer 2 and prevent loops, the Foundry devices configured for SuperSpan use a special forwarding state, Preforwarding. The Preforwarding state occurs between the Learning and Forwarding states and by default lasts for five seconds. During the Preforwarding state, the Foundry device forwards tunneled BPDUs from customers only and does not forward data traffic. This ensures that the customer's network will detect the Layer 2 loop and block a port. The SP network remains unblocked. After the Preforwarding state, the Foundry ports change to the Forwarding state and forward data traffic as well as BPDUs.

The default length of the Preforwarding state is five seconds. You can change the length of the Preforwarding state to a value from 3 – 30 seconds.

Figure 10.27 shows an example of how the Preforwarding state is used.

Figure 10.27 SuperSpan Preforwarding state



In this example, a customer has two links to the SP. Since the SP is running SuperSpan, the SP ports enter the Preforwarding state briefly to allow the customer ports connected to the SP to detect the Layer 2 loop and block one of the ports.

NOTE: If you add a new device to a network that is already running SuperSpan, you must enable SuperSpan on the new device, at least on the VLANs that will be tunneling the customer traffic. Otherwise, the new device does not use the Preforwarding state. This can cause temporary loops in the network.

Mixing Single STP and Multiple Spanning Trees

You can use SuperSpan in any of the following combinations:

- Customer and SP networks both use multiple spanning trees (a separate spanning tree in each VLAN).
- Customer uses multiple spanning trees but SP uses Single STP (all STP-enabled VLANs are in the same

spanning tree).

- Customer uses Single STP but SP uses multiple spanning trees.
- Customer and SP networks both use Single STP.

The following sections provide an example of each combination.

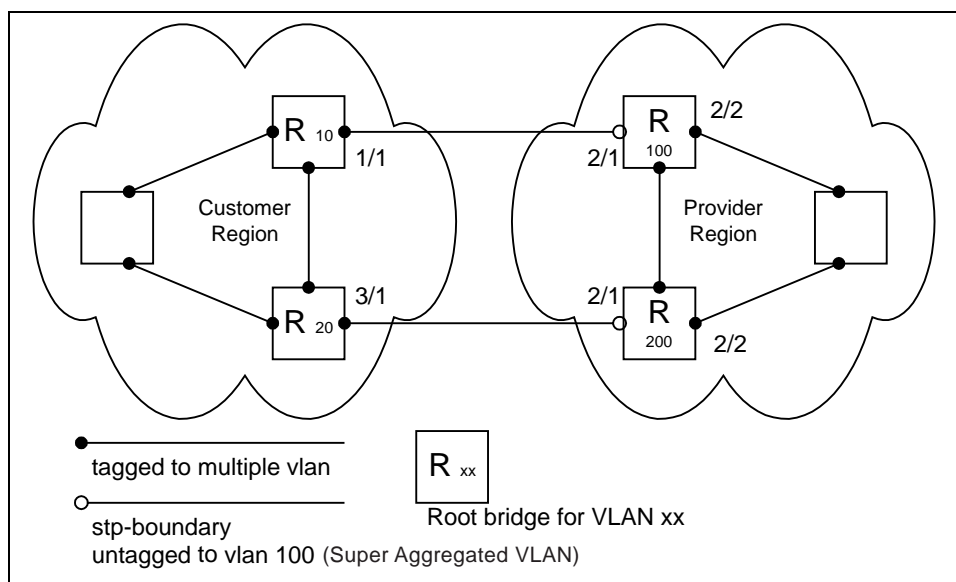
NOTE: All the combinations listed above are supported when the boundary ports joining the SP SuperSpan domain to the client spanning trees are untagged. For example, all these combinations are valid in super aggregated VLAN configurations. If the boundary ports are tagged, you cannot use Single STP in the client network in combination with multiple spanning trees in the SP SuperSpan domain.

The examples below are in super aggregated configuration scenarios.

Customer and SP Use Multiple Spanning Trees

Figure 10.28 shows an example of SuperSpan where both the customer network and the SP network use multiple spanning trees (a separate spanning tree in each port-based VLAN).

Figure 10.28 Customer and SP using multiple spanning trees



Both the customer and SP regions are running multiple spanning trees (one per port-based VLAN) in the Layer 2 switched network. The customer network contains VLANs 10 and 20 while the SP network contains VLANs 100 and 200. Customer traffic from VLAN 10 and VLAN 20 is aggregated by VLAN 100 in the SP since the boundary ports, 2/1 on R100 and R200, are untagged members of VLAN 100. By adjusting the bridge priority on VLANs 10 and 20, the customer can select a different root bridge for each spanning tree running in the customer network.

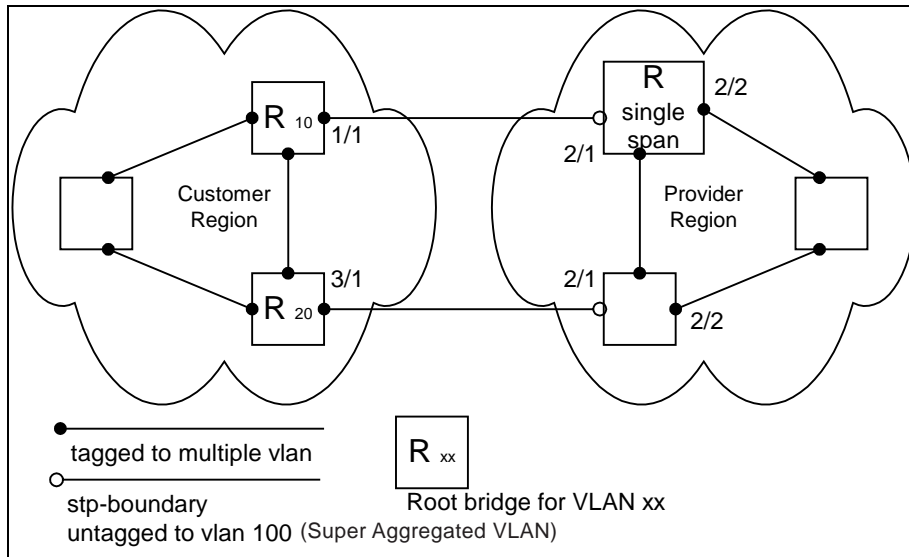
In the above example, STP in VLAN 10 will select R10 as the root bridge and make 1/1 on R10 forwarding while blocking port 3/1 on R20. The opposite occurs for STP in VLAN 20. As a result, both links connecting the customer and SP regions are fully utilized and serve as backup links at the same time, providing loop-free, non-blocking connectivity. In the SP network, multiple STP instances are running (one for VLAN 100 and one for VLAN 200) to ensure loop-free, non-blocking connectivity in each VLAN.

SuperSPAN boundaries are configured at port 2/1 of R100 and R200. Since the customer's traffic will be aggregated into VLAN 100 at the SP, the SP network appears to the customer to be a loop-free non-blocking hub to the customer network when port 2/2 on R200 is blocked by STP in VLAN 100.

Customer Uses Multiple Spanning Trees But SP Uses Single STP

Figure 10.29 shows an example of SuperSpan where the customer network uses multiple spanning trees while the SP network uses Single STP.

Figure 10.29 Customer using multiple spanning trees and SP using Single STP



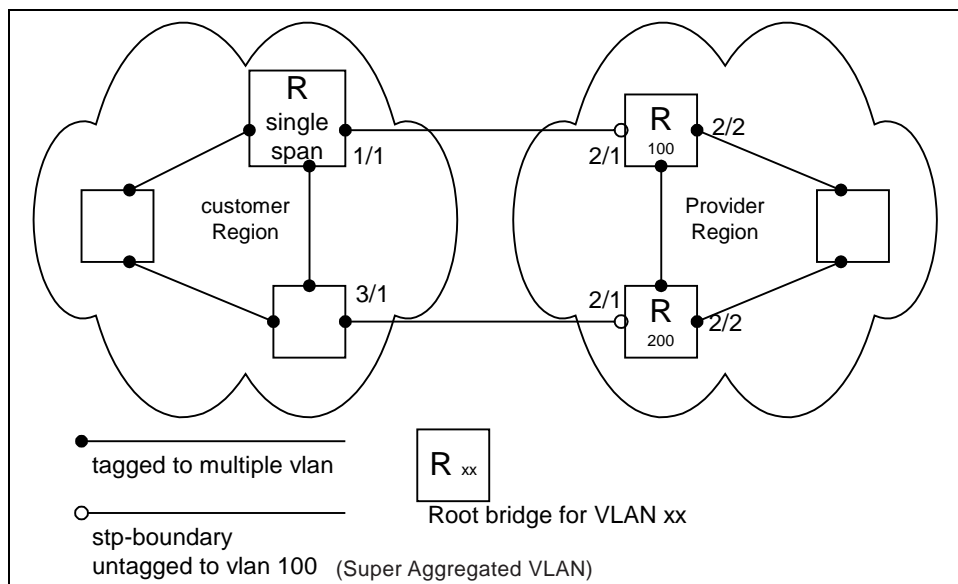
Customer traffic from different VLANs is maintained by different spanning trees, while the SP network is maintained by a single spanning tree. The SP can still use multiple VLANs at the core to separate traffic from different customers. However, all VLANs will have the same network topology because they are all calculated by the single spanning tree. The loop-free, non-blocking network acts like a hub for the customer network, with boundary ports 2/1 on each device being untagged members of VLAN 100.

Traffic from all VLANs in the customer network will be aggregated through VLAN 100 at the SP. This setup leaves the customer network's switching pattern virtually unchanged from the scenario in "Customer and SP Use Multiple Spanning Trees" on page 10-67, since the SP network still is perceived as a virtual hub, and maintenance of the hub's loop-free topology is transparent to the customer network.

Customer Uses Single STP But SP Uses Multiple Spanning Trees

Figure 10.30 shows an example of SuperSpan where the customer network uses Single STP while the SP uses multiple spanning trees.

Figure 10.30 Customer using Single STP and SP using multiple spanning trees




In this setup, the customer network is running a single spanning tree for VLANs 10 and 20. The traffic from VLAN 10 and 20 will be carried, or aggregated by VLAN 100 at the SP's network. The main difference between this scenario and the previous two scenarios is that all traffic at the customer's network now follows the same path, having the same STP root bridge in all VLANs. Therefore, the customer network will not have the ability to maximize network utilization on all its links. On the other hand, loop-free, non-blocking topology is still separately maintained by the customer network's single spanning tree and the SP's per-VLAN spanning tree on VLAN 100.

Customer and SP Use Single STP

Figure 10.31 shows an example of SuperSpan where the customer network and SP both use Single STP.

Figure 10.31 Customer and SP using Single STP



In this setup, both the customer and SP networks are running a single spanning tree at Layer 2. The traffic from VLAN 10 and 20 will be carried, or aggregated by VLAN 100 at the SP network as in the previous scenario. Loop-free, non-blocking topology is still separately maintained by the customer's single spanning tree and the SP's single spanning tree.

Configuring SuperSpan

To configure a Foundry device for SuperSpan:

- Configure each interface on the Foundry device that is connected to customer equipment as a boundary interface. This step enables the interface to convert the destination MAC address in the customer's BPDUs.

The software requires you to specify a SuperSpan customer ID when configuring the boundary interface. Use an ID from 1 – 65535. The customer ID uniquely identifies the customer. Use the same customer ID for each SP interface with the same customer. When tunneling BPDUs through the Foundry network, the devices use the customer ID to ensure that BPDUs are forwarded only to the customer's devices, and not to other customers' devices.

- Globally enable SuperSpan. This step enables the Preforwarding state.

Configuring a Boundary Interface

To configure the boundary interfaces on SP 1 in Figure 10.26 on page 10-65, enter the following commands:

```
BigIron(config)# interface 1/1
BigIron(config-if-e1000-1/1)# stp-boundary 1
BigIron(config)# interface 1/2
BigIron(config-if-e1000-1/2)# stp-boundary 2
```

These commands configure two interfaces on the Foundry device as SuperSpan boundary interfaces. Interface 1/1 is a boundary interface with customer 1. Interface 1/2 is a boundary interface with customer 2. Each boundary interface is associated with a number, which is the SuperSpan ID. The SuperSpan ID identifies the instance of SuperSpan you are associating with the interface. Use the same SuperSpan ID for each boundary interface with the same customer. Use a different SuperSpan ID for each customer. For example, use SuperSpan ID 1 for all the boundary interfaces with customer 1 and use SuperSpan ID 2 for all boundary interfaces with customer 2.

Syntax: [no] stp-boundary <num>

The <num> parameter specifies the SuperSpan ID. You can specify a number from 1 – 65535.

To configure the boundary interfaces on SP 2 in Figure 10.26 on page 10-65, enter the following commands:

```
BigIron(config)# interface 2/1
BigIron(config-if-e1000-2/1)# stp-boundary 1
BigIron(config)# interface 2/2
BigIron(config-if-e1000-2/2)# stp-boundary 2
```

Enabling SuperSpan

After you configure the SuperSpan boundary interfaces, enable SuperSpan. You can enable SuperSpan globally or on an individual VLAN level. If you enable the feature globally, the feature is enabled on all VLANs.

NOTE: If you enable the feature globally, then create a new VLAN, the new VLAN inherits the global SuperSpan state. For example, if SuperSpan is globally enabled when you create a VLAN, SuperSpan also is enabled in the new VLAN.

You also can change the length of the Preforwarding state to a value from 3 – 30 seconds. The default is 5 seconds.

To globally enable SuperSpan, enter the following command:

```
BigIron(config)# super-span-global
```

Syntax: [no] super-span-global [preforward-delay <secs>]

The <secs> parameter specifies the length of the Preforwarding state. You can specify from 3 – 30 seconds. The default is 5 seconds.

SuperSpan is enabled in all VLANs on the device. To disable SuperSpan in an individual VLAN, enter commands such as the following:

```
BigIron(config)# vlan 10
```

```
BigIron(config-vlan-10)# no super-span
```

Syntax: [no] super-span

Displaying SuperSpan Information

To display the boundary interface configuration and BPDU statistics, enter the following command:

```
BigIron(config)# show super-span
CID 1 Boundary Ports:
  Port  C-BPDU  C-BPDU  T-BPDU  T-BPDU
        Rxed   Txed   Rxed   Txed
  1/1   1       0       0       0
  1/2   0       0       0       0
  Total 1       0       0       0

CID 2 Boundary Ports:
  Port  C-BPDU  C-BPDU  T-BPDU  T-BPDU
        Rxed   Txed   Rxed   Txed
  2/1   0       0       3       0
  2/2   0       0       0       0
  Total 0       0       3       0
```

In this example, the device has two SuperSpan customer IDs.

Syntax: show superspan [cid <num>]

The **cid** <num> parameter specifies a SuperSpan customer ID. If you do not specify a customer ID, information for all the customer IDs configured on the device is shown.

This command shows the following information.

Table 10.7: CLI Display of SuperSpan Customer ID Information

This Field...	Displays...
CID	The SuperSpan customer ID number.
Port	The boundary port number.
C-BPDU Rxed	The number of BPDUs received from the client spanning tree.
C-BPDU Txed	The number of BPDUs sent to the client spanning tree.
T-BPDU Rxed	The number of BPDUs received from the SuperSpan tunnel.
T-BPDU Txed	The number of BPDUs sent to the SuperSpan tunnel.

To display general STP information, see “Displaying STP Information” on page 10-8.

STP per VLAN Group


STP per VLAN group is an STP enhancement that provides scalability while overcoming the limitations of the following scalability alternatives:

- Standard STP – You can configure only 128 instances of standard STP on a Foundry device. It is possible to need more instances of STP than this in large configurations. Using STP per VLAN group, you can aggregate STP instances.
- Single STP – Single STP allows all the VLANs to run STP, but each VLAN runs the same instance of STP,

resulting in numerous blocked ports that do not pass any Layer 2 traffic. STP per VLAN group uses all available links by load balancing traffic for different instances of STP on different ports. A port that blocks traffic for one spanning tree forwards traffic for another spanning tree.

STP per VLAN group allows you to group VLANs and apply the same STP parameter settings to all the VLANs in the group. Figure 10.32 shows an example of a STP per VLAN group implementation.

Figure 10.32 STP per VLAN Group Example



A master VLAN contains one or more member VLANs. Each of the member VLANs in a master VLAN runs the same instance of STP and uses the STP parameters configured for the master VLAN. In this example, the Foundry device is configured with VLANs 3, 4, 13, and 14. VLANs 3 and 4 are grouped in master VLAN 2, which is in STP group 1. VLANs 13 and 14 are grouped in master VLAN 12, which is in STP group 2. The VLANs in STP group 1 all share the same spanning tree. The VLANs in STP group 2 share a different spanning tree.

All the ports in the VLANs are tagged. The ports must be tagged so that they can be in both a member VLAN and the member's master VLAN. For example, ports 1/1 – 1/4 are in member VLAN 3 and also in master VLAN 2 (since master VLAN 2 contains member VLAN 3).

STP Load Balancing

Notice that the STP groups each have different STP priorities. In configurations that use the STP groups on multiple devices, you can use the STP priorities to load balance the STP traffic. By setting the STP priorities for the same STP group to different values on each device, you can cause each of the devices to be the root bridge for a different STP group. This type of configuration distributes the traffic evenly across the devices and also ensures that ports that are blocked in one STP group's spanning tree are used by another STP group's spanning tree for forwarding. See "Configuration Example for STP Load Sharing" on page 10-74 for an example using STP load sharing.

Configuring STP per VLAN Group

To configure STP per VLAN group:

- Configure the member VLANs.
- Optionally, configure master VLANs to contain the member VLANs. This is useful when you have a lot of member VLANs and you do not want to individually configure STP on each one. Each of the member VLANs in a master VLAN uses the STP settings of the master VLAN.
- Configure the STP groups. Each STP group runs a separate instance of STP.

Here are the CLI commands for implementing the STP per VLAN group configuration shown in Figure 10.32. The following commands configure the member VLANs (3, 4, 13, and 14) and the master VLANs (2 and 12). Notice that changes to STP parameters are made in the master VLANs only, not in the member VLANs.

```
BigIron(config)# vlan 2
BigIron(config-vlan-2)# spanning-tree priority 1
BigIron(config-vlan-2)# tagged ethernet 1/1 ethernet to 1/4
BigIron(config-vlan-2)# vlan 3
```



```
BigIron(config-vlan-3)# tagged ethernet 1/1 ethernet to 1/4
BigIron(config-vlan-3)# vlan 4
BigIron(config-vlan-4)# tagged ethernet 1/1 ethernet to 1/4
BigIron(config-vlan-4)# vlan 12
BigIron(config-vlan-12)# spanning-tree priority 2
BigIron(config-vlan-12)# tagged ethernet 1/1 ethernet to 1/4
BigIron(config-vlan-12)# vlan 13
BigIron(config-vlan-13)# tagged ethernet 1/1 ethernet to 1/4
BigIron(config-vlan-13)# vlan 14
BigIron(config-vlan-14)# tagged ethernet 1/1 ethernet to 1/4
BigIron(config-vlan-14)# exit
```

The following commands configure the STP groups.

```
BigIron(config)# stp-group 1
BigIron(config-stp-group-1)# master-vlan 2
BigIron(config-stp-group-1)# member-vlan 3 to 4
BigIron(config-stp-group-1)# exit
BigIron(config)# stp-group 2
BigIron(config-stp-group-2)# master-vlan 12
BigIron(config-stp-group-2)# member-vlan 13 to 14
```

Syntax: [no] stp-group <num>

This command changes the CLI to the STP group configuration level. The following commands are valid at this level. The <num> parameter specifies the STP group ID and can be from 1 – 32.

Syntax: [no] master-vlan <num>

This command adds a master VLAN to the STP group. The master VLAN contains the STP settings for all the VLANs in the STP per VLAN group. The <num> parameter specifies the VLAN ID. An STP group can contain one master VLAN.

NOTE: If you delete the master VLAN from an STP group, the software automatically assigns the first member VLAN in the group to be the new master VLAN for the group.

Syntax: [no] member-vlan <num> [to <num>]

This command adds additional VLANs to the STP group. These VLANs also inherit the STP settings of the master VLAN in the group.

Syntax: [no] member-group <num>

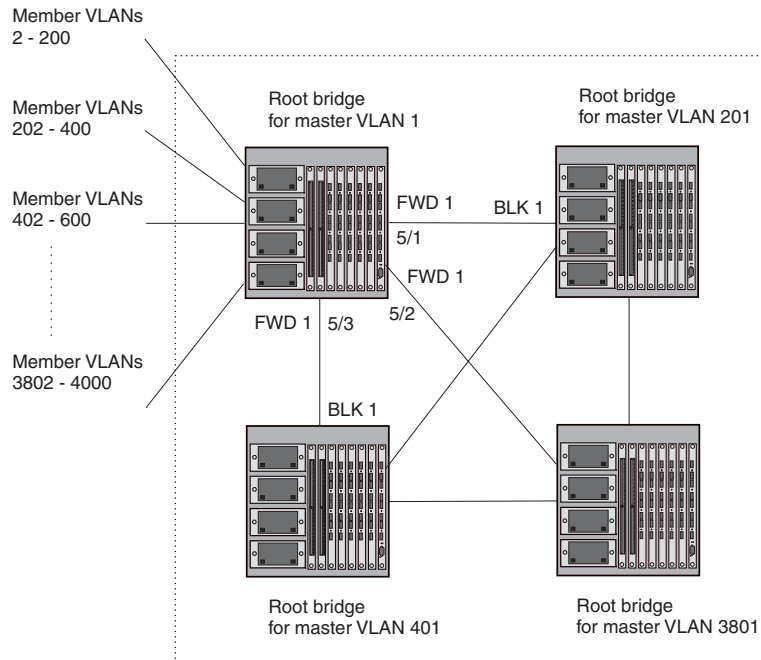
This command adds a member group (a VLAN group) to the STP group. All the VLANs in the member group inherit the STP settings of the master VLAN in the group. The <num> parameter specifies the VLAN group ID.

NOTE: This command is optional and is not used in the example above. For an example of this command, see “Configuration Example for STP Load Sharing”.

Configuration Example for STP Load Sharing

Figure 10.33 shows another example of a STP per VLAN group implementation.

Figure 10.33 More Complex STP per VLAN Group Example



In this example, each of the devices in the core is configured with a common set of master VLANs, each of which contains one or more member VLANs. Each of the member VLANs in a master VLAN runs the same instance of STP and uses the STP parameters configured for the master VLAN.

The STP group ID identifies the STP instance. All VLANs within an STP group run the same instance of STP. The master VLAN specifies the bridge STP parameters for the STP group, including the bridge priority. In this example, each of the devices in the core is configured to be the default root bridge for a different master VLAN. This configuration ensures that each link can be used for forwarding some traffic. For example, all the ports on the root bridge for master VLAN 1 are configured to forward BPDUs for master VLAN's spanning tree. Ports on the other devices block or forward VLAN 1's traffic based on STP convergence. All the ports on the root bridge for VLAN 2 forward VLAN 2's traffic, and so on.

All the ports in the VLANs are tagged. The ports must be tagged so that they can be in both a member VLAN and the member's master VLAN. For example, port 1/1 – and ports 5/1, 5/2, and 5/3 are in member VLAN 2 and master VLAN 1 (since master VLAN a contains member VLAN 2).

Here are the commands for configuring the root bridge for master VLAN 1 in figure Figure 10.32 for STP per VLAN group. The first group of commands configures the master VLANs. Notice that the STP priority is set to a different value for each VLAN. In addition, the same VLAN has a different STP priority on each device. This provides load balancing by making each of the devices a root bridge for a different spanning tree.

```
BigIron(config)# vlan 1
BigIron(config-vlan-1)# spanning-tree priority 1
BigIron(config-vlan-1)# tag ethernet 1/1 ethernet 5/1 to 5/3
BigIron(config-vlan-1)# vlan 201
BigIron(config-vlan-201)# spanning-tree priority 2
BigIron(config-vlan-201)# tag ethernet 1/2 ethernet 5/1 to 5/3
BigIron(config-vlan-201)# vlan 401
BigIron(config-vlan-401)# spanning-tree priority 3
BigIron(config-vlan-401)# tag ethernet 1/3 ethernet 5/1 to 5/3
...
BigIron(config-vlan-3601)# vlan 3801
```

```
BigIron(config-vlan-3801)# spanning-tree priority 20
BigIron(config-vlan-3801)# tag ethernet 1/20 ethernet 5/1 to 5/3
BigIron(config-vlan-3801)# exit
```

The next group of commands configures VLAN groups for the member VLANs. Notice that the VLAN groups do not contain the VLAN numbers assigned to the master VLANs. Also notice that no STP parameters are configured for the groups of member VLANs. Each group of member VLANs will inherit its STP settings from its master VLAN.

Set the bridge priority for each master VLAN to the highest priority (1) on one of the devices in the STP per VLAN group configuration. By setting the bridge priority to the highest priority, you make the device the default root bridge for the spanning tree. To ensure STP load balancing, make each of the devices the default root bridge for a different master VLAN.

```
BigIron(config)# vlan-group 1 vlan 2 to 200
BigIron(config-vlan-group-1)# tag ethernet 1/1 ethernet 5/1 to 5/3
BigIron(config-vlan-group-1)# vlan-group 2 vlan 202 to 400
BigIron(config-vlan-group-2)# tag ethernet 1/2 ethernet 5/1 to 5/3
BigIron(config-vlan-group-2)# vlan-group 3 vlan 402 to 600
BigIron(config-vlan-group-2)# tag ethernet 1/3 ethernet 5/1 to 5/3
...
BigIron(config-vlan-group-19)# vlan-group 20 vlan 3082 to 4000
BigIron(config-vlan-group-20)# tag ethernet 1/20 ethernet 5/1 to 5/3
BigIron(config-vlan-group-20)# exit
```

The following group of commands configures the STP groups. Each STP group in this configuration contains one master VLAN, which contains a VLAN group. This example shows that an STP group also can contain additional VLANs (VLANs not configured in a VLAN group).

```
BigIron(config)# stp-group 1
BigIron(config-stp-group-1)# master-vlan 1
BigIron(config-stp-group-1)# member-group 1
BigIron(config-stp-group-1)# member-vlan 4001 4004 to 4010
BigIron(config-stp-group-1)# stp-group 2
BigIron(config-stp-group-2)# master-vlan 201
BigIron(config-stp-group-2)# member-group 2
BigIron(config-stp-group-2)# member-vlan 4002 4003 4011 to 4015
BigIron(config-stp-group-2)# stp-group 3
BigIron(config-stp-group-3)# master-vlan 401
BigIron(config-stp-group-3)# member-group 3
...
BigIron(config-stp-group-19)# stp-group 20
BigIron(config-stp-group-20)# master-vlan 3081
BigIron(config-stp-group-20)# member-group 20
```

PVST/PVST+ Compatibility

The following sections describe the Per VLAN Spanning Tree (PVST) and PVST+ compatibility features on Foundry devices. Use the section that matches the software release you are using:

- For release 07.6.01 and later, see “PVST/PVST+ Compatibility – 07.6.01 and Later”.
- For releases 07.1.00 – 07.6.00x, see “PVST/PVST+ Compatibility – Earlier Than 07.6.01” on page 10-81.

PVST/PVST+ Compatibility – 07.6.01 and Later

Software release 07.6.01 enhances Foundry support for Cisco's Per VLAN Spanning Tree plus (PVST+), by allowing a Foundry device to run multiple spanning trees (MSTP) while also interoperating with IEEE 802.1Q devices¹.

Previous releases allow a Foundry device to interoperate with IEEE 802.1Q devices only when the Foundry device is configured for Single STP (SSTP). In this case, the Foundry device is operating as an IEEE 802.1Q device but

cannot run multiple spanning trees. The current release and previous releases allow the Foundry device to interoperate with PVST when the Foundry device is configured for MSTP.

NOTE: Foundry ports automatically detect PVST+ BPDUs and enable support for the BPDUs once detected. You do not need to perform any configuration steps to enable PVST+ support. However, to support the IEEE 802.1Q BPDUs, you might need to enable dual-mode support.

Foundry's support for Cisco's Per VLAN Spanning Tree plus (PVST+), allows a Foundry device to run multiple spanning trees (MSTP) while also interoperating with IEEE 802.1Q devices. Foundry ports automatically detect PVST+ BPDUs and enable support for the BPDUs once detected. The enhancement allows a port that is in PVST+ compatibility mode due to auto-detection to revert to the default MSTP mode when one of the following events occurs:

- The link is disconnected or broken
- The link is administratively disabled
- The link is disabled by interaction with the link-keepalive protocol

This enhancement allows a port that was originally interoperating with PVST+ to revert to MSTP when connected to a Foundry device.

Overview of PVST and PVST+

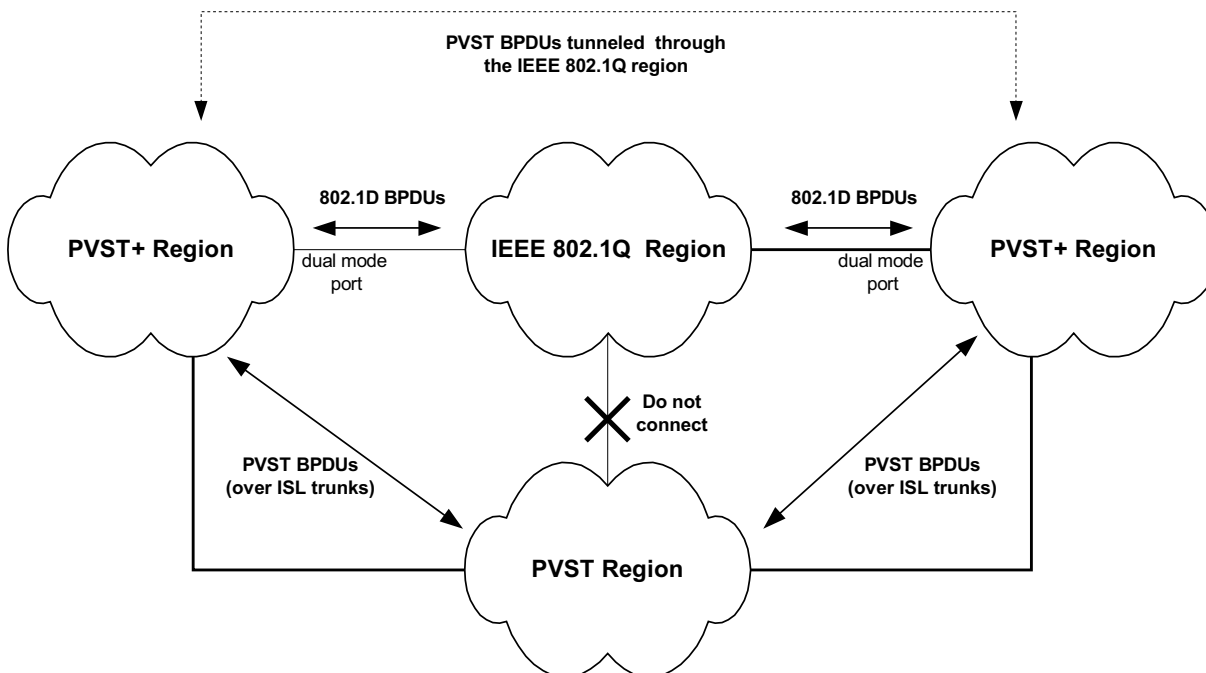
Per VLAN Spanning Tree (PVST) is a Cisco proprietary protocol that allows a Cisco device to have multiple spanning trees. The Cisco device can interoperate with spanning trees on other PVST devices but cannot interoperate with IEEE 802.1Q devices. An IEEE 802.1Q device has all its ports running a single spanning tree. **PVST+** is an extension of PVST that allows a Cisco device to also interoperate with devices that are running a single spanning tree (IEEE 802.1Q).

The enhanced PVST+ support in release 07.6.01 allows a Foundry device to interoperate with PVST spanning trees and the IEEE 802.1Q spanning tree at the same time.

IEEE 802.1Q and PVST regions cannot interoperate directly but can interoperate indirectly through PVST+ regions. PVST BPDUs are tunneled through 802.1Q regions, while PVST BPDUs for VLAN 1 (the IEEE 802.1Q VLAN) are processed by PVST+ regions. Figure 10.34 shows the interaction of IEEE 802.1Q, PVST, and PVST+ regions.

1. Cisco user documentation for PVST/PVST+ refers to the IEEE 802.1Q spanning tree as the **Common Spanning Tree (CST)**.

Figure 10.34 Interaction of IEEE 802.1Q, PVST, and PVST+ regions



VLAN Tags and Dual Mode

To support the IEEE 802.1Q (Common Spanning Tree) portion of PVST+, a port must be a member of VLAN 1. Cisco devices always use VLAN 1 to support the IEEE 802.1Q portion of PVST+.

For the port to also support the other VLANs (the PVST+ VLANs) in tagged mode, the dual-mode feature must be enabled on the port. The **dual-mode** feature enables the port to send and receive both tagged and untagged frames. When the dual-mode feature is enabled, the port is an untagged member of one of its VLANs and is at the same time a tagged member of all its other VLANs.

The untagged frames are supported on the port's **Port Native VLAN**. By default, the Port Native VLAN is the same as the device's **Default VLAN**¹, which by default is VLAN 1. Thus, to support IEEE 802.1Q in a typical configuration, the port must be able to send and receive untagged frames for VLAN 1 and tagged frames for the other VLANs.

If you want to use tagged frames on VLAN 1, you can change the default VLAN ID to an ID other than 1. You also can specify the VLAN on which you want the port to send and receive untagged frames (the Port Native VLAN). The Port Native VLAN ID does not need to be the same as the Default VLAN.

NOTE: Support for the IEEE 802.1Q spanning tree always uses VLAN 1, regardless of whether the devices are configured to use tagged or untagged frames on the VLAN.

Configuring PVST+ Support

PVST+ support is automatically enabled when the port receives a PVST BPDUs. You can manually enable the support at any time or disable the support if desired.

If you want a tagged port to also support IEEE 802.1Q BPDUs, you need to enable the dual-mode feature on the port. The dual-mode feature is disabled by default and must be enabled manually.

Starting with release 07.6.03, a port that is in PVST+ compatibility mode due to auto-detection reverts to the default MSTP mode when one of the following events occurs:

1. Cisco PVST/PVST+ documentation refers to the Default VLAN as the **Default Native VLAN**.

- The link is disconnected or broken
- The link is administratively disabled
- The link is disabled by interaction with the link-keepalive protocol

This allows a port that was originally interoperating with PVST+ to revert to MSTP when connected to a Foundry device.

Enabling PVST+ Support Manually

To immediately enable PVST+ support on a port, enter commands such as the following:

```
BigIron(config)# interface ethernet 1/1
BigIron(config-if-1/1)# pvst-mode
```

Syntax: [no] pvst-mode

NOTE: If you disable PVST+ support, the software still automatically enables PVST+ support if the port receives a BPDU with PVST+ format.

Enabling Dual-Mode Support

To enable the dual-mode feature on a port, enter the following command at the interface configuration level for the port:

```
BigIron(config-if-1/1)# dual-mode
```

Syntax: [no] dual-mode [<vlan-id>]

The <vlan-id> specifies the port's Port Native VLAN. This is the VLAN on which the port will support untagged frames. By default, the Port Native VLAN is the same as the default VLAN (which is VLAN 1 by default).

For more information about the dual-mode feature, see "Dual-Mode VLAN Ports" on page 22-69.

Displaying PVST+ Support Information

To display PVST+ information for ports on a Foundry device, enter the following command at any level of the CLI:

```
BigIron(config)# show span pvst-mode
PVST+ Enabled on:
Port      Method
1/1       Set by configuration
1/2       Set by configuration
2/10      Set by auto-detect
3/12      Set by configuration
4/24      Set by auto-detect
```

Syntax: show span pvst-mode

NOTE: This command is present in earlier releases but the output format has been changed to reflect the feature enhancements.

This command displays the following information.

Table 35: CLI Display of PVST+ Information

This Field...	Displays...
Port	The Foundry port number. Note: The command lists information only for the ports on which PVST+ support is enabled.

Table 35: CLI Display of PVST+ Information (Continued)

This Field...	Displays...
Method	<p>The method by which PVST+ support was enabled on the port. The method can be one of the following:</p> <ul style="list-style-type: none"> Set by configuration – You enabled the support. Set by auto-detect – The support was enabled automatically when the port received a PVST+ BPDU.

Configuration Examples


The following examples show configuration examples for two common configurations:

- Untagged IEEE 802.1Q BPDUs on VLAN 1 and tagged PVST+ BPDUs on other VLANs
- Tagged IEEE 802.1Q BPDUs on VLAN 1 and untagged BPDUs on another VLAN

Tagged Port Using Default VLAN 1 as its Port Native VLAN

Figure 10.36 shows an example of a PVST+ configuration that uses VLAN 1 as the untagged default VLAN and VLANs 2, 3, and 4 as tagged VLANs.

Figure 10.36 Default VLAN 1 for untagged BPDUs



To implement this configuration, enter the following commands.

Commands on the Foundry Device

```
BigIron(config)# vlan-group 1 vlan 2 to 4
BigIron(config-vlan-group-1)# tagged ethernet 1/1
BigIron(config-vlan-group-1)# exit
BigIron(config)# interface ethernet 1/1
BigIron(config-if-1/1)# dual-mode
BigIron(config-if-1/1)# pvst-mode
```

These commands configure a VLAN group containing VLANs 2, 3, and 4, add port 1/1 as a tagged port to the VLANs, and enable the dual-mode feature and PVST+ support on the port. The dual-mode feature allows the port to send and receive untagged frames for the default VLAN (VLAN 1 in this case) in addition to tagged frames for VLANs 2, 3, and 4. Enabling the PVST+ support ensures that the port is ready to send and receive PVST+ BPDUs. If you do not manually enable PVST+ support, the support is not enabled until the port receives a PVST+ BPDU.

The configuration leaves the default VLAN and the port's Port Native VLAN unchanged. The default VLAN is 1 and the port's Port Native VLAN also is 1. The dual-mode feature supports untagged frames on the default VLAN only. Thus, port 1/1 can send and receive untagged BPDUs for VLAN 1 and can send and receive tagged BPDUs for the other VLANs.


Port 1/1 will process BPDUs as follows:

- Process IEEE 802.1Q BPDUs for VLAN 1.
- Process tagged PVST BPDUs for VLANs 2, 3, and 4.
- Drop untagged PVST BPDUs for VLAN 1.

Untagged Port Using VLAN 2 as Port Native VLAN

Figure 10.37 shows an example in which a port's Port Native VLAN is not VLAN 1. In this case, VLAN 1 uses tagged frames and VLAN 2 uses untagged frames.

Figure 10.37 Port Native VLAN 2 for untagged BPDUs



To implement this configuration, enter the following commands.

Commands on the Foundry Device

```
BigIron(config)# default-vlan-id 4000
BigIron(config)# vlan 1
BigIron(config-vlan-1)# tagged ethernet 1/1
BigIron(config-vlan-1)# exit
BigIron(config)# vlan 2
BigIron(config-vlan-2)# tagged ethernet 1/1
BigIron(config-vlan-2)# exit
BigIron(config)# interface ethernet 1/1
BigIron(config-if-1/1)# dual-mode 2
BigIron(config-if-1/1)# pvst-mode
BigIron(config-if-1/1)# exit
```

These commands change the default VLAN ID, configure port 1/1 as a tagged member of VLANs 1 and 2, and enable the dual-mode feature and PVST+ support on port 1/1. Since VLAN 1 is tagged in this configuration, the default VLAN ID must be changed from VLAN 1 to another VLAN ID. Changing the default VLAN ID from 1 allows the port to process tagged frames for VLAN 1. VLAN 2 is specified with the **dual-mode** command, which makes VLAN 2 the port's Port Native VLAN. As a result, the port processes untagged frames and untagged PVST BPDUs on VLAN 2.

NOTE: Although VLAN 2 becomes the port's untagged VLAN, the CLI still requires that you add the port to the VLAN as a tagged port, since the port is a member of more than one VLAN.

Port 1/1 will process BPDUs as follows:

- Process IEEE 802.1Q BPDUs for VLAN 1.
- Process untagged PVST BPDUs for VLAN 2.
- Drop tagged PVST BPDUs for VLAN 1.

Note that when VLAN 1 is not the default VLAN, the ports must have the dual-mode featured enabled in order to process IEEE 802.1Q BPDUs.

For example, the following configuration is incorrect:

```
BigIron(config)# default-vlan-id 1000
BigIron(config)# vlan 1
BigIron(config-vlan-1)# tagged ethernet 1/1 to 1/2
BigIron(config-vlan-1)# exit
BigIron(config)# interface ethernet 1/1
BigIron(config-if-1/1)# pvst-mode
BigIron(config-if-1/1)# exit
BigIron(config)# interface ethernet 1/2
BigIron(config-if-1/2)# pvst-mode
BigIron(config-if-1/2)# exit
```


In the configuration above, all PVST BPDUs associated with VLAN 1 would be discarded. Since IEEE BPDUs associated with VLAN 1 are untagged, they are discarded because the ports in VLAN 1 are tagged. Effectively, the BPDUs are never processed by the Spanning Tree Protocol. STP assumes that there is no better bridge on the network and sets the ports to FORWARDING. This could cause a Layer 2 loop.

The following configuration is correct:

```
BigIron(config)# default-vlan-id 1000
BigIron(config)# vlan 1
BigIron(config-vlan-1)# tagged ethernet 1/1 to 1/2
BigIron(config-vlan-1)# exit
BigIron(config)# interface ethernet 1/1
BigIron(config-if-1/1)# pvst-mode
BigIron(config-if-1/1)# dual-mode
BigIron(config-if-1/1)# exit
BigIron(config)# interface ethernet 1/2
BigIron(config-if-1/2)# pvst-mode
BigIron(config-if-1/2)# dual-mode
BigIron(config-if-1/2)# exit
```

Setting the ports as dual-mode ensures that the untagged IEEE 802.1Q BPDUs reach the VLAN 1 instance.

PVST/PVST+ Compatibility – Earlier Than 07.6.01

Foundry devices that are configured to support a separate spanning tree in each port-based VLAN can interoperate with Cisco devices that are running Per VLAN Spanning Tree (PVST) or PVST+, Cisco proprietary STP implementations that support separate spanning trees in each port-based VLAN.

A Foundry device configured to run a separate spanning tree in each port-based VLAN automatically enables PVST/PVST+ support on a port if that port receives an STP BPDU with PVST/PVST+ format. You also can enable PVST/PVST+ support statically as well as display PVST/PVST+ information for each port.

The information in this section is for reference. If you are running PVST/PVST+ on the Cisco devices and the default support for separate spanning trees in each VLAN on the Foundry devices, then no configuration is necessary for the devices to share spanning tree information.

NOTE: If you plan to use the PVST/PVST+ support, do not use VLAN 1. PVST+ uses VLAN 1 as a single STP broadcast domain and thus uses a different BPDU format than for other VLANs.

PVST

Each spanning tree (that is, each instance of STP) has one device called the root bridge. The root bridge is the control point for the spanning tree, and sends STP status and topology change information to the other devices in the spanning tree by sending BPDUs to the other devices. The other devices forward the BPDUs as needed.

The format of an STP BPDU differs depending on whether it is a Cisco PVST BPDU or a Foundry BPDU. Foundry and Cisco devices also can support single STP BPDUs, which use another format.

- A Foundry device configured with a separate spanning tree in each VLAN sends BPDUs in standard IEEE 802.1D format, but includes a proprietary four-byte tag. The tag identifies the VLAN the BPDU is for.
- A Cisco device configured for PVST sends the BPDUs to multicast MAC address 01-00-0C-CC-CC-CD. If the device is configured for PVST+, then the device sends BPDUs for all VLANs except VLAN 1 to 01-00-0C-CC-CC-CD. The device sends BPDUs in VLAN 1 to 01-80-C2-00-00-00, the single STP address (see below and “PVST+”).
- A Foundry device configured for single STP (IEEE 802.1Q) sends untagged BPDUs to the well-known STP MAC address 01-80-C2-00-00-00.

NOTE: Cisco devices can be configured to interoperate with devices that support IEEE 802.1Q single STP, but the devices cannot be configured to run single STP.

Foundry's PVST support enables Foundry and Cisco devices that have separate spanning trees in each VLAN to interoperate. The Foundry PVST support is automatically enabled when a port receives a PVST BPDU and does not require configuration on the Foundry or Cisco device.

When PVST is enabled on a Foundry port, that port sends BPDUs in PVST format instead of Foundry's spanning tree format.

PVST+

Foundry devices and Cisco devices support separate spanning trees on an individual port-based VLAN basis. However, until the IEEE standard for multiple spanning trees is finalized, vendors are using different methods to support multiple spanning trees within their own products. PVST+ is an extension to PVST that enables a Cisco device to interoperate with other devices that are running a single spanning tree (IEEE 802.1Q) while still running a separate spanning tree in each VLAN.

PVST+ uses 802.1Q single STP BPDUs on VLAN 1 and PVST BPDUs (which have a proprietary format) for other VLANs. In this case, the Cisco device uses devices running 802.1Q as tunnels for PVST (non-802.1Q) traffic. The 802.1Q single STP BPDUs are addressed to the well-known STP MAC address 01-80-C2-00-00-00. The PVST BPDUs for the other VLANs are addressed to multicast address 01-00-0C-CC-CC-CD.

The PVST+ method can require manual configuration of STP parameters on the 802.1Q devices to ensure that traffic for the PVST VLANs is not blocked. In addition, the opportunities to adjust STP parameters to load balance traffic on a VLAN basis are limited when using PVST+.

Using Foundry Single STP with Cisco PVST+

Since Foundry's single STP feature complies with IEEE 802.1Q (the single STP specification), you also can use a Foundry device running single STP to interoperate with a Cisco device running PVST+. When you enable single STP on a Foundry device, the PVST compatibility feature is not enabled, even if a port receives a PVST BPDU.

Enabling PVST/PVST+ Statically

PVST/PVST+ support is automatically enabled on a port if the port receives a BPDU in PVST/PVST+ format. However, you can statically enable PVST/PVST+ support on a port if desired. In this case, the support is enabled immediately and support for Foundry tagged BPDUs is disabled at the same time. To enable the PVST/PVST+ support, use the following CLI method.

NOTE: When PVST/PVST+ support is enabled on a port, support for Foundry BPDUs is disabled.

USING THE CLI

To enable PVST/PVST+ support on a port, enter commands such as the following:

```
BigIron(config)# interface ethernet 1/1
BigIron(config-if-1/1)# pvst-mode
```

Syntax: [no] pvst-mode

NOTE: If you disable PVST/PVST+ support, the software still automatically enables PVST/PVST+ support if the port receives an STP BPDU with PVST/PVST+ format.

USING THE WEB MANAGEMENT INTERFACE

You cannot enable PVST support using the Web management interface.

Displaying PVST Information

To display PVST information, use the following CLI method.

USING THE CLI

To display PVST information for ports on a Foundry device, enter the following command at any level of the CLI:

```
BigIron(config)# show span pvst-mode

VLAN   Port   PVST   PVST
ID     Num.   Cfg.   On(by cfg. or detect)
200    10     0      1
200    11     1      1
```

This example shows that for VLAN 200, PVST support is statically enabled on port 11. PVST is not statically enabled on Port 10, but because port 10 received an incoming PVST BPDU on its interface, the port converted to using PVST mode.

Syntax: show span pvst-mode

The **show span pvst-mode** command displays the following information.

Table 10.8: CLI Display of PVST Information

This Field...	Displays...
VLAN ID	The VLAN to which the PVST/PVST+ information applies.
Port Num.	The Foundry port number.
PVST cfg.	Whether PVST support is statically enabled on the port. The value can be one of the following: <ul style="list-style-type: none"> 0 – The support has not been statically enabled. 1 – The support has been statically enabled.
PVST on (by cfg. or detect)	Whether PVST/PVST+ support is active on the port. The value can be one of the following: <ul style="list-style-type: none"> 0 – PVST/PVST+ support is not enabled. 1 – PVST/PVST+ support is enabled, either because you statically enabled the support or because the port received an STP BPDU with PVST/PVST+ format.

USING THE WEB MANAGEMENT INTERFACE

You cannot display PVST information using the Web management interface.

Chapter 11

Multiple Spanning Tree Protocol on Enterprise Releases

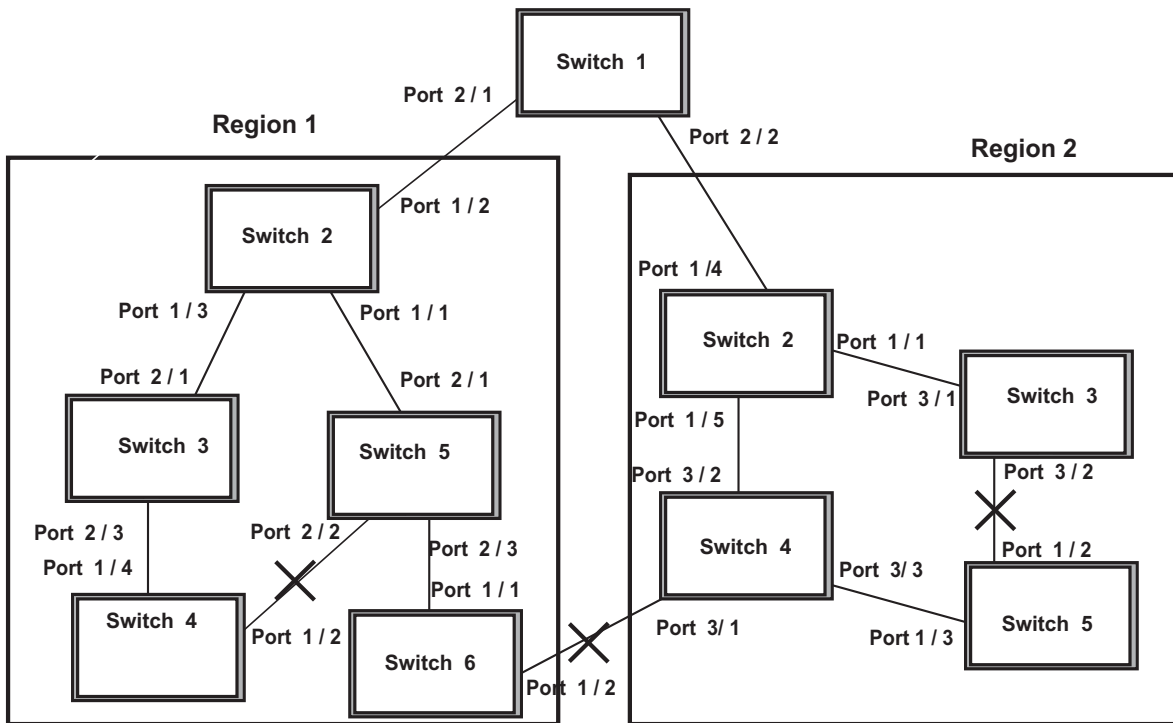
NOTE: This feature is supported on devices running Enterprise software release 08.0.00 and later. For the MSTP feature on the BigIron MG8 and NetIron 40G devices, see “Multiple Spanning Tree Protocol on BigIron MG8 and NetIron 40G” on page 14-1.

Multiple Spanning Tree Protocol (MSTP) as defined in IEEE 802.1s-2002 allows you to configure multiple STP instances. This ensures loop-free topology for 1 or more VLANs. Using MSTP, the entire network runs a common instance of RSTP. Within that common instance, one or more VLANs can be individually configured into distinct regions. The entire network runs the common spanning tree instance (CST) and the regions run a local instance. The local instance is known as Internal Spanning Tree (IST). The CST treats each instance of IST as a single bridge. Consequently, ports are blocked to prevent loops that might occur within an IST and also throughout the CST. In addition, MSTP can coexist with individual devices running STP or RSTP in the Common and Internal Spanning Trees instance (CIST). With the exception of the provisions for multiple instances, MSTP operates exactly like RSTP.

For example, in Figure 12 a network is configured with two regions: Region 1 and Region 2. The entire network is running an instance of CST. Each of the regions is running an instance of IST. In addition, this network contains Switch 1 running RSTP that isn't configured in a region and consequently is running in the CIST instance. In this configuration, the regions are each regarded as a single bridge to the rest of the network, as is Switch 1. The CST prevents loops from occurring across the network. Consequently, a port is blocked at either port 1/2 of switch 4.

Additionally loops must be prevented in each of the IST instances. Within the IST Region 1, a port is blocked at port 1/2 of switch 4 to prevent a loop in that region. Within Region 2, a port is blocked at port 3/2 of switch 3 to prevent a loop in that region.

Figure 12 MSTP Configured Network



The following definitions describe the STP instances that define an MSTP configuration:

Common Spanning (CST) – MSTP runs a single instance of spanning tree across all the bridges in a network called the Common Spanning Tree (CST). This instance treats each region as a single bridge. In all other ways, it operates exactly like Rapid Spanning Tree (RSTP).

Internal Spanning Tree (IST) – The instances of spanning tree that operate within a defined region are called ISTs (Internal Spanning Tree).

Common and Internal Spanning Trees (CIST) – This is the default MSTP instance 0. It contains all of the ISTs and all bridges that are not formally configured into a region. This instance interoperates with bridges running legacy STP and RSTP implementations.

Multiple Spanning Tree Instance (MSTI) – The MSTI is identified by an MST identifier (MSTid) value between 1 and 4094. This defines an individual instance of an IST. One or more VLANs can be assigned to an MSTI. A VLAN cannot be assigned to multiple MSTIs.

Configuring MSTP

To configure a switch for MSTP, you must first configure the name and the revision on each switch that is being configured for MSTP. This name is unique to each switch. You must then create an MSTP Instance and assign an ID. VLANs are then assigned to MSTP instances. These instances must be configured on all switches that interoperate with the same VLAN assignments. Port cost, priority and global parameters can then be configured for individual ports and instances. In addition, operational edge ports and point-to-point links can be created and MSTP can be disabled on individual ports.

Each of the commands used to configure and operate MSTP are described in the following:

- “Enabling MSTP on a Switch”
- “Setting the MSTP Revision Number and Name”
- “Configuring an MSTP Instance”

- “Disabling MSTP on a Port”
- “Configuring Port Priority and Path cost for an MSTP Instance”
- “Configuring Bridge Priority for an MSTP Instance”
- “Setting the MSTP Global Parameters”
- “Setting Ports To Be Operational Edge Ports”
- “Setting Point-to-Point Link”
- “Committing MSTP Configuration ID Changes”
- “Forcing Ports to Transmit the Configured BPDU Version”

Enabling MSTP on a Switch

To enable MSTP on your switch, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp start
```

Syntax: [no] mstp start

Setting the MSTP Revision Number and Name

Each switch that is running MSTP is configured with a name and revision number. These apply to the switch which can have many different VLANs that can belong to many different MSTP regions.

To configure an MSTP name and revision number on devices running Enterprise software release 08.0.00 and later, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp configuration name foundry revision 4
```

Syntax: [no] mstp configuration name <name> revision <revision-number>

The **name** parameter defines an ASCII name for the MSTP configuration. The default name is the MAC address of the switch expressed as a string.

The **revision** parameter specifies the revision level for MSTP that you are configuring on the switch. It can be a number between 0 and 65535.

Configuring an MSTP Instance

An MSTP instance is configured with an MSTP ID for each region. Each region can contain one or more VLANs. To configure an MSTP instance and assign a range of VLANs, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp configuration instance 7 vlan 4 to 7
```

Syntax: [no] mstp configuration instance <instance-number> vlan <vlan-id> | vlan-group <group-id>

The **instance** parameter defines the number for the instance of MSTP that you are configuring.

The **vlan** parameter assigns one or more VLANs or a range of VLANs to the instance defined in this command.

The **vlan-group** parameter assigns one or more VLAN groups to the instance defined in this command.

Disabling MSTP on a Port

To disable MSTP on a specific port, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp disable eth 2/1
```

Syntax: [no] mstp disable ethernet <slot/port>

The <slot/port> variable is location of the port that you want to disable MSTP for.

Configuring Port Priority and Path cost for an MSTP Instance

Port priority and path cost can be configured for a specified instance. To configure an MSTP instance, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp instance 1 ethernet 3/1 priority 32 path-cost 200
```

Syntax: [no] mstp instance <instance-number> [ethernet <slot/port> priority <port-priority>] path-cost <cost>

The <instance-number> variable is the number of the instance of MSTP that you are configuring port priority and path cost for.

The **ethernet** <slot/port> parameter assigns a port to this instance.

You can set a **priority** to the instance that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 240 in increments of 16. The default value is 128.

A **path-cost** can be assigned to a port to bias traffic towards or away from a path during periods of rerouting. Possible values are 1 - 200000000. The default values for Ethernet interfaces are: 100 for 10 Mbps, 19 for 100 Mbps, 4 for Gigabit Ethernet, and 2 for 10 Gigabit Ethernet.

Configuring Bridge Priority for an MSTP Instance

Bridge priority can be configured for a specified instance.

To configure priority for an MSTP instance on devices running Enterprise software release 08.0.00 and later, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp instance 1 priority 32768
```

Syntax: [no] mstp instance <instance-number> [ethernet <slot/port> priority <port-priority>] path-cost <cost>

The <instance-number> variable is the number for the instance of MSTP that you are configuring.

You can set a **priority** to the instance that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 61440 in increments of 4096. The default value is 32768.

Setting the MSTP Global Parameters

MSTP has many of the options available in RSTP as well as some unique options. The following command is used to configure MSTP Global parameters for all instances configured on a switch:

```
BigIron(config)# mstp force-version 0 forward-delay 10 hello-time 4 max-age 8 max-hops 9
```

Syntax: [no] mstp force-version <mode-number> forward-delay <value> hello-time <value> max-age <value> max-hops <value>

The **force-version** parameter forces the bridge to send BPDUs in a specific format. You can specify one of the following <mode-number> values:

- 0 – The STP compatibility mode. Only STP BPDUs will be sent.
- 2 – The RSTP compatibility mode. Only RSTP BPDUS will be sent.
- 3 – MSTP mode. In this default mode, only MSTP BPDUS will be sent.

The **forward-delay** <value> specifies how long a port waits before it forwards an RST BPDUS after a topology change. This can be a value from 4 – 30 seconds. The default is 15 seconds.

The **hello-time** <value> parameter specifies the interval between two hello packets. The parameter can have a value between 1 – 10 seconds. The default is 2 seconds.

The **max-age** <value> parameter specifies the amount of time the device waits to receive a hello packet before it initiates a topology change. You can specify a value between 6 – 40 seconds. The default value is 20 seconds.

The **max-hops** <value> parameter specifies the maximum hop count. You can specify a value between 1 – 40 hops. The default value is 20 hops.

Setting Ports To Be Operational Edge Ports

You can define specific ports as edge ports for the region that they are configured in to connect to devices that are not running STP, RSTP or MSTP such as a host. If a port is connected to an end device such as a PC, the port can be configured as an edge port. The following command is used to configure ports as operational edge ports:

```
BigIron(config)# mstp admin-edge-port ethernet 3/1
```

Syntax: [no] mstp admin-edge-port ethernet <slot/port>

The <slot/port> parameter specifies a port or range of ports as edge ports in the instance they are configured in.

Setting Point-to-Point Link

You can set a point-to-point link between ports to increase the speed of convergence. To create a point-to-point link between ports, use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp admin-pt2pt-mac ethernet 2/5 ethernet 4/5
```

Syntax: [no] mstp admin-pt2pt-mac ethernet <slot/port>

The <slot/port> parameter specifies a port or range of ports to be configured for point-to-point links to increase the speed of convergence.

Committing MSTP Configuration ID Changes

You need to use the **mstp commit** command only if changes are made to an MSTP configuration identifier while MSTP is operational. If MSTP is not operational, any configuration ID changes are applied immediately and there's no need to use "commit" command.

After you make MSTP configuration identifier changes, you must use the following command at the Global Configuration level to commit the change:

```
BigIron(config)# mstp commit
```

Syntax: [no] mstp commit

Forcing Ports to Transmit the Configured BPDU Version

To force a port to transmit a BPDU of the version that is configured on this device (3 is default), use a command such as the following at the Global Configuration level:

```
BigIron(config)# mstp force-migration-check ethernet 3/1
```

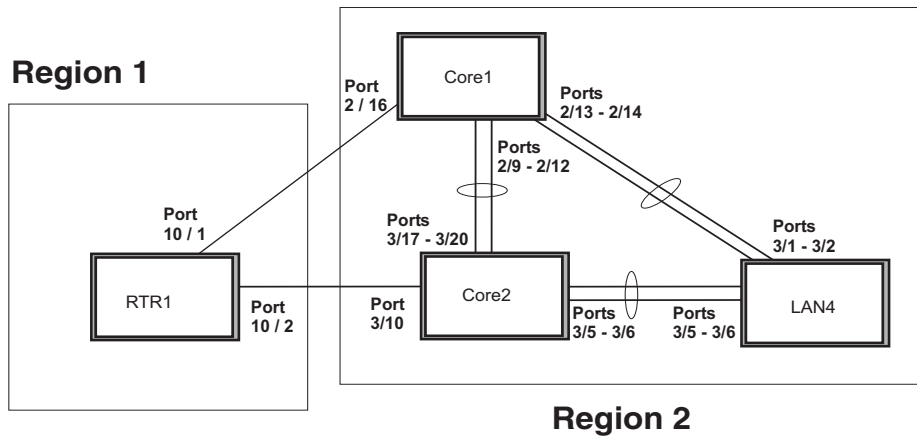
Syntax: [no] mstp force-migration-check ethernet <slot/port>

The <slot/port> variable specifies the port or ports that you want to transmit an MSTP BPDU from.

Sample Configuration

In the example shown in "SAMPLE MSTP Configuration" on page 11-6, four BigIron switch routers are configured in two regions. There are four VLANs in four instances in Region 2 and Region 1 is in the CIST.

Figure 13 SAMPLE MSTP Configuration



RTR1 Configuration

```
BigIron(config)#vlan 4093 name DEFAULT-VLAN by port
BigIron(config-vlan-4093)#tagged ethe 10/1 to 10/2
BigIron(config-vlan-4093)#exit
BigIron(config)#mstp config name Reg1 revision 1
BigIron(config)#mstp admin-pt2pt-mac ethe 10/1 to 10/2
BigIron(config)#mstp start
BigIron(config)#default-vlan-id 4093
BigIron(config)#hostname RTR1
```

Core 1 Configuration

```
BigIron(config)#trunk switch ethernet 2/9 to 2/12 ethernet 2/13 to 2/14
BigIron(config)#vlan 1 name DEFAULT-VLAN by port
BigIron(config-vlan-1)#no spanning-tree
BigIron(config-vlan-1)# exit
BigIron(config)#vlan 20 by port
BigIron(config-vlan-20)#tagged ethernet 2/9 to 2/14 ethernet 2/16
BigIron(config-vlan-20)#no spanning-tree
BigIron(config-vlan-20)# exit
BigIron(config)#vlan 21 by port
BigIron(config-vlan-21)#tagged ethernet 2/9 to 2/14 ethernet 2/16
BigIron(config-vlan-21)#no spanning-tree
BigIron(config-vlan-21)#exit
BigIron(config)#vlan 22 by port
BigIron(config-vlan-22)#tagged ethernet 2/9 to 2/14 ethernet 2/16
BigIron(config-vlan-22)#no spanning-tree
BigIron(config-vlan-22)#exit
BigIron(config)#vlan 23 by port
BigIron(config)#mstp config name HR revision 2
BigIron(config)#mstp config instance 20 vlan 20
BigIron(config)#mstp config instance 21 vlan 21
BigIron(config)#mstp config instance 22 vlan 22
BigIron(config)#mstp instance 0 priority 8192
BigIron(config)#mstp admin-pt2pt-mac ethernet 2/9 to 2/14
BigIron(config)#mstp admin-pt2pt-mac ethernet 2/16
BigIron(config)#mstp disable ethe 2/24
BigIron(config)#mstp start
BigIron(config)#hostname CORE1
```

Core2 Configuration

```
BigIron(config)#trunk switch ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
```

```
BigIron(config)#vlan 1 name DEFAULT-VLAN by port
BigIron(config-vlan-1)#no spanning-tree
BigIron(config-vlan-1)#exit
BigIron(config)#vlan 20 by port
BigIron(config-vlan-20)#tagged ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
BigIron(config-vlan-20)#no spanning-tree
BigIron(config-vlan-20)#exit
BigIron(config)#vlan 21 by port
BigIron(config-vlan-21)#tagged ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
BigIron(config-vlan-21)#no spanning-tree
BigIron(config-vlan-21)#exit
BigIron(config)#vlan 22 by port
BigIron(config-vlan-22)#tagged ethe 3/5 to 3/6 ethe 3/17 to 3/20
BigIron(config-vlan-22)#no spanning-tree
BigIron(config-vlan-22)#exit
BigIron(config)#mstp config name HR revision 2
BigIron(config)#mstp config instance 20 vlan 20
BigIron(config)#mstp config instance 21 vlan 21
BigIron(config)#mstp config instance 22 vlan 22
BigIron(config)#mstp admin-pt2pt-mac ethernet 3/17 to 3/20 ethernet 3/5 to 3/6
BigIron(config)#mstp admin-pt2pt-mac ethernet 3/10
BigIron(config)#mstp disable ethe 3/7 ethernet 3/24
BigIron(config)#mstp start
BigIron(config)#hostname CORE2
```

LAN 4 Configuration

```
BigIron(config)# trunk switch ethernet 3/5 to 3/6 ethernet 3/1 to 3/2
BigIron(config)#vlan 1 name DEFAULT-VLAN by port
BigIron(config-vlan-1)#no spanning-tree
BigIron(config-vlan-1)#exit
BigIron(config)#vlan 20 by port
BigIron(config-vlan-20)#tagged ethernet 3/1 to 3/2 ethernet 3/5 to 3/6
BigIron(config-vlan-20)#no spanning-tree
BigIron(config)#exit
BigIron(config)#vlan 21 by port
BigIron(config-vlan-21)#tagged ethernet 3/1 to 3/2 ethe 3/5 to 3/6
BigIron(config-vlan-21)#no spanning-tree
BigIron(config-vlan-21)#exit
BigIron(config)#vlan 22 by port
BigIron(config-vlan-22)#tagged ethernet 3/1 to 3/2 ethe 3/5 to 3/6
BigIron(config-vlan-22)#no spanning-tree
BigIron(config)#mstp config name HR revision 2
BigIron(config)#mstp config instance 20 vlan 20
BigIron(config)#mstp config instance 21 vlan 21
BigIron(config)#mstp config instance 22 vlan 22
BigIron(config)#mstp admin-pt2pt-mac ethernet 3/5 to 3/6 ethernet 3/1 to 3/2
BigIron(config)#mstp start
BigIron(config)#hostname LAN4
```

Displaying MSTP Information

Displaying MSTP Configuration

On devices running Enterprise software release 08.0.00 and later, MSTP information can be displayed using the commands shown below.

```
BigIron# show mstp

Multiple Spanning Tree Protocol (IEEE 802.1s) is OPERATIONAL
Configuration name: [HR], revision level: 2
Total 4 MST instance(s):
  CIST 0 mapped vlans: 1 to 19 36 to 4094
  MSTI 20 mapped vlans: 20
  MSTI 21 mapped vlans: 21
  MSTI 22 mapped vlans: 22
  Disabled ports: ethe 2/24
Force version: 3
Configured hello time 2 forward delay 15 max age 20 max hops 20
Operational hello time 2 forward delay 15 max age 20 max hops 19
```

Syntax: show mstp [<instance-number> | ethernet <slot/port> | pending-changes

The <instance-number> variable specifies the MSTP instance that you want to display information for.

The **ethernet** parameter is used to specify a port that you want to display MSTP information for.

The **pending-changes** parameter specifies that you want to display pending MSTP configuration identifier changes.

Table 13.1: Output from Show MSTP

This Field...	Displays...
Configuration Name	The MSTP configuration name that is configured on the switch.
Revision Level	The MSTP revision level that is configured on the switch.
Total MST Instances	These fields display the number of MST instances operating on the switch, and the VLANs that reside in the CST and each MST instance.
Disabled Ports	The ethernet ports that have been disabled from MSTP.
Force Version	The format that BPDUs are sent in:
Configured:	
Hello Time	The hello value configured on this device.
Forward Delay	The period of time that the device will wait before it forwards a an RST BPDU as configured on this device.
Max Age	The interval that this device will wait for a hello packet before initiating a topology change as configured.
Max Hops	The number of hops between bridges before a packet times out and is stops circulating as configured on this device.
Operational:	

Table 13.1: Output from Show MSTP (Continued)

This Field...	Displays...
Hello Time	The hello value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Hello Time.
Forward Delay	The forward delay value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Forward Delay.
Max Age	The max age value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Max Age.
Max Hops	The max hops value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Max Hops.

Displaying MSTP Information for a Specified Instance

On devices running Enterprise software release 08.0.00 and later, display information for an MSTP instance using the following command.

```
BigIron# show mstp 22
##### MST 22 vlans mapped: 22
Member ports:  ethe 2/9 to 2/14 ethe 3/1 to 3/2
Bridge ID:      address 00e0.52aa.2f00 priority 32790 (32768 sysid 22)
Reg Root ID:   address 00e0.52c2.cf60 priority 22 (0 sysid 22)
                port    2/9    int path cost: 200000
Remaining hops: 19

Interface Role State Designated ID Ext-cost Int-cost Prio.Num Pt2Pt Edge
2/9      Root FWD 001600e052c2cf60 -- 200000 128.73 Y N
2/10     Altr BLK 001600e052c2cf60 -- 200000 128.74 Y N
2/11     Altr BLK 001600e052c2cf60 -- 200000 128.75 Y N
2/12     Altr BLK 001600e052c2cf60 -- 200000 128.76 Y N
2/13     Altr BLK 8016000cdb310400 -- 200000 128.77 Y N
2/14     Altr BLK 8016000cdb310400 -- 200000 128.78 Y N
2/16     Altr BLK 801600e05202dc00 -- 20000 128.129 Y N
```

Syntax: show mstp <number>

Table 13.2: Output from Show MSTP Instance

This Field...	Displays...
Member Ports	Ports that are included in the specified MSTP instance.
Bridge ID	The ID of the bridge.
Reg Root ID	The Reg Root ID is the MAC address of the Root Bridge for the local region. This value is "self" if the device you are accessing is the Root Bridge for the local region.
Remaining Hops	Hops remaining valid for this message in relationship to the regional root bridge.
Interface	The port number of the interface.

Table 13.2: Output from Show MSTP Instance (Continued)

This Field...	Displays...
Role	The current role of the port: <ul style="list-style-type: none"> • Master • Root • Designated • Alternate • Backup
State	The port's current 802.1w state. A port can have one of the following states: <ul style="list-style-type: none"> • Forwarding • Learning • Blocking
Designated ID	The ID of the bridge that sent the best BPDU that was received on this port.
Int-cost	The configured path cost on a link connected to this port within the internal MSTP region.
Prio.Num	The configured priority and number of the port. The default value for the priority is 128.
Pt2Pt	Indicates if the port is configured with a point-to-point link: <ul style="list-style-type: none"> • Y – The port is configured in a point-to-point link • N – The port is not configured in a point-to-point link
Edge	Indicates if the port is configured as an operational edge port: <ul style="list-style-type: none"> • Y – indicates that the port is defined as an edge port. • N – indicates that the port is not defined as an edge port

Displaying MSTP Information for CIST Instance 0

Instance 0 is the Common and Internal Spanning Tree Instance (CIST). When you display information for this instance there are some differences with displaying other instances. The following example displays MSTP information for CIST Instance 0 on devices running Enterprise software release 08.0.00 and later.

```
BigIron# show mstp 0

##### MST 0 vlans mapped: 1 to 4094
Member ports:  ethe 1/1 to 1/18 ethe 1/20 to 1/28 ethe 3/1 to 3/16 ethe 4/1 to 4/2
               ethe 5/1 to 5/48 ethe 7/1 to 7/48
Bridge ID:      address 0004.8089.3300 priority 0 (0 sysid 0)
Root ID:        self
Reg Root ID:    self
Message age: 0  remaining hops: 20

Interface Role  State Designated ID      Ext-cost  Int-cost  Prio.Num  Pt2Pt  Edge
1/10      Desg  FWD  0000000480893300  200000   200000   128.10    Y      N
1/11      Desg  FWD  0000000480893300  200000   200000   128.11    Y      N
1/12      Desg  FWD  0000000480893300  200000   200000   128.12    Y      N
1/13      Desg  FWD  0000000480893300  200000   200000   128.13    N      N
4/1       Desg  FWD  0000000480893300  2000     2000     128.193   Y      N
```

Table 13.3 contains the display parameters that differ from those used for instances other than instance 0. All other parameters are described in Table 13.2.

Table 13.3: Output from Show MSTP Instance for the 0 Instance

This Field...	Displays...
Root ID	In the 0 instance, the Root ID is the MAC address of the Root Bridge for the CIST. This value is "self" if the device you are accessing is the Root Bridge for the CIST.
Reg Root ID	In the 0 instance, the Reg Root ID is the MAC address of the Root Bridge for the local region. This value is "self" if the device you are accessing is the Root Bridge for the local region.
Message age	The amount of the max age value that the message used in transit from the root bridge external to the local region.
Ext-Cost	The configured path cost on a link connected to this port to an external MSTP region.

Chapter 14

Multiple Spanning Tree Protocol on BigIron MG8 and NetIron 40G

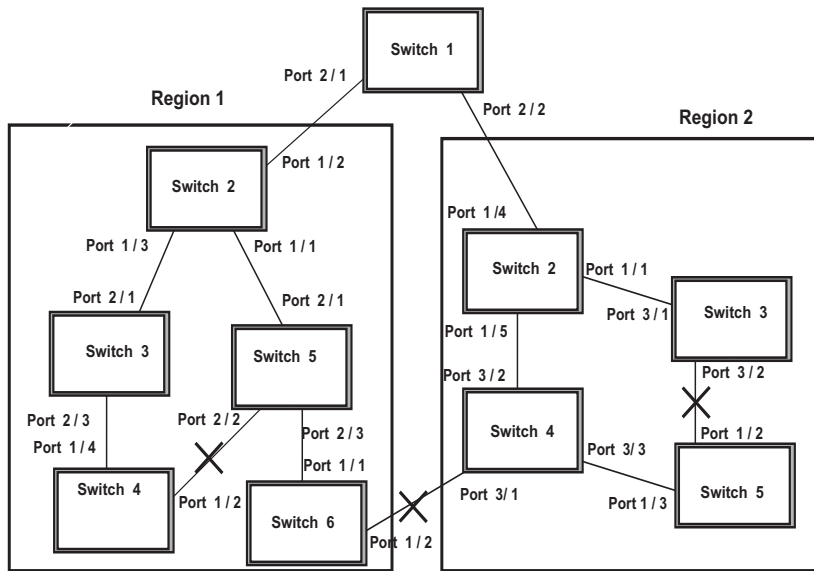
NOTE: This feature is supported on BigIron MG8 and NetIron 40G running software release 02.2.01. For the MSTP feature available on Enterprise software releases, see the “Multiple Spanning Tree Protocol on Enterprise Releases” on page 11-1.

Multiple Spanning Tree Protocol (MSTP) as defined in IEEE 802.1s-2002 allows you to configure multiple STP instances. This ensures loop-free topology for 1 or more VLANs. Using MSTP, the entire network runs a common instance of RSTP. Within that common instance, one or more VLANs can be individually configured into distinct regions. The entire network runs the common spanning tree instance (CST) and the regions run a local instance. The local instance is known as Internal Spanning Tree (IST). The CST treats each instance of IST as a single bridge. Consequently, ports are blocked to prevent loops that might occur within an IST and also throughout the CST. In addition, MSTP can coexist with individual devices running STP or RSTP in the Common and Internal Spanning Trees instance (CIST). With the exception of the provisions for multiple instances, MSTP operates exactly like RSTP.

For example, in Figure 15 a network is configured with two regions: Region 1 and Region 2. The entire network is running an instance of CST. Each of the regions is running an instance of IST. In addition, this network contains Switch 1 running RSTP that isn't configured in a region and consequently is running in the CIST instance. In this configuration, the regions are each regarded as a single bridge to the rest of the network, as is Switch 1. The CST prevents loops from occurring across the network. Consequently, a port is blocked at either port 1/2 of switch 4.

Additionally, loops must be prevented in each of the IST instances. Within the IST Region 1, a port is blocked at port 1/2 of switch 4 to prevent a loop in that region. Within Region 2, a port is blocked at port 3/2 of switch 3 to prevent a loop in that region.

Figure 15 MSTP Configured Network



The following definitions describe the STP instances that define an MSTP configuration:

Common Spanning (CST) – MSTP runs a single instance of spanning tree, called the Common Spanning Tree (CST), across all the bridges in a network. This instance treats each region as a single bridge. In all other ways, it operates exactly like Rapid Spanning Tree (RSTP).

Internal Spanning Tree (IST) – Instances of spanning tree that operate within a defined region are called ISTs (Internal Spanning Tree).

Common and Internal Spanning Trees (CIST) – This is the default MSTP instance 0. It contains all of the ISTs and all bridges that are not formally configured into a region. This instance interoperates with bridges running legacy STP and RSTP implementations.

Multiple Spanning Tree Instance (MSTI) – The MSTI is identified by an MST identifier (MSTid) value between 1 and 4094. This defines an individual instance of an IST. One or more VLANs can be assigned to an MSTI. A VLAN cannot be assigned to multiple MSTIs.

MSTP Region – These are clusters of bridges that run multiple instances of the MSTP protocol. Multiple bridges detect that they are in the same region by exchanging their configuration (instance to VLAN mapping), name, and revision-level. Therefore, if you need to have two bridges in the same region, the two bridges must have identical configurations, names, and revision-levels.

Configuring MSTP

To configure a switch for MSTP, you could configure the name and the revision on each switch that is being configured for MSTP. This name is unique to each switch. You must then create an MSTP Instance and assign an ID. VLANs are then assigned to MSTP instances. These instances must be configured on all switches that interoperate with the same VLAN assignments. Port cost, priority and global parameters can then be configured for individual ports and instances. In addition, operational edge ports and point-to-point links can be created and MSTP can be disabled on individual ports.

Each of the commands used to configure and operate MSTP are described in the following:

- “Enabling MSTP on a Switch”
- “Setting the MSTP Name”
- “Setting the MSTP Revision Number”
- “Configuring an MSTP Instance”

- “Disabling MSTP on a Port”
- “Configuring Port Priority and Port Path Cost”
- “Configuring Bridge Priority for an MSTP Instance”
- “Setting the MSTP Global Parameters”
- “Setting Ports To Be Operational Edge Ports”
- “Setting Point-to-Point Link”
- “Forcing Ports to Transmit an MSTP BPDU”

Enabling MSTP on a Switch

To enable MSTP on your switch, use a command such as the following at the Global Configuration level:

```
BigIron MG8 MG8(config)#mstp start
```

Syntax: [no] mstp start

Setting the MSTP Name

Each switch that is running MSTP is configured with a name. It applies to the switch which can have many different VLANs that can belong to many different MSTP regions. By default, the name is the MAC address of the device.

To configure an MSTP name, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp name foundry
```

Syntax: [no] mstp name <name>

The **name** parameter defines an ASCII name for the MSTP configuration. The default name is the MAC address of the switch expressed as a string.

Setting the MSTP Revision Number

Each switch that is running MSTP is configured with a revision number. It applies to the switch which can have many different VLANs that can belong to many different MSTP regions.

To configure an MSTP revision number, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp revision 4
```

Syntax: [no] mstp revision <revision-number>

The **revision** parameter specifies the revision level for MSTP that you are configuring on the switch. It can be a number from 0 and 65535.

Configuring an MSTP Instance

An MSTP instance is configured with an MSTP ID for each region. Each region can contain one or more VLANs. To configure an MSTP instance and assign a range of VLANs, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config) # mstp instance 7 vlan 4 to 7
```

Syntax: [no] mstp instance <instance-number> [vlan <vlan-id> | vlan-group <group-id>]

The **instance** parameter defines the number for the instance of MSTP that you are configuring.

The **vlan** parameter assigns one or more VLANs or a range of VLANs to the instance defined in this command.

The **vlan-group** parameter assigns one or more VLAN groups to the instance defined in this command.

Disabling MSTP on a Port

To disable MSTP on a specific port, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp disable 2/1
```

Syntax: [no] mstp disable <slot/port>

The <slot/port> variable specifies the location of the port that you want to disable MSTP for.

Configuring Port Priority and Port Path Cost

Priority and path cost can be configured for a specified instance. To configure an MSTP instance, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp instance 7 ethernet 3/1 priority 32 path-cost 200
```

Syntax: [no] mstp instance <instance-number> ethernet <slot/port> priority <port-priority> path-cost <cost>

The <instance-number> variable is the number of the instance of MSTP that you are configuring priority and path cost for.

The **ethernet** <slot/port> parameter specifies a port within a VLAN. The priority and path cost configured with this command will apply to VLAN that the port is contained within.

You can set a **priority** to the port that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 240 in increments of 16. The default value is 128.

A **path-cost** can be assigned to a port to bias traffic towards or away from a path during periods of rerouting. Possible values are 1 - 200000000.

Configuring Bridge Priority for an MSTP Instance

Priority can be configured for a specified instance. To configure priority for an MSTP instance, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp instance 1 priority 8192
```

Syntax: [no] mstp instance <instance-number> priority <priority-value>

The <instance-number> variable is the number for the instance of MSTP that you are configuring.

You can set a **priority** to the instance that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 61440 in increments of 4096. The default value is 32768.

Setting the MSTP Global Parameters

MSTP has many of the options available in RSTP as well as some unique options. To configure MSTP Global parameters for all instances on a switch:

```
BigIron MG8(config)# mstp force-version 0 forward-delay 10 hello-time 4 max-age 8 max-hops 9
```

Syntax: [no] mstp force-version <mode-number> forward-delay <value> hello-time <value> max-age <value> max-hops <value>

The **force-version** parameter forces the bridge to send BPDUs in a specific format. You can specify one of the following <mode-number> values:

- 0 – The STP compatibility mode. Only STP BPDUs will be sent. This is equivalent to single STP.
- 2 – The RSTP compatibility mode. Only RSTP BPDUS will be sent. This is equivalent to single STP.
- 3 – MSTP mode. In this default mode, only MSTP BPDUS will be sent.

The **forward-delay** <value> specifies how long a port waits before it forwards an RST BPDU after a topology change. This can be a value from 4 – 30 seconds. The default is 15 seconds.

The **hello-time** <value> parameter specifies the interval between two hello packets. The parameter can have a value from 1 – 10 seconds. The default is 2 seconds.

The **max-age** <value> parameter specifies the amount of time the device waits to receive a hello packet before it initiates a topology change. You can specify a value from 6 – 40 seconds. The default value is 20 seconds.

The **max-hops** <value> parameter specifies the maximum hop count. You can specify a value from 1 – 40 hops. The default value is 20 hops.

Setting Ports To Be Operational Edge Ports

You can define specific ports as edge ports for the region in which they are configured to connect to devices (such as a host) that are not running STP, RSTP, or MSTP. If a port is connected to an end device such as a PC, the port can be configured as an edge port. To configure ports as operational edge ports enter a command such as the following:

```
BigIron MG8(config)# mstp admin-edge-port ethernet 3/1
```

Syntax: [no] mstp admin-edge-port ethernet <slot/port>

The <slot/port> parameter specifies a port or range of ports as edge ports in the instance they are configured in.

Setting Point-to-Point Link

You can set a point-to-point link between ports to increase the speed of convergence. To create a point-to-point link between ports, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp admin-pt2pt-mac ethernet 2/5 ethernet 4/5
```

Syntax: [no] mstp admin-pt2pt-mac ethernet <slot/port>

The <slot/port> parameter specifies a port or range of ports to be configured for point-to-point links to increase the speed of convergence.

Forcing Ports to Transmit an MSTP BPDU

To force a port to transmit an MSTP BPDU, use a command such as the following at the Global Configuration level:

```
BigIron MG8(config)# mstp force-migration-check ethernet 3/1
```


Syntax: [no] mstp force-migration-check ethernet <slot/port>

The <slot/port> variable specifies the port or ports that you want to transmit an MSTP BPDU from.

Configuration Example

In Figure 16 four BigIron MG8s are configured in two regions. There are four VLANs in four instances in Region 2. Region 1 is in the CIST.

Figure 16 SAMPLE MSTP Configuration



RTR1 Configuration

```
BigIron MG8(config-vlan-4093)tagged ethe 10/1 to 10/2
BigIron MG8(config-vlan-4093)exit
BigIron MG8(config)mstp name Reg1
BigIron MG8(config)mstp revision 1
BigIron MG8(config)mstp instance 0 vlan 4093
BigIron MG8(config)mstp admin-pt2pt-mac ethernet 10/1 to 10/2
BigIron MG8(config)mstp start
BigIron MG8(config)hostname RTR1
```

Core 1 Configuration

```
BigIron MG8(config)trunk switch ethernet 2/9 to 2/12 ethernet 2/13 to 2/14
BigIron MG8(config)vlan 1 name DEFAULT-VLAN by port
BigIron MG8(config-vlan-1)no spanning-tree
BigIron MG8(config-vlan-1) exit
BigIron MG8(config)vlan 20 by port
BigIron MG8(config-vlan-20)tagged ethernet 2/9 to 2/14 ethernet 2/16
BigIron MG8(config-vlan-20)no spanning-tree
BigIron MG8(config-vlan-20) exit
BigIron MG8(config)vlan 21 by port
BigIron MG8(config-vlan-21)tagged ethernet 2/9 to 2/14 ethernet 2/16
BigIron MG8(config-vlan-21)no spanning-tree
BigIron MG8(config-vlan-21)exit
BigIron MG8(config)vlan 22 by port
BigIron MG8(config-vlan-22)tagged ethernet 2/9 to 2/14 ethernet 2/16
BigIron MG8(config-vlan-22)no spanning-tree
BigIron MG8(config-vlan-22)exit
BigIron MG8(config)vlan 23 by port
BigIron MG8(config) mstp name HR
BigIron MG8(config) mstp revision 2
BigIron MG8(config) mstp instance 20 vlan 20
BigIron MG8(config) mstp instance 21 vlan 21
BigIron MG8(config) mstp instance 22 vlan 22
BigIron MG8(config) mstp instance 0 priority 8192
BigIron MG8(config) mstp admin-pt2pt-mac ethernet 2/9 to 2/14
BigIron MG8(config) mstp admin-pt2pt-mac ethernet 2/16
BigIron MG8(config) mstp disable ethe 2/24
BigIron MG8(config) mstp start
```

```
BigIron MG8(config) hostname CORE1
```

Core2 Configuration

```
BigIron MG8(config)trunk switch ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
BigIron MG8(config) vlan 1 name DEFAULT-VLAN by port
BigIron MG8(config-vlan-1)no spanning-tree
BigIron MG8(config-vlan-1) exit
BigIron MG8(config)vlan 20 by port
BigIron MG8(config-vlan-20)tagged ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
BigIron MG8(config-vlan-20)no spanning-tree
BigIron MG8(config-vlan-20)exit
BigIron MG8(config)vlan 21 by port
BigIron MG8(config-vlan-21)tagged ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
BigIron MG8(config-vlan-21)no spanning-tree
BigIron MG8(config-vlan-21)exit
BigIron MG8(config) vlan 22 by port
BigIron MG8(config-vlan-22)tagged ethe 3/5 to 3/6 ethe 3/17 to 3/20
BigIron MG8(config-vlan-22)no spanning-tree
BigIron MG8(config-vlan-22)exit
BigIron MG8(config)mstp name HR
BigIron MG8(config)mstp revision 2
BigIron MG8(config)mstp instance 20 vlan 20
BigIron MG8(config)mstp instance 21 vlan 21
BigIron MG8(config)mstp instance 22 vlan 22
BigIron MG8(config)mstp admin-pt2pt-mac ethernet 3/17 to 3/20 ethernet 3/5 to 3/6
BigIron MG8(config)mstp admin-pt2pt-mac ethernet 3/10
BigIron MG8(config)mstp disable ethe 3/7 ethernet 3/24
BigIron MG8(config)mstp start
BigIron MG8(config)hostname CORE2
```

LAN 4 Configuration

```
BigIron MG8(config) trunk switch ethernet 3/5 to 3/6 ethernet 3/1 to 3/2
BigIron MG8(config)vlan 1 name DEFAULT-VLAN by port
BigIron MG8(config-vlan-1)no spanning-tree
BigIron MG8(config-vlan-1)exit
BigIron MG8(config)vlan 20 by port
BigIron MG8(config-vlan-20)tagged ethernet 3/1 to 3/2 ethernet 3/5 to 3/6
BigIron MG8(config-vlan-20)no spanning-tree
BigIron MG8(config)exit
BigIron MG8(config)vlan 21 by port
BigIron MG8(config-vlan-21)tagged ethernet 3/1 to 3/2 ethe 3/5 to 3/6
BigIron MG8(config-vlan-21)no spanning-tree
BigIron MG8(config-vlan-21)exit
BigIron MG8(config)vlan 22 by port
BigIron MG8(config-vlan-22)tagged ethernet 3/1 to 3/2 ethe 3/5 to 3/6
BigIron MG8(config-vlan-22)no spanning-tree
BigIron MG8(config)mstp config name HR
BigIron MG8(config)mstp revision 2
BigIron MG8(config)mstp instance 20 vlan 20
BigIron MG8(config)mstp instance 21 vlan 21
BigIron MG8(config)mstp instance 22 vlan 22
BigIron MG8(config)mstp admin-pt2pt-mac ethernet 3/5 to 3/6 ethernet 3/1 to 3/2
BigIron MG8(config)mstp start
BigIron MG8(config)hostname LAN4
```

Displaying MSTP Information

Displaying MSTP Statistics

MSTP statistics can be displayed using the commands shown below.

To display all general MSTP information, enter the following command:

```
BigIron MG8(config)#show mstp
MSTP Instance 0 (CIST) - VLANs: 1
-----
Bridge          Bridge Bridge Bridge Bridge Root   Root   Root   Root
Identifier      MaxAge Hello  FwdDly Hop    MaxAge Hello FwdDly Hop
hex            sec   sec   sec   cnt   sec   sec   sec   cnt
8000000cdb80af01 20    2     15    20    20    2     15    19

Root           ExtPath  RegionalRoot      IntPath  Designated      Root
Bridge         Cost     Bridge            Cost     Bridge          Port
hex            hex
8000000480bb9876 2000    8000000cdb80af01 0         8000000480bb9876 3/1

Port  Pri  PortPath  P2P Edge Role      State      Designa-  Designated
Num   Cost  Mac Port   ROOT      FORWARDING  ted cost  bridge
3/1   128  2000     T   F        ROOT        0         8000000480bb9876

MSTP Instance 1 - VLANs: 2
-----
Bridge          Max RegionalRoot      IntPath  Designated      Root  Root
Identifier      Hop Bridge            Cost     Bridge          Port  Hop
hex            cnt hex
8001000cdb80af01 20  8001000cdb80af01 0         8001000cdb80af01 Root  20

Port  Pri  PortPath  Role      State      Designa-  Designated
Num   Cost  MASTER   FORWARDING  ted cost  bridge
3/1   128  2000     MASTER   FORWARDING  0         8001000cdb80af01
```

Syntax: show mstp <instance-number>

The <instance-number> variable specifies the MSTP instance that you want to display information for.

Table 16.1: Output from Show MSTP

This Field...	Displays...
MSTP Instance	The ID of the MSTP instance whose statistics are being displayed. For the CIST, this number is 0.
VLANs:	The number of VLANs that are included in this instance of MSTP. For the CIST this number will always be 1.
Bridge Identifier	The MAC address of the bridge.
Bridge MaxAge sec	Displays configured Max Age.
Bridge Hello sec	Displays configured Hello variable.
Bridge FwdDly sec	Displays configured FwdDly variable.

Table 16.1: Output from Show MSTP (Continued)

This Field...	Displays...
Bridge Hop cnt	Displays configured Max Hop count variable.
Root MaxAge sec	Max Age configured on the root bridge.
Root Hello sec	Hello interval configured on the root bridge.
Root FwdDly sec	FwdDly interval configured on the root bridge.
Root Hop Cnt	Current hop count from the root bridge.
Root Bridge	Bridge identifier of the root bridge.
ExtPath Cost	The configured path cost on a link connected to this port to an external MSTP region.
Regional Root Bridge	The Regional Root Bridge is the MAC address of the Root Bridge for the local region.
IntPath Cost	The configured path cost on a link connected to this port within the internal MSTP region.
Designated Bridge	The MAC address of the bridge that sent the best BPDU that was received on this port.
Root Port	Port indicating shortest path to root. Set to "Root" if this bridge is the root bridge.
Port Num	The port number of the interface.
Pri	The configured priority of the port. The default is 128.
PortPath Cost	Configured or auto detected path cost for port.
P2P Mac	Indicates if the port is configured with a point-to-point link: <ul style="list-style-type: none"> • T – The port is configured in a point-to-point link • F – The port is not configured in a point-to-point link
Edge	Indicates if the port is configured as an operational edge port: <ul style="list-style-type: none"> • T – indicates that the port is defined as an edge port. • F – indicates that the port is not defined as an edge port
Role	The current role of the port: <ul style="list-style-type: none"> • Master • Root • Designated • Alternate • Backup • Disabled

Table 16.1: Output from Show MSTP (Continued)

This Field...	Displays...
State	The port's current 802.1w state. A port can have one of the following states: <ul style="list-style-type: none"> • Forwarding • Discarding • Learning • Disabled
Designated Cost	Port path cost to the root bridge.
Max Hop cnt	The maximum hop count configured for this instance.
Root Hop cnt	Hop count from the root bridge.

Displaying MSTP Information for a Specified Instance

The following example displays MSTP information specified for an MSTP instance.

```
BigIron MG8(config)#show mstp 1
MSTP Instance 1 - VLANs: 2
-----
Bridge          Max RegionalRoot   IntPath   Designated   Root   Root
Identifier      Hop Bridge          Cost      Bridge       Port   Hop
hex             cnt hex
8001000cdb80af01 20 8001000cdb80af01 0      8001000cdb80af01 Root 20

Port  Pri PortPath  Role      State      Designa-  Designated
Num   Cost
3/1   128 2000     MASTER    FORWARDING 0         8001000cdb80af01
```

See Table 16.1 for details about the display parameters.

Displaying MSTP Information for CIST Instance 0

Instance 0 is the Common and Internal Spanning Tree Instance (CIST). When you display information for this instance there are some differences with displaying other instances. The following example displays MSTP information for CIST Instance 0.

```
BigIron MG8(config)#show mstp 0
MSTP Instance 0 (CIST) - VLANs: 1
-----
Bridge          Bridge Bridge Bridge Bridge Root   Root  Root  Root
Identifier      MaxAge Hello  FwdDly Hop    MaxAge Hello FwdDly Hop
hex             sec   sec   sec   cnt   sec   sec  sec   cnt
8000000cdb80af01 20    2    15   20    20    2    15   19

Root           ExtPath  RegionalRoot  IntPath  Designated  Root
Bridge         Cost     Bridge        Cost     Bridge      Port
hex            hex      hex            hex      hex
8000000480bb9876 2000    8000000cdb80af01 0        8000000480bb9876 3/1

Port  Pri  PortPath  P2P Edge Role      State      Designa-  Designated
Num   Cost Mac Port      State      ted cost  bridge
3/1  128 2000    T  F  ROOT      FORWARDING 0        8000000480bb9876
```

Syntax: show mstp 0

Displaying MSTP Configuration

To display details about the MSTP configuration, enter the following command:

```
BigIron MG8(config)#show mstp conf
MSTP CONFIGURATION
-----
Name       : Reg1
Revision  : 1
Version   : 3 (MSTP mode)
Status    : Started

Instance VLANs
-----
0          4093
```

To display details about the MSTP that is configured on the device, enter the following command:

```
BigIron MG8(config)#show mstp detail
MSTP Instance 0 (CIST) - VLANs: 4093
-----
Bridge: 800000b000c00000 [Priority 32768, SysId 0, Mac 00b000c00000]
FwdDelay 15, HelloTime 2, MaxHops 20, TxHoldCount 6

Port 6/54 - Role: DESIGNATED - State: FORWARDING
PathCost 20000, Priority 128, OperEdge T, OperPt2PtMac F, Boundary T
Designated - Root 800000b000c00000, RegionalRoot 800000b000c00000,
Bridge 800000b000c00000, ExtCost 0, IntCost 0
ActiveTimers - helloWhen 1
MachineState - PRX-DISCARD, PTX-IDLE, PPM-SENDING_RSTP, PIM-CURRENT
PRT-ACTIVE_PORT, PST-FORWARDING, TCM-INACTIVE
BPDUs - Rcvd MST 0, RST 0, Config 0, TCN 0
Sent MST 6, RST 0, Config 0, TCN 0
```

See Table 16.1 for explanation about the parameters in the output.

Syntax: show mstp [<mstp-id> | configuration | detail] [| begin <string> | exclude <string> | include <string>]

Enter an MSTP ID for <mstp-id>.

Chapter 17

Configuring Trunk Groups and Dynamic Link Aggregation

This chapter describes how to configure trunk groups and 802.3ad link aggregation.

- Trunk groups are manually-configured aggregate links containing multiple ports.
- 802.3ad link aggregation is a protocol that dynamically creates and manages trunk groups.

NOTE: You can use both types of trunking on the same device. However, you can use only one type of trunking for a given port. For example, you can configure port 1/1 as a member of a static trunk group or you can enable 802.3ad link aggregation on the port, but you cannot do both.

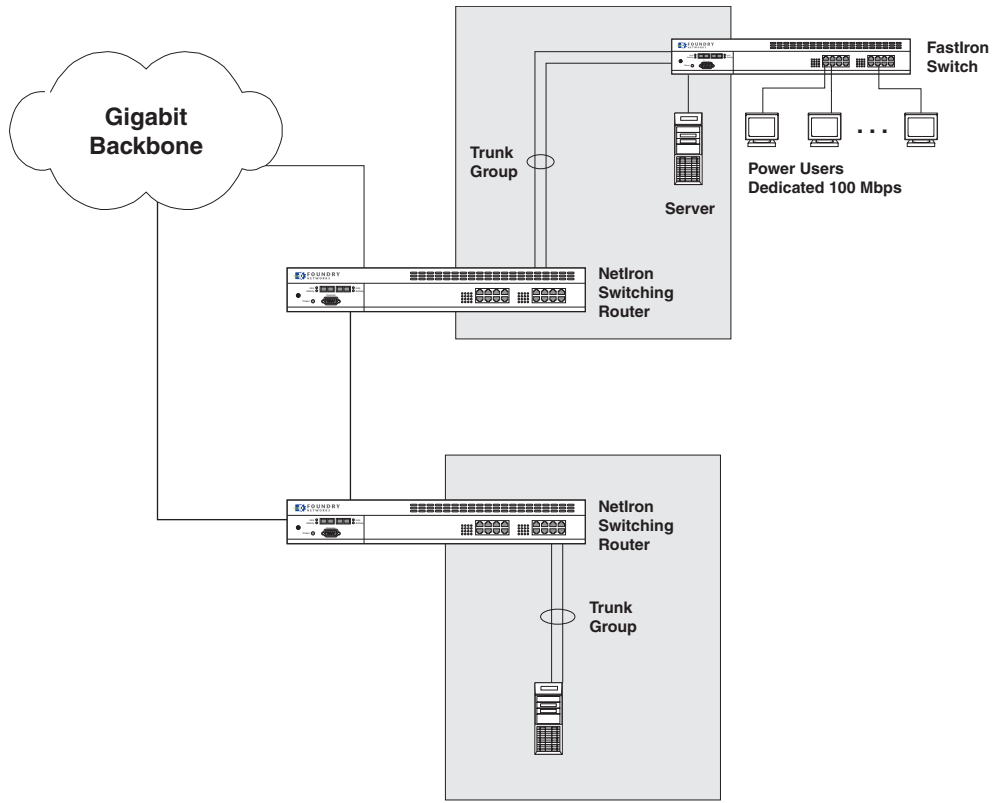
Configuring Trunk Groups

The Trunk Group feature allows you to manually configure multiple high-speed load-sharing links between two Foundry Layer 2 Switches or Layer 3 Switches or between a Foundry Layer 2 Switch and Layer 3 Switch and a server. You can configure up to 8 ports as a trunk group, supporting transfer rates of up to 8 Gbps of bi-directional traffic.

In addition to enabling load sharing of traffic, trunk groups provide redundant, alternate paths for traffic if any of the segments fail.

Figure 17.1 shows an example of a configuration that uses trunk groups.

Figure 17.1 Trunk Group application within a NetIron and FastIron network




NOTE: The ports in a trunk group make a single logical link. Therefore, all the ports in a trunk group must be connected to the same device at the other end.

Trunk Group Connectivity to a Server

To support termination of a trunk group, the server must have either multiple network interface cards (NICs) or either a dual or quad interface card installed. The trunk server is designated as a server with multiple adapters or a single adapter with multiple ports that share the same MAC and IP address. Figure 17.2 shows an example of a trunk group between a server and a Foundry device.

Figure 17.2 Trunk group between a server and a Foundry Stackable Layer 2 Switch or Layer 3 Switch



Trunk Group Rules

- You cannot configure a port as a member of a trunk group if 802.3ad link aggregation is enabled on the port.
- The following table lists the maximum number of trunk groups you can configure on a Foundry device, the valid number of ports in a trunk group, and the port ranges on a device.

Table 17.1: Trunk Group Support

Model	Maximum Number of Trunk Groups		Valid Number of Ports in a Group		Port Ranges and Primary Ports	
	10/100	Gigabit	10/100	Gigabit	10/100	Gigabit
Chassis devices	64		2, 4, or 8		1 – 8, 9 – 16, 17 – 241 – 8, 9 – 16, 17 – 24	
FastIron 4802	8	1	2 or 4	2	1 – 4, 5 – 8, 9 – 12, 13 – 16, 17 – 20, 21 – 24, 25 – 28, 29 – 32, 33 – 36, 37 – 40, 41 – 44, 45 – 48	49 – 50
FES9604	12	2	2, 3, 4, 5, 6, 7, or 8	2, 3, or 4	1 – 8, 9 – 16, 17 – 24, 25 – 32, 33 – 40, 41 – 48, 49 – 56, 57 – 64, 65 – 72, 73 – 80, 81 – 88, 89 – 96	97 – 100
FES4802 and FES4802-POE	6	1	2, 3, 4, 5, 6, 7, or 8	2	1 – 8, 9 – 16, 17 – 24, 25 – 32, 33 – 40, 41 – 48	49 – 50
FES2402 and FES2402-POE	3	1	2, 3, 4, 5, 6, 7, or 8	2	1 – 8, 9 – 16, 17 – 24	25 – 26
FES12GCF		6		2 – 8		
FESX424		13		2, 3, or 4		
Other Stackable Devices	4		2 or 4		1 – 4, 5 – 8, 9 – 12, 13 – 16, 17 – 20, 21 – 24, 25 – 26	
TurboIron/4	3		2		1 – 2, 3 – 4, 5 – 6	
TurboIron/8	4		2 or 4		1 – 4 and 5 – 8	

- Each trunk group must start with a primary port. The primary port is always the lowest number in the port range. For example, trunk groups on a Chassis device can have the following primary ports:
 - 2-port trunk groups: 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, 23
 - 4-port trunk groups: 1, 5, 9, 13, 17, 21
 - 8-port trunk groups: 1, 9, 17

NOTE: On the FastIron Edge Switch (FES) and FES X-Series, you can select any port within a range to be the primary port of the trunk group.


- All ports in a trunk group must be in the same port range. See Table 17.1. Note that this restriction does not apply to the FES12GCF or FESX424, which do not have port ranges.
- On the FastIron Edge Switch (FES) device, you can configure only one trunk group within a port range. The exception is that you can configure two Gigabit trunk groups in the port range 97 – 100 on the 9604.
- You cannot combine 10/100 ports and Gigabit ports in the same trunk group.
- You cannot combine Gigabit and 10-Gigabit ports in the same trunk group.
- Except for the FES X-Series, port assignment on a module must be contiguous. The port range on the module cannot contain gaps. For example, you can configure ports 1, 2, 3, and 4 on a module together as a trunk group but not ports 1, 3, and 4 (excluding 2).
- Port assignment cannot be across multiple trunk group boundaries. For example, on a FastIron 4802, ports 4 and 5 cannot be in the same trunk group.
- All the ports must be connected to the same device at the other end.
- All trunk group member properties must match the lead port of the trunk group with respect to the following parameters:
 - Port tag type (untagged or tagged port)
 - Port speed and duplex
 - QoS priority

To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.

- Make sure the device on the other end of the trunk link can support the same number of ports in the link. For example, if you configure a five-port trunk group on the FastIron Edge Switch switch and the other end is a different type of switch, make sure the other switch can support a five-port trunk group .
- Trunking is supported on POS OC-3, OC-12, and OC-48 ports. Switch and server trunking of POS ports is supported only on Layer 2 software images. Server trunking for POS ports requires software release 07.6.01 or later.
- You can trunk two 10 Gigabit Ethernet ports together. The first port must be in an odd-numbered chassis slot and the second port must be in the following even-numbered slot. Trunking of 10-Gigabit Ethernet ports requires software release 07.6.01 or later. See “Configuring a Trunk Group of 10-Gigabit Ethernet Ports” on page 17-25.

Figure 17.3 shows some examples of valid 2-port trunk group links between devices. The trunk groups in this example are switch trunk groups, between two Foundry devices. Ports in a valid 2-port trunk group on one device are connected to two ports in a valid 2-port trunk group on another device. The same rules apply to 4-port trunk groups.

Figure 17.3 Examples of 2-port trunk groups






Figure 17.4 shows examples of two Layer 3 Switches connected by multi-slot trunk groups.

Figure 17.4 Examples of multi-slot trunk groups





Figure 17.5 shows the valid 2-port and 4-port trunk groups on chassis 10/100 modules.

Figure 17.5 Valid 2-port and 4-port trunk groups on chassis 10/100 modules



Additional Trunk Group Rules for Multi-Slot Trunk Groups

- Multi-slot trunk groups are supported only on Chassis devices.
- You can configure a multi-slot trunk group on two Gigabit Ethernet modules.
- You can configure a maximum of eight ports in the trunk group.
- You can configure up to two groups of ports to make the trunk group and the groups must be alike. For example, you can group two sets of two ports together or two sets of four ports together but you cannot group a set of two ports with a set of four ports. Each group of ports can contain two or four ports.
- Each group of ports must begin with a primary port. On Gigabit Ethernet modules, the primary ports are 1, 3, 5, and 7.
- When you specify the ports in the trunk group, you must specify them in ascending numerical order, beginning with the primary port. For example, to specify a group containing ports 1/1 – 1/4 and 3/1 – 3/4, you must specify them in the order shown. You cannot specify 3/1 – 3/4 first.
- Port configuration for each trunk group is based on the configuration of the primary port. To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.
- If you plan to configure ports on a module into a server trunk group, use the following additional guidelines:
 - In software releases prior to 07.5.04, the management module(s) and the module that has the server trunk group's ports must be in the same set of slots (slots 1 – 7 or 9 – 15). Do not place the management module(s) and the module containing the trunk ports in separate sets of slots.

This restriction was removed in software release 07.5.04 and later. Software releases 07.5.04 and later support placing the management module and forwarding module in different sets of slots. For single-slot server trunk group, the management module and forwarding module can be on any slot of the chassis.

- Do not place the management module(s) or the module that has the server trunk group's ports in slot 8. The modules must both be in the same set or slots (slots 1–7 or 9–15).

These guidelines apply to a server trunk group that is configured on a single module or on a pair of modules (multi-slot trunk group). You do not need to follow these guidelines for a switch trunk group.

Specifying a Minimum Number of Ports for a Trunk Group

Beginning with Enterprise software release 07.8.01, you can configure the Foundry device to disable all of the ports in a trunk group when the number of active member ports drops below a specified threshold value. For example, if a trunk group has 10 ports, and the threshold for the trunk group is 5, then the trunk group is disabled if the number of available ports in the trunk group drops below 5. If the trunk group is disabled, then traffic is forwarded over a different link or trunk group.

For example, the following commands establish a trunk group consisting of 4 ports, then establish a threshold for this trunk group of 3 ports.

```
BigIron(config)# trunk e 3/31 to 3/34
BigIron(config-trunk-3/31-3/34)# threshold 3
```

In this example, if the number of active ports drops below 3, then all the ports in the trunk group are disabled.

Syntax: [no] threshold <number>

You can specify a threshold from 1 (the default) up to the number of ports in the trunk group.

Notes:

- The **disable module** command can be used to disable the ports on a module. However, on 10 Gigabit modules, the **disable module** command does not cause the remote connection to be dropped. If a trunk group consists of 10 Gigabit ports, and you use the **disable module** command to disable ports in the trunk group, which then causes the number of active ports in the trunk group to drop below the threshold value, the trunk group is not disabled.
- If you establish a threshold for a trunk used in conjunction with the Metro Ring Protocol (MRP), then you must

also enable Remote Fault Notification (RFN) for 1 Gigabit interfaces, or Link Fault Signalling (LFS) for 10 Gigabit interfaces.

BigIron MG8 and NetIron 40G Trunk Formation Rules (Release 02.2.01 and Later)

The trunk formation rules for BigIron MG8 and NetIron 40G have been changed in software release 02.2.01 and later. The new trunk formation rules are:

- Trunks can be formed from any number of ports, as long as they contain at a minimum of 2 ports and a maximum of 8 ports.
- Ports in a trunk must have the same speed, the same negotiation mode, and the same QoS priority; otherwise, the trunk is rejected.
- Rate Limiting and PBR requirements:
Primary port policy will apply to all secondary ports. No trunk is rejected.
- Mirroring/Monitoring requirements:
The trunk is rejected if any trunk port has mirroring or monitoring configured.
- VLAN and inner-VLAN translation:
The trunk is rejected if any trunk port has vlan or inner-vlan translation configured.
- Layer 2 requirements:
The trunk is rejected if the trunk ports:
 - do not have the same untagged VLAN component.
 - do not share the same SuperSpan customer id (or cid).
 - do not share the same vlan membership
 - do not share the same uplink vlan membership
 - do not share the same protocol-vlan configuration
 - are configured as mrp primary and secondary interfaces
- Layer 3 requirements:
The trunk is rejected if any of the secondary trunk port has any Layer 3 configurations, such as Ipv4 or Ipv6 address, ospf, rip, ripng, isis, etc.
- Layer 4 (ACL) requirements:
All trunk ports must have the same ACL configurations; otherwise, the trunk is rejected.
- You can have a maximum of 31 trunks.

BigIron MG8 and NetIron 40G Trunk Formation Rules (Prior to Release 02.2.01)

If your BigIron MG8 and NetIron 40G device is running a release prior to 02.2.xx, the trunk rules in this section apply to your device. Read the rules carefully before configuring trunks.

Trunks can be formed in the BigIron MG8 and NetIron 40G chassis from multiple interface modules of either 2, 4, or 8 ports. Within this limit, not all combinations of ports can form a trunk. Also, all ports of a trunk must have the same attributes, and each network layer may have its specific attribute requirements. This section describes the port architecture and provides procedures to pick valid ports when forming a trunk.

Interface Module Packet Processor to Port Architecture

As stated earlier, not all combinations of ports can be used to form a trunk. This is because of how the ports are associated with Packet Processors (PPCRs) within each module. Each port on an Interface Module is associated with a PPCR. Each PPCR can interface to either 1, 10, or 20 ports depending on the interface module. There can

be between one and four PPCRs on an interface module. Table describes the number of PPCRs on each of the Interface Modules supported by the MG8 and the interface module ports supported by each PPCR.

Table 17.2: PPCR to Port layouts

module type	Number of Packet Processors (PPCR)	Module Port Range Belonging to each PPCR			
		PPCR 1	PPCR 2	PPCR 3	PPCR 4
2 x 10G	2	1	2	N/A	N/A
4 x 10G	2	1	2	3	4
10 x 1G	1	1 - 10	N/A	N/A	N/A
20 x 1G	2	1 - 10	11 - 20	N/A	N/A
40 x 1G	4	1 - 10	11 - 20	21 -30	31 - 40
60 x 1G	3	1 - 20	21 - 40	41 - 60	N/A

Figure 17.6 shows a 40 x 10G Interface Module containing 4 PPCRs. The first PPCR is shown associated with 10 ports and the last 6 ports of the fourth PPCR are shown.

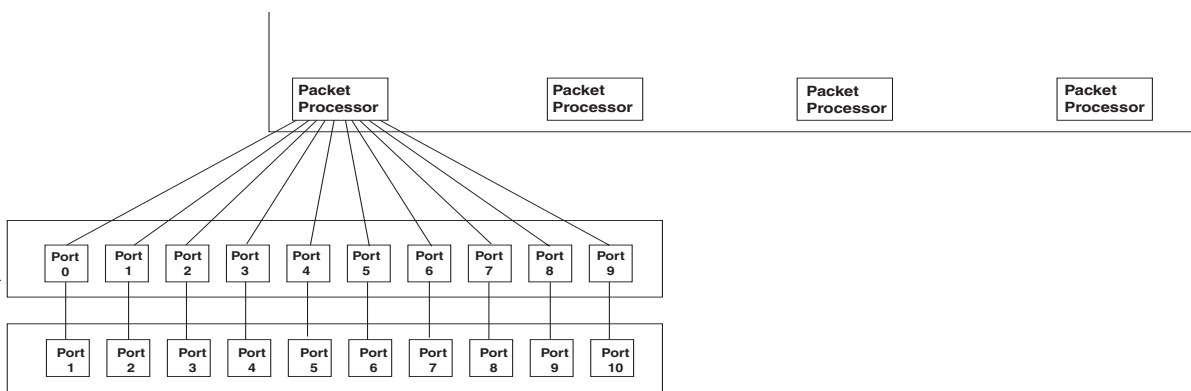


Figure 17.6 Packet Processors on 1G x 40 Interface Module

The ten ports associated with PPCR 1 have an identity as Interface Module ports and as PPCR ports. The Interface Module ports (shown in this figure) are identified as ports 1 to 10 in the first PPCR and 35 to 40 in the fourth. The PPCR ports are numbered from 0 to 9 for each PPCR. For example, interface port 1 of the module has a PPCR port number of 0 and that interface port 40 of the module, which is the last interface module port, (associated with PPCR 4), has a PPCR port number of 9. The Interface Module port numbers are how you identify them for mechanical connection and system configuration. The PPCR port numbers are essential to determining whether a port can be validly used in a trunk.

Either 2, 4, or 8 ports from an individual packet processor can be used in constructing a trunk. If 8 ports are used from a specific packet processor, this will constitute all the ports in an 8 port trunk (the largest possible). If either 2 ports or 4 ports are used from a specific packet processor, they may be combined with ports from another port-set or another PPCR (on the same interface module or from a different interface module) to create a bigger trunk. Since each 10 Gbps port has its own PPCR, ports for 10 Gbps modules can be used in any order as long as you follow all of the other rules for creating a trunk.

Either of the following sections; “Determining Valid Ports Using the Trunk Mask Test” and “Determining valid ports using Valid Port Tables” can be used to determine valid port-sets for use in forming MG8 trunks.

Determining Valid Ports for Trunking

As described earlier, either 2, 4, or 8 ports from an individual PPCR can be used in constructing a trunk. Depending on the number of ports used, you can obtain all of the trunk ports you need from a single packet processor or use ports from up to 4 PPCRs on one or multiple Interface Modules. If 8 ports are used from a specific PPCR, this will constitute all the ports in an 8 port trunk (the largest possible). If either 2 ports or 4 ports are used from a specific PPCR, they may be combined with ports from another PPCR (on the same interface module or from a different interface module) to create a bigger trunk. For example, from the 40 x 1 Gbps Interface Module shown in Figure 17.6, you could use ports 1 and 2 (PPCR ports 0 and 1) from PPCR 1 and ports 39 and 40 (PPCR ports 8 and 9) from PPCR 4 to construct a 4 port trunk.

Trunk ports from a particular PPCR must be added to a trunk by following a very specific set of rules. As shown in Figure 18, five different two-port sets can be used from a PPCR as long as they start with an even PPCR port number. For example, you can use PPCR ports 6 and 7. Two different four-port sets can be used: PPCR ports 0 - 3 or 4 - 7. If you use 8 ports from a single processor, you must use PPCR ports 0 - 7.

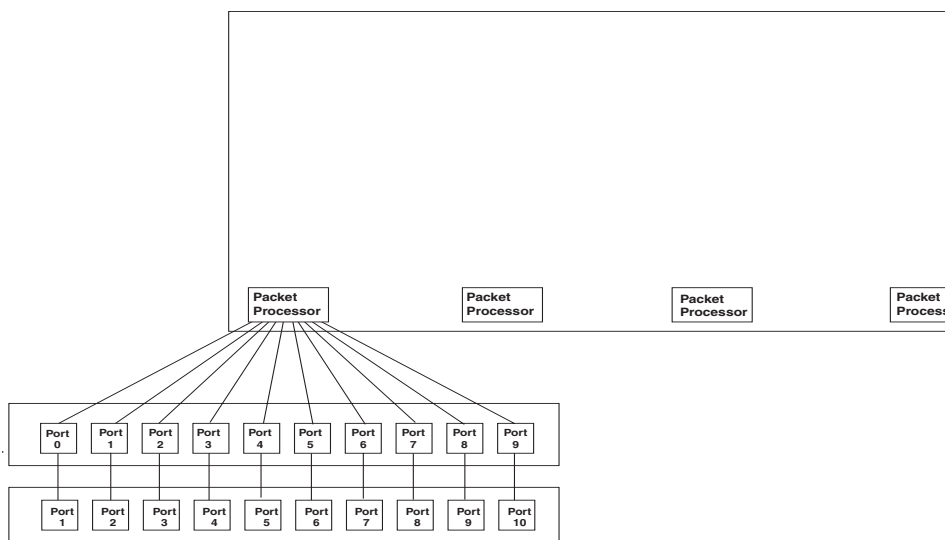


Figure 18 Port Numbering Detail

The following sections: “Determining Valid Ports Using the Trunk Mask Test” and “Determining valid ports using Valid Port Tables” provide methods for determining valid port sets that can be used to form a trunk.

NOTE: Because all 10 Gbps ports are associated with a single packet processor, you can use any combination of 10 Gbps ports within an MG8 chassis to form a trunk of 2, 4, or 8 ports.

Determining Valid Ports Using the Trunk Mask Test

As described in “Determining Valid Ports for Trunking”, you must use particular PPCR ports when forming a trunk. You can do this by applying a Trunk Mask Test as described in this section or by using the valid port tables as described in “Determining valid ports using Valid Port Tables”.

Use the following formula to determine if a set of ports can be used to form a trunk:

```
<PPCR_Port_Number> & ~<trunk_mask> == First PPCR port in range
```

PPCR_Port_Number - This is the PPCR port number as described earlier in this section. You must account for the fact that while the Interface Module ports start at 1 and end with the last port on the Module, PPCR ports are specific to a PPCR and always start on an PPCR at zero, run up to the last port of that PPCR and start at zero for the next PPCR.

For example: To determine the A 40 x 1 Gbps Interface Module has four PPCRs. Each PPCR supports 10 Interface Module ports. Port 38 of the Interface Module is associated with the fourth PPCR. Consequently, port 31 of the

Interface Module is associated with PPCR port 0 and Interface Port 38 is PPCR port 7 of the fourth packet processor.

Trunk_Mask - This variable depends on the number of ports you intend to use from a single PPCR. Table describes the number of PPCRs on each of the Interface Modules supported by the MG8 and the interface module ports supported by each PPCR.

- 0x1 for two ports
- 0x3 for four ports
- 0x7 for eight ports

EXAMPLE:

Interface Module 4 in an MG8 chassis is a 40 x 1 Gbps module. Trunk ports 4/11 to 4/14 and ports 4/25 to 4/28 are able to pass the Trunk Mask Test as shown in the following.

Interface ports 4/11, 4/12, 4/13, and 4.14 are associated with PPCR ports 0, 1, 2, and 3. The Trunk_Mask for four ports is 0x3. The mask is applied to each of the ports as shown in the following:

```
<PPCR_Port_Number> & ~<trunk_mask> == First PPCR port in range
0 & ~0x3 == 0 (This is the 1st PPCR port number in the range)
1 & ~0x3 == 0
2 & ~0x3 == 0
3 & ~0x3 == 0
```

Since each PPCR port belongs to a port set that begins with the same PPCR port (0). Therefore, this is a valid set of four ports to use in a trunk.

Interface ports 4/25, 4/26, 4/27, and 4/28 are associated with PPCR ports 4, 5, 6, and 7. The Trunk_Mask for four ports is 0x3. The mask is applied to each of the ports as shown in the following:

```
<PPCR_Port_Number> & ~<trunk_mask> == First PPCR port in range
4 & ~0x3 == 4 (This is the 1st PPCR port number in the range)
5 & ~0x3 == 4
6 & ~0x3 == 4
7 & ~0x3 == 4
```

Since each PPCR port belongs to a port set that begins with the same PPCR port (4) this is a valid set of four ports to use in a trunk.

Using these two sets of four ports, you can create a valid 8 port trunk.

Determining valid ports using Valid Port Tables

As described in “Determining Valid Ports for Trunking”, you must use particular PPCR ports when forming a trunk. Valid trunking ports can be determined by applying the Trunk Mask Test as described in “Determining Valid Ports Using the Trunk Mask Test” or by using the following tables. The valid port-sets for each of the current Interface Modules are in one the following tables. Port-sets can be used alone to create a trunk of the desired size, or they can be mixed between Interface Modules in a chassis and PPCRs on an individual module, as long as the trunk meets all of the other rules as described in “Other Rules for Forming a Trunk on a BigIron MG8 or NetIron 40G”.

Table 3: 10 x 1 Gbps & 20 X 1 Gbps Interface Module Port-sets for trunking

Module Type	PPCR Number	Valid Port Sets 2-port	Valid Port Sets 4-port	Valid Port Sets 8-port
10 x 1Gbps	PPCR 1	1 - 2 3 - 4 5 - 6 7 - 8 9 - 10	1 - 4 5 - 8	1 - 8

Table 3: 10 x 1 Gbps & 20 X 1 Gbps Interface Module Port-sets for trunking (Continued)

Module Type	PPCR Number	Valid Port Sets 2-port	Valid Port Sets 4-port	Valid Port Sets 8-port
20 x 1G	PPRC 1	1 - 2 3 - 4 5 - 6 7 - 8 9 - 10	1 - 4 5 - 8	
	PPRC 2	11 - 12 13 - 14 15 - 16 17 - 18 19 - 20	11 - 14 15 - 18	11 - 18

Table 4: 40 x 1 Gbps Interface Module Port-sets for trunking

module type	Packet Processor #	Valid Port Sets 2-port	Valid Port Sets 4-port	Valid Port Sets 8-port
40 x 1G	PPRC 1	1 - 2 3 - 4 5 - 6 7 - 8 9 - 10	1 - 4 5 - 8	
	PPRC 2	11 - 12 13 - 14 15 - 16 17 - 18 19 - 20	11 - 14 15 - 18	
	PPRC 3	21 - 22 23 - 24 25 - 26 27 - 28 29 - 30	21 - 24 25 - 28	
	PPRC 4	31 - 32 33 - 34 35 - 36 37 - 38 39 - 40	31 - 34 35 - 38	

Table 5: 60 X 1 Gbps Interface Module Port-sets for trunking

module type	Packet Processor #	Valid Port Sets 2-port	Valid Port Sets 4-port	Valid Port Sets 8-port
60 x 1G	PPRC 1	1 - 2 3 - 4 5 - 6 7 - 8 9 - 10 11 - 12 13 - 14 15 - 16 17 - 18 19 - 20	1 - 4 5 - 8 9 - 12 13 - 16 17 - 20	1 - 8 9 - 16
	PPRC 2	21 - 22 23 - 24 25 - 26 27 - 28 29 - 30 31 - 32 33 - 34 35 - 36 37 - 38 39 - 40	21 - 24 25 - 28 29 - 32 33 - 36 37 - 40	21 - 28 29 - 36
	PPRC 3	41 - 42 43 - 44 45 - 46 47 - 48 49 - 50 51 - 52 53 - 54 55 - 56 57 - 58 59 - 60	41 - 44 45 - 48 49 - 52 53 - 56 57 - 60	41 - 48 49 - 56

Other Rules for Forming a Trunk on a BigIron MG8 or NetIron 40G

Once you have determined the ports you intend to use for your trunk, you must make sure that they meet the requirements defined in the following list.

1. Physical port requirements
All trunk ports must have the same physical port attributes; otherwise, the trunk is rejected.
2. Rate Limiting and PBR requirements
Primary port policy will apply to all secondary ports. No trunk is rejected.
3. Mirroring/Monitoring requirements
The trunk is rejected if any trunk port has mirroring or monitoring configured.
4. VLAN and inner-VLAN translation
The trunk is rejected if any trunk port has vlan or inner-vlan translation configured.
5. Layer 2 requirements
The trunk is rejected if the trunk ports:

- do not have the same untagged VLAN component.
 - do not share the same superspan customer id (or cid).
 - do not share the same vlan membership
 - do not share the same uplink vlan membership
 - do not share the same protocol-vlan configuration
 - are configured as mrp primary and secondary interfaces
6. Layer 3 requirements
- The trunk is rejected if any of the secondary trunk port has any layer 3 configurations, such as Ipv4 or Ipv6 address, ospf, rip, ripng, isis, etc.
7. Layer 4 (ACL) requirements
- All trunk ports must have the same ACL configurations; otherwise, the trunk is rejected.
8. On BigIron MG8 and software release 02.0.02 and later, and NetIron 40G software release 02.0.04 and later, you can have a maximum of 64 server trunks. Prior to these releases, only up to 16 server trunks can be configured.

Trunk Group Load Sharing

Except on the FES X-Series device, when you configure a trunk group, you specify whether the trunk group is a “switch” trunk group or a “server” trunk group:

- Switch trunk group – Use this type of trunk group to connect one Foundry Layer 2 Switch or Layer 3 Switch to another Foundry Layer 2 Switch or Layer 3 Switch.
- Server trunk group – Use this type of trunk group to connect a Foundry Layer 2 Switch or Layer 3 Switch to a file server or single host device.

NOTE: Trunking is not supported for ATM ports.

The Foundry device load shares across the ports in the trunk group. The method used for the load sharing depends on the following:

- Device type – Chassis device or Stackable device
- Traffic type – Layer 2 or Layer 3
- Trunk type – Switch or server
- For certain traffic, port type on which the traffic enters the Foundry device (Gigabit or 10/100)

NOTE: The port type applies only to Layer 2 traffic on a server trunk group configured on a .

NOTE: On a device managed by a VM1, you can optimize server trunk load sharing on individual ports. See “Enabling Optimized Server Trunk Load Balancing (VM1 only)” on page 17-31.

NOTE: Starting in software release 07.7.00, JetCore devices load balance IP traffic on server trunks based on source and destination TCP and UDP application ports (Layer 4 information), as well as on source and destination IP addresses (Layer 3 information). In addition, software release 07.7.00 enables you to configure server trunk load balancing per packet and to specify the maximum number of hash buckets per server trunk. See “Server Trunk Group Load Sharing Enhancements and Options (Release 07.7.00 and Higher)” on page 17-29.

NOTE: Foundry devices that support IPv6 take a packet's IPv6 address into account when sharing traffic across a trunk group. The load sharing is performed in the same way it is for IPv4 addresses; that is, trunk types whose traffic load is shared based on IPv4 address information can now use IPv6 addresses to make the load sharing decision.

For switch trunk groups, load sharing occurs as follows:

- For switched traffic, load sharing is based on the destination MAC address
- For routed traffic, load sharing is based on the destination IPv6 address

To select a port, the device determines the port in the trunk group that has the fewest number of flows. Traffic for the destination address is sent over the selected port.

For server trunk groups, which connect a Foundry Layer 2 Switch or Layer 3 Switch to a file server or single host device, load sharing occurs as follows:

- For switched traffic, load sharing is based on source MAC address
- For routed traffic, load sharing is based on the combination of source and destination IPv6 addresses

To select a port, the device calculates a hash value derived from source and destination IPv6 addresses, then uses the hash value to select a port from the available ports in the trunk group.

Trunk Load Sharing with JetCore Modules

Table 18.1 lists how JetCore devices load balance traffic across the ports in a trunk group.

NOTE: The load sharing methods for server trunk groups also apply to trunks dynamically configured by 802.3ad link aggregation.

Table 18.1: Foundry Trunk Group Load Sharing – JetCore devices

Traffic Type	Trunk Type	Input Port Type	Load Balancing Method
Layer 2	Switch	10/100 Ethernet	Destination MAC address
		Gigabit Ethernet	Destination MAC address
		10 Gigabit Ethernet	Destination MAC address
	Server	10/100 Ethernet	Source MAC address
		Gigabit Ethernet	Source MAC address
		10 Gigabit Ethernet	Source MAC address

Table 18.1: Foundry Trunk Group Load Sharing – JetCore devices (Continued)

Traffic Type	Trunk Type	Input Port Type	Load Balancing Method
Layer 2 IP	Switch	10/100 Ethernet	Destination MAC address
		Gigabit Ethernet	Destination MAC address
		10 Gigabit Ethernet	Destination MAC address
	Server	10/100 Ethernet	Source and destination IP addresses
		Non-fragmented traffic	Source and destination IP addresses and source and destination TCP and UDP application port numbers
		Gigabit Ethernet	Source and destination IP addresses
	10 Gigabit Ethernet	Source and destination IP addresses	
Layer 3 IP	Switch ^a	10/100 Ethernet	Destination IP address
		Gigabit Ethernet	Destination IP address
		10 Gigabit Ethernet	Destination IP address
	Server	10/100 Ethernet	Source and destination IP addresses
		Non-fragmented traffic ^b	Source and destination IP addresses and source and destination TCP and UDP application port number ^b
		Gigabit Ethernet	Source and destination IP addresses
		10 Gigabit Ethernet	Source and destination IP addresses

a.By default, Layer 3 IP traffic uses ip load-sharing by-host. Refer to the *Foundry Switch and Router Command Line Interface Reference* for details on this command.

b.New in 07.6.01

Trunk Load Sharing with IronCore Modules

Table 18.2 lists how IronCore s load balance traffic across the ports in a trunk group.

Table 18.2: Foundry Trunk Group Load Sharing – IronCore Chassis devices

Traffic Layer	Trunk Group Type	Traffic Type	Load-Sharing Basis
Layer 2	Switch	All traffic types	Destination MAC address
	Server	IP received on 10/100 port	Hash value derived from source and destination IP addresses
		IPX received on 10/100 port	Hash value derived from source and destination IPX addresses
		AppleTalk received on 10/100 port	Hash value derived from source and destination AppleTalk addresses
		Other traffic types received on 10/100 port	Hash value derived from source and destination MAC address
		All traffic types received on Gigabit port	Gigabit Port number on which traffic was received
Layer 3	Switch	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other traffic types	Destination MAC address
	Server	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other traffic types	Destination MAC address

Table 18.3 lists how IronCore Stackable devices load balance traffic across the ports in a trunk group.

Table 18.3: Foundry Trunk Group Load Sharing – IronCore Stackable devices

Traffic Layer	Trunk Group Type	Traffic Type	Load-Sharing Basis
Layer 2	Switch	All traffic types	Destination MAC address
	Server	IP	Hash value derived from source and destination IP addresses
		IPX	Hash value derived from source and destination IPX addresses
		AppleTalk	Hash value derived from source and destination AppleTalk addresses
		Other traffic types	Hash value derived from source and destination MAC address
Layer 3	Switch	IP	Source and destination IP addresses
		IPX	Source and destination IPX addresses
		AppleTalk	Source and destination AppleTalk addresses
		Other traffic types	Source and destination MAC address
	Server	IP	Source and destination IP addresses
		IPX	Source and destination IPX addresses
		AppleTalk	Source and destination AppleTalk addresses
		All other	Source and destination MAC address

Trunk Load Sharing with FES Devices

Table 18.2 lists how FES devices load balance traffic across the ports in a switch trunk group.

NOTE: Server trunk groups load balance Layer 2 switched IP and IPX traffic. Other traffic is forwarded the same as on a switch trunk.

Table 18.4: Foundry Switch Trunk Group Load Sharing

Traffic Layer	Traffic Type	Load-Sharing Basis
Layer 2	All traffic types	Destination MAC address
Layer 3	IP	Destination IP address
	IPX	Destination IPX address
	AppleTalk	Destination AppleTalk address
	All other traffic types	Destination MAC address

Trunk Load Sharing with FES X-Series Devices

Table 18.2 lists how FES X-Series devices load balance traffic across the ports in a trunk group. Note that load balancing on the FES X-Series is hardware-based.

Table 1: FES X-Series Trunk Group Load Sharing

Traffic Layer	Traffic Type	Load-Sharing Basis
Layer 2	All traffic types	Destination MAC address Source MAC address
Layer 3	IP	Destination IP address Source MAC address Protocol
	All other traffic types	Destination MAC address Source MAC address

Configuring a Trunk Group

1. Disconnect the cables from those ports on both systems that will be connected by the trunk group. Do not configure the trunk groups with the cables connected.

NOTE: If you connect the cables before configuring the trunk groups and then rebooting, the traffic on the ports can create a spanning tree loop.

2. Configure the trunk group on one of the two Layer 2 Switches or Layer 3 Switches involved in the configuration.
 3. Save the configuration changes to the startup-config file.
 4. Dynamically place the new trunk configuration into effect by entering the **trunk deploy** command at the global CONFIG level of the CLI.
-

NOTE: If you are running a software release earlier than 07.5.00, you must reload the software to place a trunk configuration change into effect.

5. If the device at the other end of the trunk group is another Layer 2 Switch or Layer 3 Switch, repeat Steps 2 – 4 for the other device.
6. When the trunk groups on both devices are operational, reconnect the cables to those ports that are now configured as trunk groups, starting with the first port (lead port) of each trunk group.
7. To verify the link is operational, use the **show trunk** command.

Example 1: Configuring the Trunk Groups Shown in Figure 17.1

To configure the trunk groups shown in Figure 17.1, enter the following commands. Notice that the commands are entered on multiple devices.

USING THE CLI

To configure the trunk group link between NetIron1 and the FastIron:

NOTE: The text shown in italics in the CLI example below shows messages echoed to the screen in answer to the CLI commands entered.

```
NetIron1(config)# trunk switch e 5 to 8
Trunk 2 is created for next power cycle.
Please save configuration to flash and reboot.
NetIron1(config)# write memory
Write startup-config in progress.
.Write startup-config done.
NetIron1(config)# exit
NetIron1# reload
```

NOTE: This example uses devices that are not running software release 07.5.00 or later. Devices running software earlier than 07.5.00 must be reloaded in order to place trunk configuration changes into effect. On devices running 07.5.00 or later, you can dynamically place trunk configuration changes into effect by entering the **trunk deploy** command at the global CONFIG level of the CLI.

To configure the trunk group link between NetIron2 and the server:

```
NetIron2(config)# trunk server e 2 to 4
Trunk 0 is created for next power cycle.
Please save configuration to flash and reboot.
NetIron2(config)# write memory
Write startup-config in progress.
.Write startup-config done.
NetIron2(config)# exit
NetIron2# reload
```

You then configure the trunk group on the FastIron.

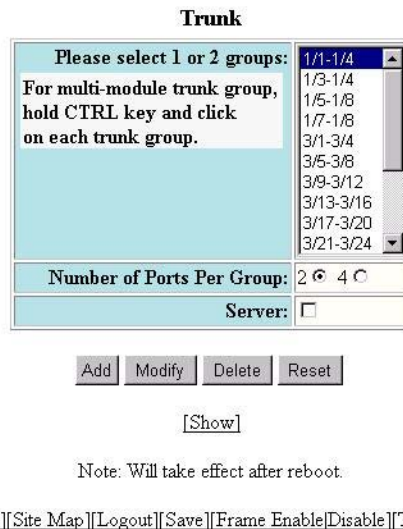
```
FastIron(config)# trunk switch ethernet 17 to 18
```

```
FastIron(config)# write memory
Write startup-config in progress.
.Write startup-config done.
FastIron(config)# exit
FastIron# reload
```

USING THE WEB MANAGEMENT INTERFACE

To configure ports 5 – 8 as a trunk group between two Layer 2 Switches, two Layer 3 Switches, or a Layer 2 Switch or Layer 3 Switch and a server:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the Trunk link.
 - If the device does not have any trunk groups configured, the Trunk configuration panel is displayed, as shown in the following example.
 - If a trunk group is already configured and you are adding a new one, click on the Add Trunk Group link to display the Trunk configuration panel, as shown in the following example.
 - If you are modifying an existing trunk group, click on the Modify button to the right of the row describing the trunk group to display the Trunk configuration panel, as shown in the following example.



NOTE: This panel lists port ranges only for the slots that contain an active module. In addition, only the ranges that are valid for the module are listed.

The port ranges listed by the panel contain four ports, but the default number of ports in a group is two. If you select a group and leave the number of ports in a group at two, the software assigns the first two ports in the group you select to the trunk group. The last two ports do not become members of the trunk group.

4. Select a port range (for example, 5 – 8). On Chassis devices, the port numbers include the slot numbers. For example, you can select 1/5 – 1/8.
5. Select the number of ports you want to use in the trunk group. You can select 2 or 4.
6. Click in the checkbox next to Server to place a checkmark in the box if the other end of the trunk group is a server. If the other end of the connection is a Foundry Layer 2 Switch or Layer 3 Switch, do not click this checkbox.

- Click Apply to save the changes to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
- Click on the plus sign next to Command in the tree view to list the command options.
- Select the [Reload](#) link and select Yes when the Web management interface asks you whether you really want to reload the software.

NOTE: This example uses devices that are not running software release 07.5.00 or later. Devices running software earlier than 07.5.00 must be reloaded in order to place trunk configuration changes into effect. On devices running 07.5.00 or later, you can dynamically place trunk configuration changes into effect by entering the **trunk deploy** command at the global CONFIG level of the CLI.

- If the other end of the trunk group is a Layer 2 Switch or Layer 3 Switch, log in to the other device and follow the steps above.

NOTE: Turbolron/4 Layer 2 and Layer 3 Switches support a maximum of six ports—four standard plus two expansion ports. Therefore, Turbolron/4 Layer 2 and Layer 3 Switches support a maximum of three trunk groups of two ports each. The possible trunk groups are ports 1-2, 3-4 and 5-6.

Example 2: Configuring a Trunk Group That Spans Multiple Gigabit Ethernet Modules in a Chassis Device

To configure a trunk group that spans two modules in a BigIron Layer 3 Switch, use one of the following methods.

USING THE CLI

To configure a trunk group consisting of two groups of ports, 1/1 – 1/4 on module 1 and 4/5 – 4/8 on module 4, enter the following commands:

```
BigIron(config)# trunk ethernet 1/1 to 1/4 ethernet 4/5 to 4/8
BigIron(config-trunk-1/1-4/8)# write memory
BigIron(config-trunk-1/1-4/8)# exit
BigIron(config)# trunk deploy
```

NOTE: The **trunk deploy** command dynamically places trunk configuration changes into effect, without a software reload. This command is supported only in software release 07.5.00 and later. If you are running a release earlier than 07.5.00, you must reload the software to place trunk configuration changes into effect.

Syntax: [no] trunk [server | switch] ethernet | pos <primary-portnum> to <portnum>
ethernet | pos <primary-portnum> to <portnum>

Syntax: trunk deploy

The **server** | **switch** parameter specifies whether the trunk ports will be connected to a server or to another Layer 2 Switch or Layer 3 Switch. This parameter affects the type of load balancing performed by the Foundry device. See “Trunk Group Load Sharing” on page 17-15. The default is **switch**. Note that this parameter is not supported on the FES X-Series.

Each **ethernet** or **pos** parameter introduces a port group.

The <primary-portnum> to <portnum> parameters specify a port group. Notice that each port group must begin with a primary port. After you enter this command, the primary port of the first port group specified (which must be the group with the lower port numbers) becomes the primary port for the entire trunk group. For Gigabit Ethernet modules, the primary ports are 1, 3, 5, and 7.

To configure a trunk group consisting of two groups of two ports each, enter commands such as the following:

```
BigIron(config)# trunk ethernet 1/1 to 1/2 ethernet 3/3 to 3/4
BigIron(config)# write memory
BigIron(config)# trunk deploy
```

Notice that the groups of ports meet the criteria for a multi-slot trunk group. Each group contains the same number of ports (two) and begins on a primary port (1/1 and 3/3).

NOTE: The **trunk deploy** command dynamically places trunk configuration changes into effect, without a software reload. This command is supported only in software release 07.5.00 and later. If you are running a release earlier than 07.5.00, you must reload the software to place trunk configuration changes into effect.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the Trunk link.
 - If the device does not have any trunk groups configured, the Trunk configuration panel is displayed, as shown in the following example.
 - If a trunk group is already configured and you are adding a new one, click on the Add Trunk Group link to display the Trunk configuration panel, as shown in the following example.
 - If you are modifying an existing trunk group, click on the Modify button to the right of the row describing the trunk group to display the Trunk configuration panel, as shown in the following example.

Trunk

Please select 1 or 2 groups:
 For multi-module trunk group,
 hold CTRL key and click
 on each trunk group.

	1/1-1/4
	1/3-1/4
	1/5-1/8
	1/7-1/8
	3/1-3/4
	3/5-3/8
	3/9-3/12
	3/13-3/16
	3/17-3/20
	3/21-3/24

Number of Ports Per Group: 2 4

Server:

[\[Show\]](#)

Note: Will take effect after reboot.

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

4. Select a port range (for example, 5 – 8). On Chassis devices, the port numbers include the slot numbers. For example, you can select 1/5 – 1/8.
5. Select 2 or 4 to indicate the number of ports in each group. Each group must have the same number of ports.
6. Select the port groups. Each group begins with the primary port number for that group. To select two groups, click on the first group, then hold down the CTRL key and click on the second group. Do not select more than two groups.
7. Select Server if you are connecting the trunk group ports to a server. Otherwise, the software assumes you are connecting the trunk group ports to another Layer 2 Switch or Layer 3 Switch and uses the default value Switch.

8. Click **Apply** to save the changes to the device's running-config file.
9. Select the **Save** link at the bottom of the dialog. Select **Yes** when prompted to save the configuration change to the startup-config file on the device's flash memory.
10. Click on the plus sign next to **Command** in the tree view to list the command options.
11. Select the **Reload** link and select **Yes** when the Web management interface asks you whether you really want to reload the software.
12. If the other end of the trunk group is a Layer 2 Switch or Layer 3 Switch, log in to the other device and follow the steps above.

NOTE: Turbolron/4 Layer 2 and Layer 3 Switches support a maximum of six ports—four standard plus two expansion ports. Therefore, Turbolron/4 Layer 2 and Layer 3 Switches support a maximum of three trunk groups of two ports each. The possible trunk groups are ports 1-2, 3-4 and 5-6.

NOTE: Foundry Networks recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

Configuring a Trunk Group of 10-Gigabit Ethernet Ports

Software release 07.6.01 enables you to configure 10 Gigabit Ethernet ports together in a trunk group (aggregate link).

To configure a trunk group containing two 10 Gigabit Ethernet ports, enter commands such as the following:

```
BigIron(config)# trunk ethernet 1/1 to 2/1
BigIron(config-trunk-1/1-2/1)# write memory
BigIron(config-trunk-1/1-2/1)# exit
BigIron(config)# trunk deploy
```

These commands configure a trunk group consisting of 10 Gigabit Ethernet ports 1/1 and 2/1, then deploy the trunk group. The trunk configuration does not take effect until you deploy it.

Syntax: [no] trunk [server | switch] ethernet <primary-portnum> to <secondary-portnum>

Syntax: trunk deploy

The **server** | **switch** parameter specifies whether the trunk ports will be connected to a server or to another Layer 2 Switch or Layer 3 Switch. This parameter affects the type of load balancing performed by the Foundry device. See "Trunk Group Load Sharing" on page 17-15. The default is **switch**. Note that this parameter is not supported on the FES X-Series.

The <primary-portnum> parameter specifies the trunk group's primary port. Except on the FES and FES X-Series, you must specify an odd-numbered slot. See the table "Trunk Group Support" on page 17-3 for valid primary ports.

The <secondary-portnum> parameter specifies the secondary port in the trunk group. You must specify a port that is in the next slot number up from the primary port. For example, if the primary port is 1/1, specify 2/1 as the secondary port.

NOTE: Two-port trunk groups are supported for 10 Gigabit Ethernet. You cannot specify more than two ports.

To display configuration information and load-sharing statistics for the trunk group, enter the **show trunk** command. See "Displaying Trunk Group Configuration Information" on page 17-34.

Additional Trunking Options

The CLI contains commands for doing the following:

- Naming a trunk port
- Disabling or re-enabling a trunk port

- Modifying trunk group membership (ServerIron only)
- Deleting a trunk group

NOTE: To monitor the traffic on a trunk port, see “Monitoring an Individual Trunk Port” on page 9-59.

Naming a Trunk Port

To name an individual port in a trunk group, enter a command such as the following at the trunk group configuration level:

```
BigIron(config-trunk-4/1-4/4)# port-name customer1 ethernet 4/2
```

Syntax: [no] port-name <text> ethernet | pos <portnum>

The <text> parameter specifies the port name. The name can be up to 50 characters long.

This command assigns the name “customer1” to port 4/2 in the trunk group consisting of ports 4/1 – 4/4.

Disabling or Re-Enabling a Trunk Port

You can disable or re-enable individual ports in a trunk group. To disable an individual port in a trunk group, enter commands such as the following at the trunk group configuration level:

```
BigIron(config-trunk-4/1-4/4)# config-trunk-ind
BigIron(config-trunk-4/1-4/4)# disable ethernet 4/2
```

Syntax: [no] config-trunk-ind

Syntax: [no] disable ethernet | pos <portnum>

The **config-trunk-ind** command enables configuration of individual ports in the trunk group. If you do not use this command, the **disable** command will be valid only for the primary port in the trunk group and will disable all ports in the trunk group. You need to enter the **config-trunk-ind** command only once in a trunk group. After you enter the command, all applicable port configuration commands apply to individual ports only.

NOTE: If you enter **no config-trunk-ind**, all port configuration commands are removed from the individual ports and the configuration of the primary port is applied to all the ports. Also, once you enter the **no config-trunk-ind** command, the **enable**, **disable**, and **monitor** commands are valid only on the primary port and apply to the entire trunk group.

The **disable** command disables the port. The states of other ports in the trunk group are not affected.

If you have configured a name for the trunk port, you can specify the port name, as shown in the following example:

```
BigIron(config-trunk-4/1-4/4)# config-trunk-ind
BigIron(config-trunk-4/1-4/4)# disable customer1
```

Syntax: disable <portname>

To enable an individual port in a trunk group, enter commands such as the following at the trunk group configuration level:

```
BigIron(config-trunk-4/1-4/4)# config-trunk-ind
BigIron(config-trunk-4/1-4/4)# enable ethernet 4/2
```

Syntax: enable ethernet | pos <portnum>

Syntax: enable <portname>

Disabling or Re-Enabling a Range or List of Trunk Ports

To disable a range of ports in a trunk group, enter commands such as the following:

```
BigIron(config)# trunk switch ethernet 2/1 to 2/8
BigIron(config-trunk-2/1-2/8)# config-trunk-ind
BigIron(config-trunk-2/1-2/8)# disable ethernet 2/2 to 2/5
```

This command disables ports 2/2 – 2/5 in trunk group 2/1 – 2/8.

To disable a list of ports, enter a command such as the following:

```
BigIron(config-trunk-2/1-2/8)# disable ethernet 2/2 ethernet 2/4 ethernet 2/7
```

This command disables ports 2/2, 2/4, and 2/7 in the trunk group.

You can specify a range and a list on the same command line. For example, to re-enable some trunk ports, enter a command such as the following:

```
BigIron(config-trunk-2/1-2/8)# enable ethernet 2/2 to 2/5 ethernet 2/7
```

Syntax: [no] disable ethernet <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] disable pos <portnum> [to <portnum> | pos <portnum>]

Syntax: [no] enable ethernet <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] enable pos <portnum> [to <portnum> | pos <portnum>]

The **to** <portnum> parameter indicates that you are specifying a range. Specify the lower port number in the range first, then **to**, then the higher port number in the range.

The **ethernet** <portnum> or **pos** <portnum> parameter specifies an individual port. You can enter this parameter multiple times to specify a list, as shown in the examples above. You cannot specify Ethernet and POS ports on the same command line.

Modifying Trunk Group Membership

You can change port membership by removing individual ports from the trunk group. To remove a port from a trunk group, use one of the following methods.

NOTE: Removing individual trunk ports is supported only on the ServerIron.

USING THE CLI

To remove ports 1/3 and 1/4 from the trunk group, enter the following command:

```
BigIron(config)# no trunk ethernet 1/3 to 1/4
```

Syntax: no trunk ethernet | pos <portnum> [to <portnum>]

The <portnum> parameter indicates the port you are removing.

NOTE: Make sure you enter the lower port in the range before the “to” and the higher port in the range after the “to”.

As a shortcut, you also can enter just the lower port in the range. The software automatically removes all higher ports in addition to the specified port. For example, to remove ports 1/3 and 1/4, you can enter the following command:

```
BigIron(config)# no trunk ethernet 1/3
```

The rules regarding trunk group membership are the same as in earlier software releases.

Therefore, for trunk group 1/1 – 1/4, the following commands are not valid:

```
BigIron(config)# no trunk ethernet 1/2
```

Or

```
BigIron(config)# no trunk ethernet 1/2 to 1/4
```

These commands are invalid because the trunk group cannot contain only a single port. These commands, if the software allowed them, would result in a trunk group consisting only of port 1/1.

On most devices, trunk groups can contain two ports or four ports but cannot contain only three ports. Therefore, the following command also is invalid for trunk group 1/1 – 1/4:

```
BigIron(config)# no trunk ethernet 1/4
```

This command is invalid because it would result in a trunk group containing three ports, 1/1 – 1/3.

USING THE WEB MANAGEMENT INTERFACE

1. Disconnect the ports to the server, Layer 2 Switch, or Layer 3 Switch at the other end of the trunk.
2. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
3. Click on the plus sign next to Configure in the tree view to display the configuration options.
4. Select the Trunk link to display a table listing the configured trunk groups.
5. Click the Modify button next to the trunk group you want to modify. The Trunk configuration panel is displayed. The panel contains the settings for the trunk group you selected.
6. Select 2 or 4 to indicate the number of ports.
7. Select Server if you are connecting the trunk group ports to a server. Otherwise, the software assumes you are connecting the trunk group ports to another Layer 2 Switch or Layer 3 Switch and uses the default value Switch.
8. Click the Modify button.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
10. Click on the plus sign next to Command in the tree view to list the command options.
11. Select the Reload link and select Yes when the Web management interface asks you whether you really want to reload the software.

NOTE: Foundry Networks recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

NOTE: If you accidentally select a different port range by selecting a value in the Trunk Group field's pulldown menu, the software creates a new trunk group with the range and other values you select.

Deleting a Trunk Group

To delete a trunk group, use either of the following methods.

USING THE CLI

To delete a trunk group, use “**no**” in front of the command you used to create the trunk group. For example, to remove one of the trunk groups configured in the examples above, enter the following command:

```
BigIron(config)# no trunk ethernet 1/1 to 1/2 ethernet 3/3 to 3/4
```

Syntax: no trunk ethernet | pos <portnum> to <portnum>

USING THE WEB MANAGEMENT INTERFACE

To delete a trunk group:

1. Disconnect the ports to the server, Layer 2 Switch, or Layer 3 Switch at the other end of the trunk.
2. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
3. Click on the plus sign next to Configure in the tree view to display the configuration options.
4. Select the Trunk link to display a table listing the configured trunk groups.
5. Click the Delete button next to the trunk group you want to delete.
6. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

- Click on the plus sign next to Command in the tree view to list the command options.
- Select the **Reload** link and select Yes when the Web management interface asks you whether you really want to reload the software.

NOTE: If the other end of the trunk group is a Layer 2 Switch or Layer 3 Switch, log in to the other system and follow the applicable steps above.

Server Trunk Group Load Sharing Enhancements and Options (Release 07.7.00 and Higher)

Software release 07.7.00 introduced the following server trunk load balancing enhancements and options on Foundry's JetCore devices:

- Load balancing IP traffic based on TCP and UDP application ports
- Ability to configure load balancing per packet
- Ability to configure the maximum number of hash buckets per server trunk

These new features improve the overall performance of server trunk load balancing for Layer 2 and Layer 3 IP traffic.

NOTE: These enhancements apply to server trunks only.

NOTE: These enhancements apply to IP traffic only.

Server Trunk Load Balancing Based on Application Ports

In software release 07.7.00 and later, Foundry's JetCore devices load balance IP traffic on server trunks based on source and destination TCP and UDP application ports (Layer 4 information), as well as on source and destination IP addresses (Layer 3 information). Previous releases supported server trunk load balancing based on Layer 3 information only. Adding Layer 4 information to the load balancing scheme enables the Foundry device to efficiently forward traffic to server trunk group ports.

You do not need to perform any configuration steps to enable this support. If you configure the device for server trunking, it automatically load balances traffic over server trunks based on Layer 3 and Layer 4 information.

For information about configuring server trunks, see the "Configuring Trunk Groups and Dynamic Link Aggregation" chapter in the *Foundry Enterprise Configuration and Management Guide*.

Note Regarding Fragmented Packets

The descriptions above apply to non-fragmented packets. For fragmented packets, the Foundry device uses the source and destination IP addresses only, so that all fragments in a session are forwarded over the same port.

Configuring Server Trunk Load Balancing Per Packet

NOTE: The configuration commands in this section apply to incoming (trunk) ports only.

Starting in release 07.7.00, you can configure the ports on a Foundry device to load balance IP traffic based on individual packets received on the interface. When you enable this feature, the device uses the IP packet headers to load balance the traffic among all the ports in the trunk group.

You configure this feature at the Interface level of the CLI, and not globally (on the entire device). When you configure this feature on the primary port of the trunk group, the software automatically applies it to the other ports in the trunk group.

To enable this feature, enter commands such as the following:

```
BigIron(config)# interface e 1/1
BigIron(config-if-e10000-1/1)# serv-trunk-per-pkt-lb
```

When interface e 1 receives IP packets destined for a trunk port, it uses information in the IP packet header to select the trunk port on which to forward the traffic.

Syntax: [no] serv-trunk-per-pkt-lb

Configuring the Maximum Number of Hash Buckets for Server Trunks

NOTE: This section applies to Foundry Layer 3 Switches only.

Server trunks use hash buckets to implement packet forwarding and load balancing. The hash buckets enable forwarding of packets in hardware, as opposed to forwarding them in software (sending them to the CPU). Packets forwarded in hardware travel faster in comparison to packets sent to the CPU for processing.

When the Foundry device learns that a specific packet has to go through an outgoing port, it places an entry in the hash bucket. The entry defines the data path from the incoming port to the outgoing port. When the device receives subsequent packets destined for the same path, it retrieves the entry in the hash bucket and forwards the packets accordingly.

In releases prior to 07.7.00, the Foundry device allocates a fixed number of hash buckets for each server trunk. This number is not configurable.

In software release 07.7.00, depending on the number of server trunks configured on the Foundry device, you can specify the maximum number of hash buckets per server trunk, up to a maximum of 256. In addition, the total number of hash buckets for all server trunks combined has increased. The BigIron 15000 supports a total of 1024 hash buckets, and the BigIron 4000 and BigIron 8000 support a total of 8192 hash buckets.

Increasing the number of hash buckets per server trunk enhances the speed and efficiency at which the Foundry device forwards and load balances IP packets on server trunk ports.

Table 18.5 shows the hash bucket configurations supported on the BigIron 15000.

Table 18.5: Configurable Hash Buckets on the BigIron 15000

Maximum Number of Hash Buckets per Server Trunk	Number of Server Trunks	Total Number of Hash Buckets for all Server Trunks Combined
16 – default value	15	240
32	15	480
64	15	960
128	8	1024
256	4	1024

Table 18.6 shows the hash bucket configurations supported on the BigIron 4000 and BigIron 8000.

Table 18.6: Configurable Hash Buckets on the BigIron 4000 and BigIron 8000

Maximum Number of Hash Buckets per Server Trunk	Number of Server Trunks	Total Number of Hash Buckets for all Server Trunks Combined
16	32	512
32	32	1024
64	32	2048
128	32	4096
256 – default value	32	8192

To configure the maximum number of hash buckets per server trunk, enter commands such as the following:

```
BigIron(config)# system hash-per-server-trunk 64
BigIron(config)# write mem
BigIron(config)# end
BigIron # reload
```

NOTE: You must reload the software to place this configuration in effect.

Syntax: [no] system hash-per-server-trunk <maximum number of hash buckets>

where **maximum number of hash buckets** can be 32, 64, 128, or 256. On the BigIron 15000, the default is 16. On the BigIron 4000 and BigIron 8000, the default is 256.

Enabling Optimized Server Trunk Load Balancing (VM1 only)

You can optimize individual ports for server trunk load balancing. An optimized port load balances based on source and destination IP address but uses a smaller session table, which enables the port to more quickly forward traffic received on the port to the server trunk group ports.

NOTE: This enhancement applies only to the VM1.

NOTE: This enhancement applies to server trunk groups only, not to switch trunk groups.

NOTE: On Foundry devices that support IPv6, optimized trunk ports can load balance traffic based on source and destination IPv6 addresses.

Without optimization, the device performs the following types of load balancing for IP traffic.

Layer 2

The load balancing occurs at Layer 2 if the traffic is being forwarded in hardware. IP traffic on a server trunk group is load balanced as follows:

- On a Layer 2 Switch:
 - IP traffic received on a 10/100 port is load balanced based on source and destination IP address.
 - IP traffic received on a Gigabit port is load balanced based on the port number of the Gigabit port that received the traffic.
- On a Layer 3 Switch:
 - IP traffic received on a 10/100 port or Gigabit port is load balanced based on destination IP address.

Layer 3

If any of the following features are enabled on a port, load balancing occurs in software using the entries in the session table. In this case, the IP traffic is load balanced based on source and destination IP address.

- ACLs
- Rate limiting (Fixed Rate Limiting or Adaptive Rate Limiting)
- NetFlow
- sFlow Export
- Network Address Translation (NAT)
- Policy-Based Routing (PBR)

If you do not have any of these features enabled on the port but you still want to load balance the traffic based on source and destination IP address, you can do so by enabling the server trunk load balancing optimization feature. Even if you do have one of the features above configured on the port, you can enhance load balancing

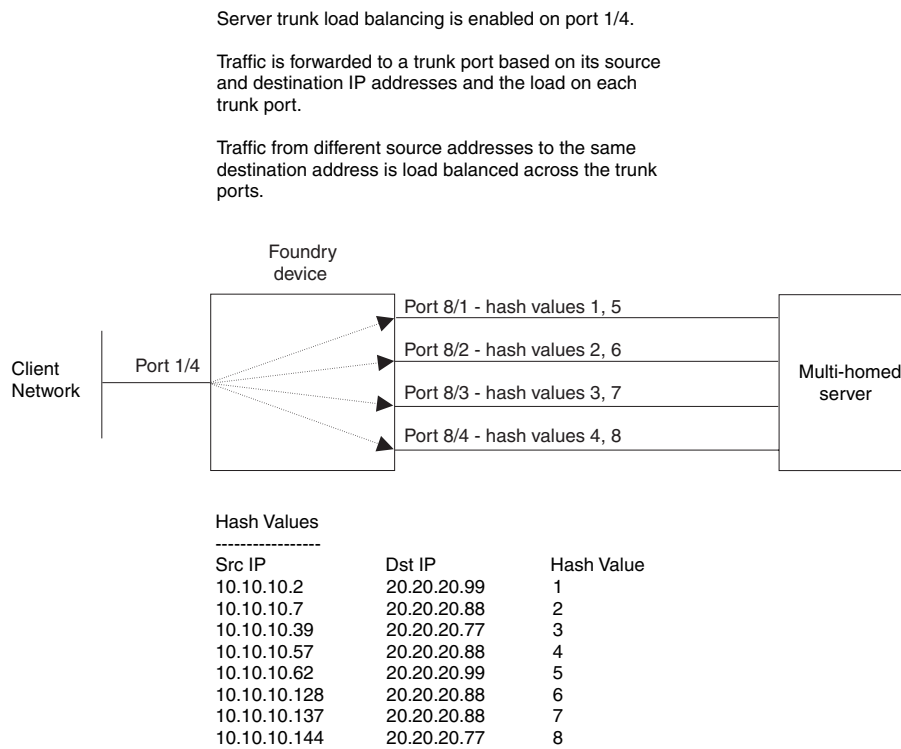
performance by enabling the optimization feature. The optimization feature uses a smaller session table, which allows forwarding to occur more quickly.

NOTE: When you enable the server trunk load balancing optimization feature on a port, the feature listed above are disabled on that port. This occurs because the features use the session table, but the optimization feature uses a smaller session table than the other features. The configuration information for the other features is retained in the device's configuration file, but the features are disabled.

Example of Server Trunk Load Balancing at Layer 3

Figure 18.2 shows an example of how IP traffic is load balanced to server trunk ports when the traffic is forwarded at Layer 3. In this example, server trunk load balancing based on source and destination IP addresses is enabled on a Gigabit Ethernet port connected to a network containing multiple clients. Four other Ethernet ports are configured in a server trunk group that is connected to a multi-homed server. The server can have multiple network adapters or a single adapter with multiple ports that have unique MAC and IP addresses.

Figure 18.2 Server trunk load balancing based on source and destination IP addresses



When the port connected to the client network receives traffic that needs to be forwarded to the server, the Foundry device selects one of the ports in the trunk group, and forwards the traffic on the selected port.

The Foundry device selects the trunk port based on a hash value, which can be a number from 1 – 256. The Foundry device calculates a hash value for traffic that enters the device through the server trunk load balancing port and exits the device through a trunk group. The hash value is calculated based on the source and destination IP addresses in the traffic.

After the Foundry device calculates the hash value for the traffic, the device examines the trunk ports connected to the destination address and selects the port with the fewest hash values already assigned. After calculating a hash value and assigning the value to a port, the device always uses the same port to forward traffic for the same source and destination IP addresses.

For example, the first time the Foundry device receives traffic from 10.10.10.7 addressed to 20.20.20.88, the device calculates the hash value 2 for the traffic. The device then checks the trunk ports to see whether a port is assigned to hash value 2.

- If a trunk port is assigned to hash value 2, the device uses that port to forward the traffic.
- Otherwise, the device assigns hash value 2 to the trunk port with the fewest hash values already assigned to it. The device continues to use this port for traffic with hash value 2, until a state change occurs on a trunk port or a trunk port is added or removed.

Trunk ports keep the hash values that are assigned to them until a trunk port's state changes or a trunk port is added or removed. When any of these changes occurs, the Foundry device clears the hash values from all of the trunk ports and begins calculating and assigning hash values again for new traffic.

Configuration Considerations

- You can enable the server trunk load balancing optimization feature on an individual port basis only. You cannot enable the feature on a virtual routing interface basis. This is true even if you have assigned a virtual routing interface to the trunk ports.
- Each VSP CPU has a separate hash bucket for the ports managed by the CPU. The buckets are independent of one another. Thus, if you enable the feature on more than one port and the ports are not managed by the same CPU, it is possible for the same hash values to be assigned to more than one trunk port, because the values are assigned separately by each CPU.
- When you enable the server trunk load balancing optimization feature on a port, the following features are disabled on the port:
 - ACLs
 - Rate limiting (Fixed Rate Limiting or Adaptive Rate Limiting)
 - NetFlow
 - sFlow Export
 - Network Address Translation (NAT)
 - Policy-Based Routing (PBR)

The features are disabled because the server trunk load balancing optimization feature uses a simpler session table whose forwarding entries are keyed by source and destination IP addresses only. The features listed above require use of the standard session table, which also includes keys for the IP protocol and the source and destination TCP or UDP application ports (when the IP protocol is TCP or UDP).

The configuration information for these features remains in the device's configuration file but the features are disabled on the port.

Enabling Server Trunk Load Balancing Optimization

To enable server trunk load balancing optimization, you enable the feature on the ports that will receive the traffic that needs to be load balanced. To enable the optimization feature on a port, enter the following command at the configuration level for the port:

```
BigIron(config-if-e1000-1/4)# stlb
```

Syntax: [no] stlb

Displaying Server Trunk Load Balancing Information

To display the current hash assignments for server trunk ports, log on to the VSP CPU that is managing the ports, then enter the **show trunk** command. Here is an example.

```
BigIron# rconsole 2 1
BigIron2/1 # show trunk
BigIron2/1 #Number of trunk groups: 1
Note: Value in ( ) is for server trunk hashing.

TRUNK ID: 71      server:1      multi-slot:0
           configured ports: 8/1  8/2  8/3  8/4
           active ports   : 8/1  (2) 8/2  (2) 8/3  (2) 8/4  (1)
BigIron2/1 # rconsole-exit
```

The **rconsole 2 1** command logs on to VSP CPU 1 on the VM1 module in slot 2.

The **show trunk** command displays the trunk information for the ports managed by the CPU. The server trunk load balancing information is shown in bold type in this example. The number in parentheses indicates how many hash values are assigned to the port. The CPU assigns the hash values evenly to the trunk ports managed by the CPU. In this example, the next time the device needs to assign a hash value, the device will assign the value to port 8/4.

The **rconsole-exit** command logs out of the VSP CPU.

Syntax: show trunk

For information about the VM1, including how the module distributes management of the ports in the chassis, see "Using the Velocity Management Module" on page 4-1.

Displaying Trunk Group Configuration Information

To display configuration information for the trunk groups configured on the Layer 3 Switch, use one of the following methods. Each method displays information for configured trunk groups and operational trunk groups. A configured trunk group is one that has been configured in the software but has not been placed into operation by a reset or reboot. An operational trunk group is one that has been placed into operation by a reset or reboot.

USING THE CLI

Enter the following command at any CLI level:

```
BigIron(config)# show trunk
Configured trunks:
Trunk Type  Ports
   1  Switch 1/1 1/2 1/3 1/4 2/1 2/2 2/3 2/4
Operational trunks:
Trunk Type  Ports                                Duplex Speed Tag Priority
   1  Switch 1/1 1/2 1/3 1/4 2/1 2/2 2/3 2/4  None  None  No  level0
```

Syntax: show trunk [ethernet | pos <portnum> to <portnum>]

The following table describes the information displayed by the **show trunk** command.

Table 18.7: CLI Trunk Group Information

This Field...	Displays...
Trunk	The trunk group number. The software numbers the groups in the display to make the display easy to use.

Table 18.7: CLI Trunk Group Information (Continued)

This Field...	Displays...
Type	The type of trunk group, which can be one of the following: <ul style="list-style-type: none"> • Server – The trunk group is connected to a server. • Switch – The trunk group is connected to another Layer 2 Switch or Layer 3 Switch.
Ports	The ports in the trunk group.
Duplex	The mode of the port, which can be one of the following: <ul style="list-style-type: none"> • None – The link on the primary trunk port is down. • Full – The primary port is running in full-duplex. • Half – The primary port is running in half-duplex. <p>Note: This field and the following fields apply only to operational trunk groups.</p>
Speed	The speed set for the port. The value can be one of the following: <ul style="list-style-type: none"> • None – The link on the primary trunk port is down. • 10 – The port speed is 10 Mbps. • 100 – The port speed is 100 Mbps. • 1G – The port speed is 1000 Mbps.
Tag	Indicates whether the ports have 802.1q VLAN tagging. The value can be Yes or No.
Priority	Indicates the Quality of Service (QoS) priority of the ports. The priority can be a value from 0 – 7.

To display trunk group information for specific ports, enter a command such as the following:

```
BigIron(config)# show trunk ethernet 1/1 to 1/8

Configured trunks:

Trunk ID: 1
Type: Switch
Ports_Configured: 8
Primary Port Monitored: Jointly

Ports      1/1      1/2      1/3      1/4      1/5      1/6      1/7      1/8
Port Names none     none     none     none     none     longna   test     none
Port_Status enable   enable   enable   enable   disable  disable  enable   enable
Monitor     on       on       off      on       off      off      off      off
Mirror Port 3/3     3/4     N/A     3/5     N/A     N/A     N/A     N/A
Monitor Dir both     in       N/A     out     N/A     N/A     N/A     N/A

Operational trunks:

Trunk ID: 1
Type: Switch
Duplex: Full
Speed: 1G
Tag: No
Priority: level0
Active Ports: 6

Ports      1/1      1/2      1/3      1/4      1/5      1/6      1/7      1/8
Link_Status active   active   active   active   down     down     active   active
LACP_Status ready    ready    ready    expired  down     down     ready    ready
Load Sharing
Mac Address 3        2        2        2        0        0        6        1
IP          0        0        0        0        0        0        0        0
IPX        0        2        1        0        0        0        0        1
Apple Talk 1        2        0        4        0        0        0        3
```

The display is divided into sections for configured trunks and operational trunks. A configured trunk group is one that has not been activated yet.

Table 18.8 describes the information displayed by the **show trunk** command.

Table 18.8: CLI Trunk Group Information

This Field...	Displays...
Trunk ID	The trunk group number. The software numbers the groups in the display to make the display easy to use.
Type	The type of trunk group, which can be one of the following: <ul style="list-style-type: none"> Server – The trunk group is connected to a server. Switch – The trunk group is connected to another Layer 2 Switch or Layer 3 Switch.

Table 18.8: CLI Trunk Group Information (Continued)

This Field...	Displays...
Duplex	<p>The mode of the port, which can be one of the following:</p> <ul style="list-style-type: none"> None – The link on the primary trunk port is down. Full – The primary port is running in full-duplex. Half – The primary port is running in half-duplex. <p>Note: This field and the following fields apply only to operational trunk groups.</p>
Speed	<p>The speed set for the port. The value can be one of the following:</p> <ul style="list-style-type: none"> None – The link on the primary trunk port is down. 10 – The port speed is 10 Mbps. 100 – The port speed is 100 Mbps. 1G – The port speed is 1000 Mbps.
Tag	Indicates whether the ports have 802.1q VLAN tagging. The value can be Yes or No.
Priority	Indicates the Quality of Service (QoS) priority of the ports. The priority can be a value from 0 – 7.
Active Ports	The number of ports in the trunk group that are currently active.
Ports	The ports in the trunk group.
Link_Status	The link status or each port in the trunk group.
LACP_Status	<p>This field appears in software releases 07.6.03 and later. For more information about this feature, see the section “Displaying and Determining the Status of Aggregate Links” on page 17-50.</p> <ul style="list-style-type: none"> Ready - The port is functioning normally in the trunk group and is able to transmit and receive LACP packets. Expired - The time has expired (as determined by timeout values) and the port has shut down because the port on the other side of the link has stopped transmitting packets. Down - The port’s physical link is down.
Load Sharing	The number of traffic flows currently being load balanced on the trunk ports. All traffic exchanged within the flow is forwarded on the same trunk port. For information about trunk load sharing, see “Trunk Group Load Sharing” on page 17-15.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the Trunk link to display a table listing the configured trunk groups.

This display shows the following information.

Table 18.9: Web Management Trunk Group Information

This Field...	Displays...
Connection Type	The type of trunk group, which can be one of the following: <ul style="list-style-type: none"> • Server – The trunk group is connected to a server. • Switch – The trunk group is connected to another Layer 2 Switch or Layer 3 Switch.
Port Members	The ports in the trunk group.

Dynamic Link Aggregation

The software supports the IEEE 802.3ad standard for link aggregation. This standard describes the Link Aggregation Control Protocol (LACP), a mechanism for allowing ports on both sides of a redundant link to configure themselves into a trunk link (aggregate link), without the need for manual configuration of the ports into trunk groups.

When you enable link aggregation on a group of Foundry ports, the Foundry ports can negotiate with the ports at the remote ends of the links to establish trunk groups.

Usage Notes

- You cannot use 802.3ad link aggregation on a port configured as a member of a static trunk group.
- This feature is supported only for 10/100 and Gigabit Ethernet ports.
- When the feature dynamically adds or changes a trunk group, the **show trunk** command displays the trunk as both configured and active. However, the **show running-config** or **write terminal** command does not contain a trunk command defining the new or changed trunk group.
- If the feature places a port into a trunk group as a secondary port, all configuration information except information related to link aggregation is removed from the port. For example, if port 1/3 has an IP interface, and the link aggregation feature places port 1/3 into a trunk group consisting of ports 1/1 – 1/4, the IP interface is removed from the port.
- If you use this feature on a Layer 3 Switch that is running OSPF or BGP4, the feature causes these protocols to reset when a dynamic link change occurs. The reset includes ending and restarting neighbor sessions with OSPF and BGP4 peers, and clearing and relearning dynamic route entries and forwarding cache entries. Although the reset causes a brief interruption, the protocols automatically resume normal operation.
- You can enable link aggregation on 802.1q tagged ports (ports that belong to more than one port-based VLAN) in software release 07.7.00 and later, and in FES software releases 03.4.00 and later.
- Dynamic Operation of Allocation Keys (section 43.6.2 in the 802.3ad specification) is supported in release 07.7.00 and later. FES devices do not support dynamically allocated keys for tagged ports. These devices support statically assigned keys only.

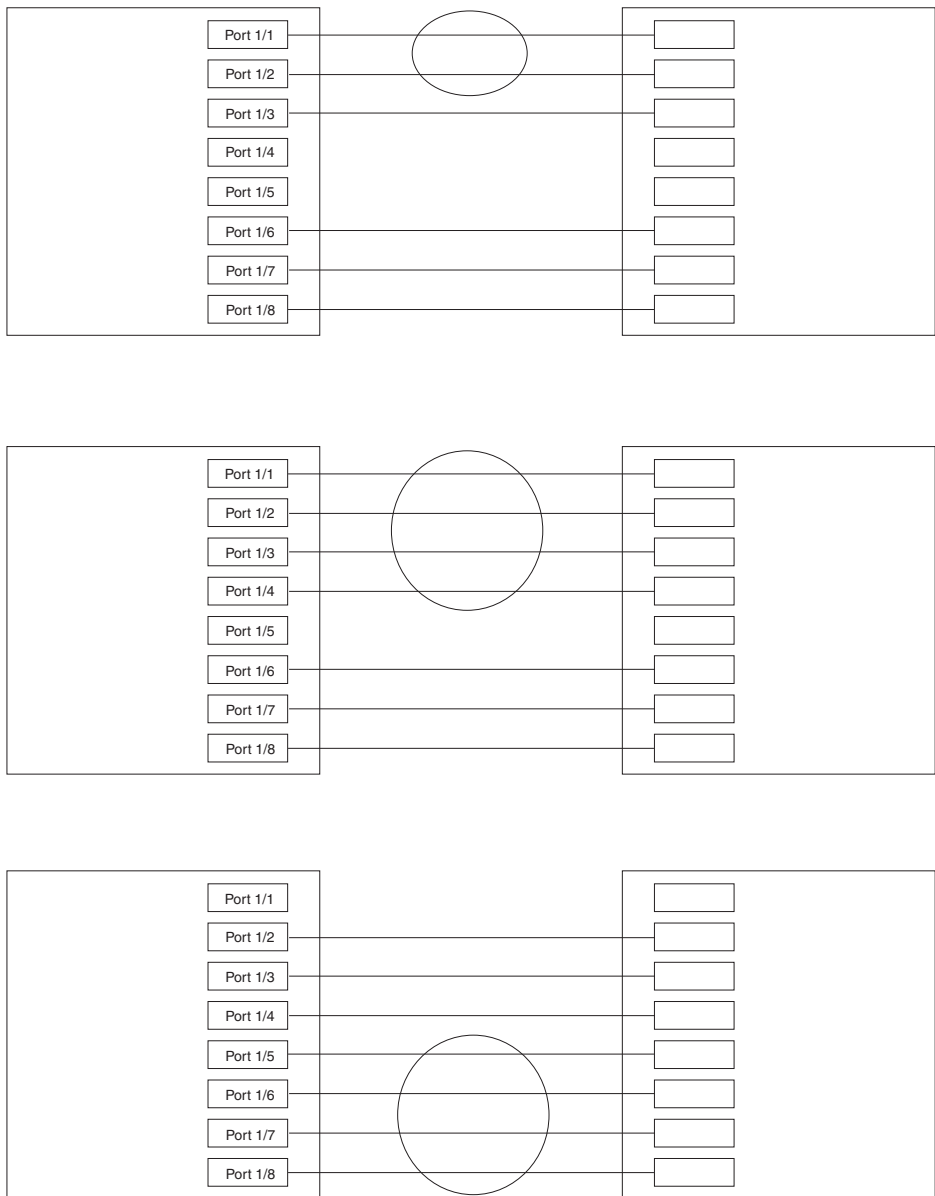
Configuration Rules

Foundry ports follow the same configuration rules for dynamically created aggregate links as they do for statically configured trunk groups. See “Trunk Group Rules” on page 17-3 and “Trunk Group Load Sharing” on page 17-15.

Figure 18.3 on page 17-39 shows some examples of valid aggregate links.

Figure 18.3 Examples of valid aggregate links

Foundry ports enabled for link aggregation follow the same rules as ports configured for trunk groups.



In this example, assume that link aggregation is enabled on all of the links between the Foundry device on the left and the device on the right (which can be either a Foundry device or another vendor's device). Notice that some ports are not able to join an aggregate link even though link aggregation is enabled on them. The ports that are not members of aggregate links in this example are not following the configuration rules for trunk links on Foundry devices.

The Foundry rules apply to a Foundry device even if the device at the other end is from another vendor and uses different rules. See "Trunk Group Rules" on page 17-3.

The link aggregation feature automates trunk configuration but can coexist with Foundry's trunk group feature. Link aggregation parameters do not interfere with trunk group parameters.

NOTE: Use the link aggregation feature only if the device at the other end of the links you want to aggregate also supports IEEE 802.3ad link aggregation. Otherwise, you need to manually configure the trunk links.

Link aggregation support is disabled by default. You can enable the feature on an individual port basis, in active or passive mode.

- Active mode – When you enable a port for active link aggregation, the Foundry port can exchange standard LACP Protocol Data Unit (LACPDU) messages to negotiate trunk group configuration with the port on the other side of the link. In addition, the Foundry port actively sends LACPDU messages on the link to search for a link aggregation partner at the other end of the link, and can initiate an LACPDU exchange to negotiate link aggregation parameters with an appropriately configured remote port.
- Passive mode – When you enable a port for passive link aggregation, the Foundry port can exchange LACPDU messages with the port at the remote end of the link, but the Foundry port cannot search for a link aggregation partner or initiate negotiation of an aggregate link. Thus, the port at the remote end of the link must initiate the LACPDU exchange.

NOTE: Foundry recommends that you disable or remove the cables from the ports you plan to enable for dynamic link aggregation. Doing so prevents the possibility that LACP will use a partial configuration to talk to the other side of a link. A partial configuration does not cause errors, but does sometimes require LACP to be disabled and re-enabled on both sides of the link to ensure that a full configuration is used. It's easier to disable a port or remove its cable first. This applies both for active link aggregation and passive link aggregation.

802.3ad Enhancements in Release 07.6.01

Software release 07.6.01 contains the following enhancements to 802.3ad support:

- Adaptation to trunk disappearance. The Foundry device will tear down an aggregate link if the device at the other end of the link reboots or brings all the links down. Tearing the aggregate link down prevents a mismatch if the other device has a different trunk configuration following the reboot or re-establishment of the links.
- The criteria for being eligible to be in an aggregate link are more flexible. A range of ports can contain down ports and still be eligible to become an aggregate link.


Adaptation to Trunk Disappearance

Release 07.6.01 prevents trunk mismatches caused when one device changes the number of ports in group of ports that has become part of an 802.3 aggregate link. In 07.6.01 and later, if a device changes the number of ports in an active aggregate link, the Foundry device on the other end of the link tears down the link. Once the other device recovers, 802.3 can renegotiate the link without a mismatch.

In previous releases, it is possible for a trunk mismatch to occur between two devices that have established an aggregate link. This can occur if one of the devices reboots or brings the trunk links down, then re-establishes the links but with a different number of trunk ports. Figure 18.4 shows an example.

Figure 18.4 Trunk port mismatch


Four ports on each device are eligible for link aggregation. The device negotiates a four-port trunk using the ports.



One device reloads, after which only two of its ports are eligible for link aggregation.

However, the first device is still configured with the four-port trunk group. The trunks are mismatched.

This type of mismatch does not occur in release 06.7.01 and later.



Flexible Trunk Eligibility

Software release 07.6.01 also increases the tolerance for down ports during link negotiation. In previous releases, all the ports in a valid trunk configuration (2-port, 4-port, or 8-port trunk starting on a valid primary port number) need to be up. Thus, in previous releases, if you enable link aggregation on four ports but one of the ports is down, the device will negotiate based only on a valid two-port trunk group consisting of two of the up ports. For example, if you enable link aggregation on ports 1/1 - 1/4 and port 1/3 is down, 802.3ad will negotiate only for a two-port link consisting of ports 1/1 and 1/2.

In release 07.6.01 and later, the device groups the device's ports into 2-port groups consisting of an odd-numbered port and the next even-numbered port. For example, ports 1/1 and 1/2 are a two-port group, as are ports 1/3 and 1/4, 9/1 and 9/10, and do on. If either of the ports in a two-port group is up, the device considers both ports to be eligible to be in an aggregate link.

Figure 18.5 shows an example of 2-port groups in a range of eight ports on which link aggregation is enabled. Based on the states of the ports, some or all of them will be eligible to be used in an aggregate link.

Figure 18.5 Two-port groups used to determine aggregation eligibility

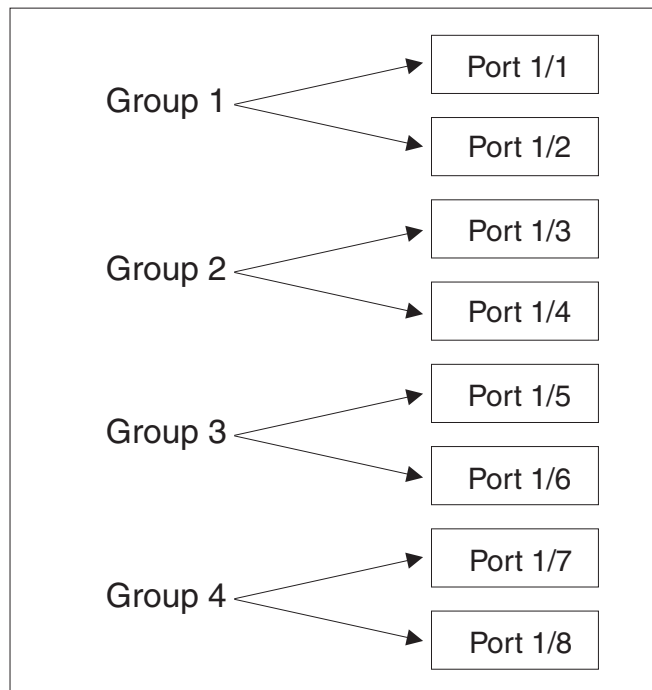


Table 18.10 shows examples of the ports from Figure 18.5 that will be eligible for an aggregate link based on individual port states.

Table 18.10: Port Eligibility for Link Aggregation

	Port Group 1		Port Group 2		Port Group 3		Port Group 4		Trunk Eligibility
	1/1	1/2	1/3	1/4	1/5	1/6	1/7	1/8	
Link State	Up	Up	Up	Up	Up	Up	Up	Up	8-port 1/1 – 1/8
	Up	Up	Up	Up	Up	Down	Up	Up	8-port 1/1 – 1/8
	Up	Up	Up	Up	Up	Down	Up	Down	8-port 1/1 – 1/8
	Up	Up	Up	Up	Down	Down	Down	Up	4-port 1/1 – 1/4
	Down	Down	Down	Up	Up	Up	Up	Up	4-port 1/5 – 1/8
	Up	Down	Down	Down	Up	Down	Down	Down	2-port 1/1 – 1/2

As shown in these examples, all or a subset of the ports within a port range will be eligible for formation into an aggregate link based on port states. Notice that the sets of ports that are eligible for the aggregate link must be valid static trunk configurations. For example, a 4-port link consisting of ports 1/4 – 1/7 is not valid because this port configuration is not valid for static trunk groups on the Foundry device.

Enabling Link Aggregation

By default, link aggregation is disabled on all ports. To enable the feature, use one of the following CLI methods.

USING THE CLI

To enable link aggregation on a set of ports, enter commands such as the following at the interface configuration level of the CLI.

NOTE: Configuration commands for link aggregation differ depending on whether you are using the default link aggregation key automatically assigned by the software, or if you are assigning a different, unique key. Follow the commands below, according to the type of key you are using. For more information about keys, see “Key” on page 17-44.

Using the Default Key Assigned by the Software

```
BigIron(config)# interface ethernet 1/1
BigIron(config-if-e1000-1/1)# link-aggregate active
BigIron(config)# interface ethernet 1/2
BigIron(config-if-e1000-1/2)# link-aggregate active
```

The commands in this example enable the active mode of link aggregation on ports 1/1 and 1/2. The ports can send and receive LACPDU messages. Note that these ports will use the default key, since one has not been explicitly configured.

Assigning a Unique Key

```
BigIron(config)# interface ethernet 1/1
BigIron(config-if-e1000-1/1)# link-aggregate configure key 10000
BigIron(config-if-e1000-1/1)# link-aggregate active
BigIron(config)# interface ethernet 1/2
BigIron(config-if-e1000-1/2)# link-aggregate configure key 10000
BigIron(config-if-e1000-1/2)# link-aggregate active
```

The commands in this example assign the key 10000 and enable the active mode of link aggregation on ports 1/1 and 1/2. The ports can send and receive LACPDU messages.

NOTE: As shown in this example, when configuring a key, it is pertinent that you assign the key prior to enabling link aggregation.

The following commands enable passive link aggregation on ports 1/5 – 1/8:

```
BigIron(config)# interface ethernet 1/5 to 1/8
BigIron(config-mif-1/5-1/8)# link-aggregate passive
```

The commands in this example enable the passive mode of link aggregation on ports 1/5 – 1/8. These ports wait for the other end of the link to contact them. After this occurs, the ports can send and receive LACPDU messages.

To disable link aggregation on a port, enter a command such as the following:

```
BigIron(config-if-e1000-1/8)# link-aggregate off
```

Syntax: [no] link-aggregate active | passive | off

Syntax: [no] link-aggregate configure [system-priority <num>] | [port-priority <num>] | [key <num>] | [type server | switch]

NOTE: For more information about keys, including details about the syntax shown above, see “Key” on page 17-44.

Link Aggregation Parameters

You can change the settings for the following link aggregation parameters, on an individual port basis:

- System priority
- Port priority
- Link type
- Key

System Priority

The system priority specifies the Foundry device's link aggregation priority relative to the devices at the other ends of the links on which link aggregation is enabled. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

NOTE: If you are connecting the Foundry device to another vendor's device and the link aggregation feature is not working, set the system priority on the Foundry device to a lower priority (a higher priority value). In some cases, this change allows the link aggregation feature to operate successfully between the two devices.

Port Priority

The port priority determines the active and standby links. When a group of ports is negotiating with a group of ports on another device to establish a trunk group, the Foundry port with the highest priority becomes the default active port. The other ports (with lower priorities) become standby ports in the trunk group. You can specify a priority from 0 – 65535. A higher value indicates a lower priority. The default is 1.

NOTE: This parameter is not supported in the current software release. The primary port in the port group becomes the default active port. The primary port is the lowest-numbered port in a valid trunk-port group.

Link Type

The link type specifies whether the trunk is connecting to a server (server link) or to another networking device (switch link). The default link type is switch.

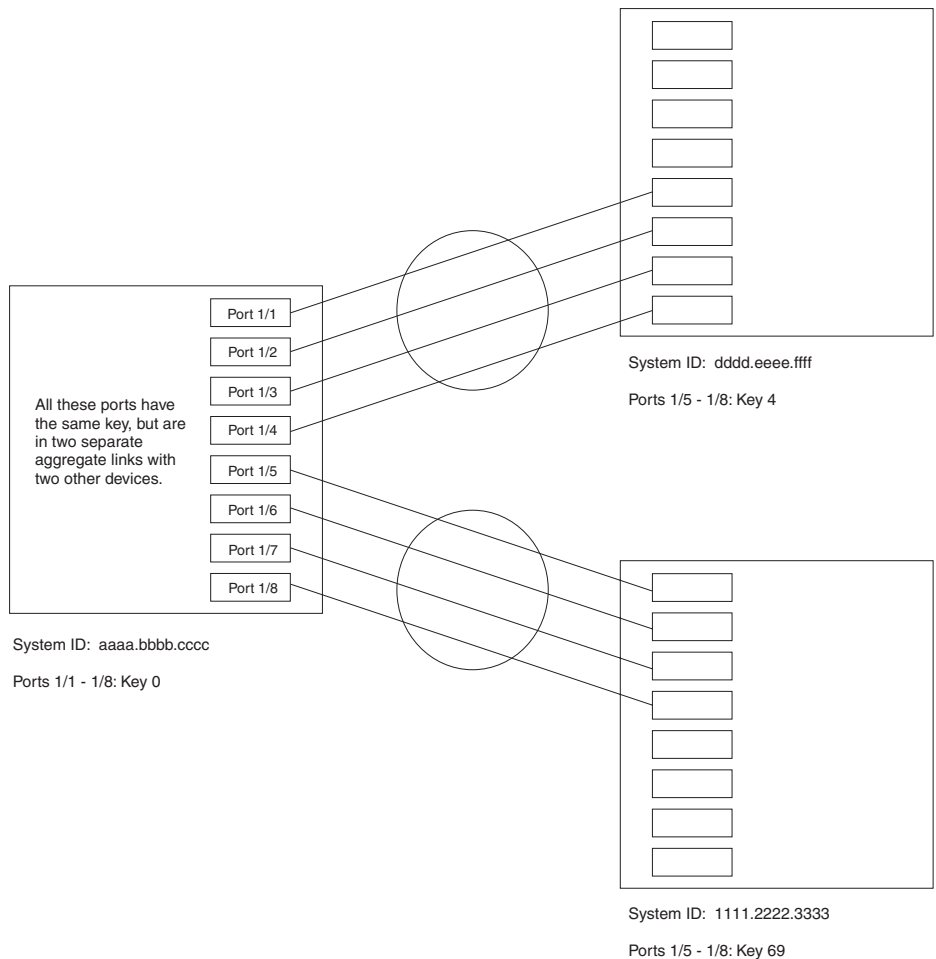
Key

Every port that is 802.3ad-enabled has a key. The key identifies the group of potential trunk ports to which the port belongs. Ports with the same key are called a key group and are eligible to be in the same trunk group.

When you enable link-aggregation on a tagged or untagged port, Foundry's software assigns a default key to the port. The default key is based on the position of the port within an eight-port group (the maximum number of ports in a trunk group on a Layer 3 Switch). The software assigns the keys in ascending numerical order, beginning with key 0 for the first group of eight ports. For example, a 24-port module in chassis slot 1 contains keys 0, 1, and 2 by default. Ports 1/1 – 1/8 have key 0, ports 1/9 – 1/16 have key 1, and so on.

All ports within an aggregate link must have the same key. However, if the device has ports that are connected to two different devices, and the port groups allow the ports to form into separate aggregate links with the two devices, then each group of ports can have the same key while belonging to separate aggregate links with different devices. Figure 18.6 on page 17-45 shows an example.


Figure 18.6 Ports with the same key in different aggregate links



Notice that the keys between one device and another do not need to match. The only requirement for key matching is that all the ports within an aggregate link on a given device must have the same key.

Devices that support multi-slot trunk groups can form multi-slot aggregate links using link aggregation. However, the link aggregation keys for the groups of ports on each module must match. For example, if you want to allow link aggregation to form an aggregate link containing ports 1/1 – 1/4 and 3/5 – 3/8, you must change the link aggregation key on one or both groups of ports so that the key is the same on all eight ports. Figure 18.7 on page 17-46 shows an example.

Figure 18.7 Multi-slot aggregate link



By default, the device's ports are divided into 4-port groups. The software dynamically assigns a unique key to each 4-port group. If you need to divide a 4-port group into two 2-port groups, change the key in one of the groups so that the two 2-port groups have different keys. For example, if you plan to use ports 1/1 and 1/2 in VLAN 1, and ports 1/3 and 1/4 in VLAN 2, change the key for ports 1/3 and 1/4.

NOTE: If you change the key for a port group, Foundry recommends that you use the value 10000 or higher, to avoid potential conflicts with dynamically created keys.

Dynamic Operation of Allocation Keys

Starting with software release 07.7.00, the Foundry device dynamically changes a port's key based on changes to the port's VLAN membership.

When you change a port's VLAN membership, the device searches through existing key groups for a port with matching port properties. Specifically, it searches for a match on all three of the following properties:

- VLAN ID
- default key
- port tag type (tagged or untagged)

If it finds a match, the port (whose VLAN membership you are changing) gets the matching port's key. If it does not find a match, the port gets a new key.

NOTE: For multi-slot trunk groups, you must manually configure the keys in the trunk group(s) to match. For instructions on configuring keys manually, see "Configuring Keys For Ports with Link Aggregation Enabled" on page 17-48.

How Changing a Port's VLAN Membership Affects Trunk Groups and Dynamic Keys

When you change a port's VLAN membership, and the port is currently a member of a trunk group, the following changes occur:

- The Foundry device tears down the existing trunk group.
- All ports in the trunk group get a new key.
- The new key group aggregates into a new trunk group.

When you change a port's VLAN membership, and the port is not a member of a trunk group, the following changes occur:

- The port gets a new key depending on changes to the port's VLAN tag type, as follows:
 - Tagged to Tagged VLAN – The primary port of the trunk group gets a new key.

- Tagged to Untagged VLAN –The port gets the default key for untagged ports.
- Untagged to Tagged VLAN – If the Foundry device finds a port with matching port properties, the port gets that port’s key. If it doesn’t find one, the port gets a new key.
- Untagged to Untagged VLAN – The port gets a new key depending on whether it’s in the default VLAN or not. If there is a trunk group associated with the key, it is not affected.
- All other ports keep their existing key.
- The new key groups try to aggregate into trunk groups.

Viewing Keys for Tagged Ports

To display link aggregation information, including the key for a specific port, enter a command such as the following at any level of the CLI:

```
BigIron# show link-aggregation ethernet 1/1
System ID: 00e0.52a9.bb00
Port [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp]
1/1      0      0      0  No  L  No  No  No  No  No  No
```

The command in this example shows the key and other link aggregation information for port 1/1.

To display link aggregation information, including the key for all ports on which link aggregation is enabled, enter the following command at any level of the CLI:

```
BigIron# sh link-agg
System ID: 0004.8055.b200
Long timeout: 90, default: 90
Short timeout: 3, default: 3
Port [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
1/1      1      1  10000  Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
1/2      1      1  10000  Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
2/1      1      1  10000  Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
2/2      1      1  10000  Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
4/1      1      1    480   Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
4/2      1      1    480   Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
4/3      1      1    480   Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
4/4      1      1    480   Yes  S  Agg  Syn  Col  Dis  Def  No  Dwn
4/17     1      1    481   Yes  S  Agg  Syn  Col  Dis  Def  No  Ope
4/18     1      1    481   Yes  S  Agg  Syn  Col  Dis  Def  No  Ope
4/19     1      1    481   Yes  S  Agg  Syn  Col  Dis  Def  No  Ope
4/20     1      1    481   Yes  S  Agg  Syn  Col  Dis  Def  No  Ope
```

For information about the fields in this display, see Table 18.11 on page 17-52.

Syntax: show link-aggregation [ethernet <portnum>]

Possible values: N/A

Default value: N/A

Configuring Link Aggregation Parameters

You can configure one or more parameters on the same command line, and you can enter the parameters in any order.

NOTE: For key configuration only, configuration commands differ depending on whether or not link aggregation is enabled on the port(s). Follow the appropriate set of commands below, according to your system's configuration.

For example, to change a port group's key from the one assigned by the software to another value, enter commands such as the following:

NOTE: Use this command sequence to change the key for ports that do not have link aggregation enabled, and for all other link aggregation parameters (i.e., system priority, port priority, and link type).

```
BigIron(config)# interface ethernet 1/1 to 1/4
BigIron(config-mif-1/1-1/4)# link-aggregate configure key 10000
BigIron(config-mif-1/1-1/4)# interface ethernet 3/5 to 3/8
BigIron(config-mif-3/5-3/8)# link-aggregate configure key 10000
```

Configuring Keys For Ports with Link Aggregation Enabled

NOTE: As shown in this command sequence, to change the key on ports that already have link aggregation enabled, you must first turn OFF link aggregation, configure the new key, then re-enable link aggregation.

```
BigIron(config)# interface ethernet 1/1 to 1/4
BigIron(config-mif-1/1-1/4)# link-aggregate off
BigIron(config-mif-1/1-1/4)# link-aggregate configure key 10000
BigIron(config-mif-1/1-1/4)# link-aggregate active
BigIron(config-mif-1/1-1/4)# interface ethernet 3/5 to 3/8
BigIron(config-mif-3/5-3/8)# link-aggregate off
BigIron(config-mif-3/5-3/8)# link-aggregate configure key 10000
BigIron(config-mif-3/5-3/8)# link-aggregate active
```

These commands change the key for ports 1/1 – 1/4 and 3/5 – 3/8 to 10000. Since all ports in an aggregate link must have the same key, the command in this example enables ports 1/1 – 1/4 and 3/5 – 3/8 to form a multi-slot aggregate link.

Syntax: [no] link-aggregate configure [system-priority <num>] | [port-priority <num>] | [key <num>] | [type server | switch]

The **system-priority** <num> parameter specifies the Foundry device's link aggregation priority. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

NOTE: On the BigIron MG8 and NetIron 40G running software release 02.2.01 and later, system-priority is configured at the global level.

The **port-priority** <num> parameter specifies an individual port's priority within the port group. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

The **key** <num> parameter identifies the group of ports that are eligible to be aggregated into a trunk group. The software automatically assigns a key to each group of ports. The software assigns the keys in ascending numerical order, beginning with 0. You can change a port group's key to a value from 0 – 65535.

NOTE: If you change the key for a port group, Foundry recommends that you use the value 10000 or higher, to avoid potential conflicts with dynamically created keys.

The **type server | switch** parameter specifies whether the port group is connected to a server (**server**) or to another networking device (**switch**). The default is **switch**.

You can enter one or more of the command's parameters on the same command line, in any order.

Configuring Link Aggregation System Priority on BigIron MG8 and NetIron 40G

On BigIron MG8 and NetIron 40G devices running software releases prior to release 02.2.01, link Aggregation system priority is configured as described in the section “Configuring Keys For Ports with Link Aggregation Enabled” on page 17-48. However, on BigIron MG8 and NetIron 40G devices running software releases 02.2.01 and later, system priority is configured at the GLOBAL level. For example:

```
BigIron MG8(config)#lacp system-priority 10
```

Syntax: lacp system-priority <value>

You can specify a priority from 0 – 65535. The default is 1

Link Aggregation on Tagged Ports for the FES

NOTE: This feature is supported on the FES in releases 03.4.00 and later.

Foundry software supports the IEEE 802.3ad standard for link aggregation. This standard describes the Link Aggregation Control Protocol (LACP), a mechanism for allowing ports on both sides of a redundant link to configure themselves into a trunk link (aggregate link), without the need for manual configuration of the ports into trunk groups.

When you enable link aggregation on a group of Foundry ports, the Foundry ports can negotiate with the ports at the remote ends of the links to establish trunk groups.

In releases prior to 03.4.00, link aggregation works with untagged ports (ports that belong to a single port-based VLAN). In release 03.4.00, link aggregation also works with 802.1q tagged ports (ports that belong to more than one port-based VLAN).

Configuration Rules

You do not need to perform any extra configuration steps to enable support of link aggregation on tagged ports. If you configure the port(s) for link aggregation, support for 802.1q tagged ports is automatically enabled.

NOTE: FES devices do not support dynamically allocated keys for tagged ports. These devices support statically assigned keys only.

Viewing Keys for Tagged Ports

To display link aggregation information, including the key for a *specific port*, enter a command such as the following at any level of the CLI:

```
FES Router# show link-aggregation ethernet 1/1

System ID: 00e0.52a9.bb00
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp]
1/1   0       0       0   No  L   No  No  No  No  No  No
```

The command in this example shows the key and other link aggregation information for port 1/1.

To display link aggregation information, including the key for *all ports* on which link aggregation is enabled, enter the following command at any level of the CLI:

```
FES Router# sh link-agg

System ID: 0004.8055.b200
Long timeout: 90, default: 90
Short timeout: 3, default: 3

Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
1/1    1        1    10000  Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
1/2    1        1    10000  Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
2/1    1        1    10000  Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
2/2    1        1    10000  Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
4/1    1        1     480   Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
4/2    1        1     480   Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
4/3    1        1     480   Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
4/4    1        1     480   Yes  S   Agg  Syn  Col  Dis  Def  No   Dwn
4/17   1        1     481   Yes  S   Agg  Syn  Col  Dis  Def  No   Ope
4/18   1        1     481   Yes  S   Agg  Syn  Col  Dis  Def  No   Ope
4/19   1        1     481   Yes  S   Agg  Syn  Col  Dis  Def  No   Ope
4/20   1        1     481   Yes  S   Agg  Syn  Col  Dis  Def  No   Ope
```

Syntax: show link-aggregation [ethernet <portnum>]

Possible values: N/A

Default value: N/A

Displaying and Determining the Status of Aggregate Links

Software release 07.6.03 and later provides the ability to determine the status of ports that are members of an aggregate link, and whether or not LACPDU messages are being transmitted between the ports. In releases prior to 07.6.03, this level of detail was not readily available. With the link aggregation enhancement, the **show link-aggregation** command provides the ability to view the status of dynamic links.

The following section provides details about the events that can affect the status of ports in an aggregate link and the status of LACP messages exchanged between the ports. Later sections provide instructions for viewing these status reports.

About Blocked Ports

Foundry devices can block traffic on a port or shut down a port that is part of a trunk group or aggregate link for the following reasons:

- For the purpose of link aggregation, the ports on Foundry devices are grouped into pairs of two; one odd-numbered port, and the next even-numbered port. When you configure link aggregation on a port (for instance, on an odd-numbered port), this port will be blocked and unable to join a trunk group until you configure the adjacent port (the even-numbered port) as part of the aggregate link. When you configure both ports with link aggregation and assign both ports the same key, both ports are able to join a trunk group. Once the ports become part of a trunk group, they can transmit and receive LACP packets.

NOTE: Ports that are configured as part of an aggregate link must also have the same key. For more information about assigning keys, see the section titled “Configuring Link Aggregation Parameters” in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

- When a port joins a trunk group and the port on the other end of the link shuts down or stops transmitting LACP packets, the Foundry device blocks the port. Depending on the timeout value set on the port, the link aggregation information expires.

NOTE: For more information about timeout values, see the section titled “Displaying Link Aggregation Information” in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

If either of these events occur, the Foundry device shuts down the port and notifies all the upper layer protocols that the port is down.

Foundry devices can also block traffic on a port that is initially configured with link aggregation. The port is blocked until it joins a trunk group. In this case, traffic is blocked, but the port is still operational.

A port remains blocked until one of the following events occur:

- Link aggregation is enabled on the adjacent port (the paired port) and both ports have the same key
- LACP brings the port back up
- The port joins a trunk group

Displaying Link Aggregation and Port Status Information

Use the **show link-aggregation** command to determine the operational status of ports associated with aggregate links.

To display the link aggregation information for a specific port, enter a command such as the following at any level of the CLI:

```
BigIron(config-mif-1/1-1/8)# show link-aggregation ethernet 1/1
System ID: 00e0.52a9.bb00
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp] [Ope]
1/1      0      0      0  No  L  No  No  No  No  No  No  No  Ope
```

The command in this example shows the link aggregation information for port 1/1.

NOTE: The **Ope** column displays in software releases 07.6.03 and later.

To display the link aggregation information for all ports on which link aggregation is enabled, enter the following command at any level of the CLI:

```
BigIron(config)# show link-aggregation
System ID: 00e0.52a9.bb00
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp] [Ope]
1/1      1      1      0  No  L  Agg  Syn  No  No  Def  Exp  Ope
1/2      1      1      0  No  L  Agg  Syn  No  No  Def  Exp  Ina
1/3      1      1      0  No  L  Agg  Syn  No  No  Def  Exp  Ina
1/4      1      1      0  No  L  Agg  Syn  No  No  Def  Exp  Blo
1/5      1      1      1  No  L  Agg  No   No  No  Def  Exp  Ope
1/6      1      1      1  No  L  Agg  No   No  No  Def  Exp  Ope
1/7      1      1      1  No  L  Agg  No   No  No  Def  Exp  Dwn
1/8      1      1      1  No  L  Agg  No   No  No  Def  Exp  Dwn
```

NOTE: The **Ope** column displays in software releases 07.6.03 and later.

Syntax: show link-aggregation [ethernet <portnum>]

Use **ethernet <portnum>** to display link-aggregation information for a specific port.

NOTE: Ports that are configured as part of an aggregate link must also have the same key. For more information about assigning keys, see the section titled "Configuring Link Aggregation Parameters" in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

The **show link aggregation** command shows the following information.

Table 18.11: CLI Display of Link Aggregation Information

This Field...	Displays...
System ID	Lists the base MAC address of the device. This is also the MAC address of port 1 (or 1/1).
Port	Lists the port number.
Sys P	Lists the system priority configured for this port.
Port P	Lists the port's link aggregation priority.
Key	Lists the link aggregation key.
Act	<p>Indicates the link aggregation mode, which can be one of the following:</p> <ul style="list-style-type: none"> No – The mode is passive or link aggregation is disabled (off) on the port. <p>If link aggregation is enabled (and the mode is passive), the port can send and receive LACPDU messages to participate in negotiation of an aggregate link initiated by another port, but cannot search for a link aggregation port or initiate negotiation of an aggregate link.</p> <ul style="list-style-type: none"> Yes – The mode is active. The port can send and receive LACPDU messages.
Tio	<p>Indicates the timeout value of the port. The timeout value can be one of the following:</p> <ul style="list-style-type: none"> L – Long. The trunk group has already been formed and the port is therefore using a longer message timeout for the LACPDU messages exchanged with the remote port. Typically, these messages are used as confirmation of the health of the aggregate link. S – Short. The port has just started the LACPDU message exchange process with the port at the other end of the link. The S timeout value also can mean that the link aggregation information received from the remote port has expired and the ports are starting a new information exchange.
Agg	<p>Indicates the link aggregation state of the port. The state can be one of the following:</p> <ul style="list-style-type: none"> Agg – Link aggregation is enabled on the port. No – Link aggregation is disabled on the port.

Table 18.11: CLI Display of Link Aggregation Information (Continued)

This Field...	Displays...
Syn	<p>Indicates the synchronization state of the port. The state can be one of the following:</p> <ul style="list-style-type: none"> • No – The port is out of sync with the remote port. The port does not understand the status of the LACPDU process and is not prepared to enter a trunk link. • Syn – The port is in sync with the remote port. The port understands the status of the LACPDU message exchange process, and therefore knows the trunk group to which it belongs, the link aggregation state of the remote port, and so on.
Col	<p>Indicates the collection state of the port, which determines whether the port is ready to send traffic over the trunk link.</p> <ul style="list-style-type: none"> • Col – The port is ready to send traffic over the trunk link. • No – The port is not ready to send traffic over the trunk link.
Dis	<p>Indicates the distribution state of the port, which determines whether the port is ready to receive traffic over the trunk link.</p> <ul style="list-style-type: none"> • Dis – The port is ready to receive traffic over the trunk link. • No – The port is not ready to receive traffic over the trunk link.
Def	<p>Indicates whether the port is using default link aggregation values. The port uses default values if it has not received link aggregation information through LACP from the port at the remote end of the link. This field can have one of the following values:</p> <ul style="list-style-type: none"> • Def – The port has not received link aggregation values from the port at the other end of the link and is therefore using its default link aggregation LACP settings. • No – The port has received link aggregation information from the port at the other end of the link and is using the settings negotiated with that port.
Exp	<p>Indicates whether the negotiated link aggregation settings have expired. The settings expire if the port does not receive an LACPDU message from the port at the other end of the link before the message timer expires. This field can have one of the following values:</p> <ul style="list-style-type: none"> • Exp – The link aggregation settings this port negotiated with the port at the other end of the link have expired. The port is now using its default link aggregation settings. • No – The link aggregation values that this port negotiated with the port at the other end of the link have not expired, so the port is still using the negotiated settings.

Table 18.11: CLI Display of Link Aggregation Information (Continued)

This Field...	Displays...
Ope	<ul style="list-style-type: none"> • Ope (operational) - The port is operating normally. • Ina (inactive) - The port is inactive because the port on the other side of the link is down or has stopped transmitting LACP packets. • Blo (blocked) - The port is blocked because the adjacent port is not configured with link aggregation or because it is not able to join a trunk group. To unblock the port and bring it to an operational state, enable link aggregation on the adjacent port and ensure that the ports have the same key.

Displaying Trunk Group and LACP Status Information

Use the **show trunk** command to determine the status of LACP. See “Displaying Trunk Group Configuration Information” on page 17-34.

Clearing the Negotiated Link Aggregations

When a group of ports negotiates a trunk group configuration, the software stores the negotiated configuration in a table. You can clear the negotiated link aggregation configurations from the software. When you clear the information, the software does not remove link aggregation parameter settings you have configured. Only the configuration information negotiated using LACP is removed.

NOTE: The software automatically updates the link aggregation configuration based on LACPDU messages. However, clearing the link aggregation information can be useful if you are troubleshooting a configuration.

To clear the link aggregation information, use the following CLI method.

USING THE CLI

To clear the link aggregation information, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# clear link-aggregate
```

Syntax: clear link-aggregate

Chapter 19


Configuring Uni-Directional Link Detection (UDLD)

Uni-directional Link Detection (UDLD) monitors a link between two Foundry devices and brings the ports on both ends of the link down if the link goes down at any point between the two devices. This feature is useful for links that are individual ports and for trunk links. Figure 19.1 shows an example.

Figure 19.1 UDLD example

Without link keepalive, the Foundry ports remain enabled. Traffic continues to be load balanced to the ports connected to the failed link.

When link keepalive is enabled, the feature brings down the Foundry ports connected to the failed link.



Normally, a Foundry device load balances traffic across the ports in a trunk group. In this example, each Foundry device load balances traffic across two ports. Without the UDLD feature, a link failure on a link that is not directly attached to one of the Foundry devices is undetected by the Foundry devices. As a result, the Foundry devices continue to send traffic on the ports connected to the failed link.

When UDLD is enabled on the trunk ports on each Foundry device, the devices detect the failed link, disable the ports connected to the failed link, and use the remaining ports in the trunk group to forward the traffic.

Ports enabled for UDLD exchange proprietary health-check packets once every second (the keepalive interval). If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for two more intervals. If the port still does not receive a health-check packet after waiting for three intervals, the port concludes that the link has failed and takes the port down.

NOTE: On FastIron Edge Switches, this feature is supported on devices running software release 03.1.00 and later.

Configuration Considerations

- The feature is supported only on Ethernet ports.
- To configure UDLD on a trunk group, you must configure the feature on each port of the group individually. Configuring UDLD on a trunk group's primary port enables the feature on that port only.
- Dynamic trunking is not supported. If you want to configure a trunk group that contains ports on which UDLD is enabled, you must remove the UDLD configuration from the ports. After you create the trunk group, you can re-add the UDLD configuration.

Configuring UDLD

To enable UDLD on a port, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# link-keepalive ethernet 1/1
```

Syntax: [no] link-keepalive ethernet <portnum> [ethernet <portnum>]

To enable the feature on a trunk group, enter commands such as the following:

```
BigIron(config)# link-keepalive ethernet 1/1 ethernet 1/2  
BigIron(config)# link-keepalive ethernet 1/3 ethernet 1/4
```

These commands enable UDLD on ports 1/1 – 1/4. You can specify up to two ports on the same command line.

Changing the Keepalive Interval

By default, ports enabled for UDLD send a link health-check packet once every 500 ms. You can change the interval to a value from 1 – 60, where 1 is 100 ms, 2 is 200 ms, and so on. To change the interval, enter a command such as the following:

```
BigIron(config)# link-keepalive interval 3
```

Syntax: [no] link-keepalive interval <num>

The <num> parameter specifies how often the ports send a UDLD packet. You can specify from 1 – 60, in 100 ms increments. The default is 5 (500 ms).

Changing the Keepalive Retries

By default, a port waits one second to receive a health-check reply packet from the port at the other end of the link. If the port does not receive a reply, the port tries four more times by sending up to four more health-check packets. If the port still does not receive a reply after the maximum number of retries, the port goes down.

You can change the maximum number of keepalive attempts to a value from 3 – 10. To change the maximum number of attempts, enter a command such as the following:

```
BigIron(config)# link-keepalive retries 4
```

Syntax: [no] link-keepalive retries <num>

The <num> parameter specifies the maximum number of times the port will try the health check. You can specify a value from 3 – 10. The default is 5.

UDLD for Tagged Ports

The default implementation of UDLD sends the packets untagged, even across tagged ports. If the untagged UDLD packet is received by a third-party switch, that switch may reject the packet. As a result, UDLD may be limited only to Foundry devices, since UDLD may not function on third-party switches.

Beginning with Enterprise software release 07.6.06, you can configure ports to send out UDLD control packets that are tagged with a specific VLAN ID as tagged UDLD control packets. The enhancement also allows third party switches to receive the control packets that are tagged with the specified VLAN.

To allow ports to receive and send UDLD control packets tagged with a specific VLAN ID, enter commands such as the following:

```
BigIron(config)# link-keepalive ethernet 1/18 vlan 22
```

This commands enables UDLD on port 1/18 and allows UDLD control packet tagged with VLAN 22 to be received and sent on port 1/18.

Syntax: [no] link-keepalive ethernet <portnum> [vlan <vlan-ID>]

Enter the slot number (if applicable) and the port number of the Ethernet port.

Enter the ID of the VLAN that the UDLD control packets can contain to be received and sent on the port. If a VLAN ID is not specified, then UDLD control packets are sent out of the port as untagged packets.

NOTE: You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.

Displaying UDLD Information

Displaying Information for All Ports

To display UDLD information for all ports, enter the following command:

```
BigIron(config)# show link-keepalive
Total link-keepalive enabled ports: 4
Keepalive Retries: 3      Keepalive Interval: 1 Sec.
```

Port	Physical Link	Logical Link	State
4/1	up	up	FORWARDING
4/2	up	up	FORWARDING
4/3	down	down	DISABLED
4/4	up	down	DISABLED

Syntax: show link-keepalive [ethernet <portnum>]

Table 19.1: CLI Display of UDLD Information

This Field...	Displays...
Total link-keepalive enabled ports	The total number of ports on which UDLD is enabled.
Keepalive Retries	The number of times a port will attempt the health check before concluding that the link is down.
Keepalive Interval	The number of seconds between health check packets.
Port	The port number.
Physical Link	The state of the physical link. This is the link between the Foundry port and the directly connected device.
Logical Link	The state of the logical link. This is the state of the link between this Foundry port and the Foundry port on the other end of the link.
State	The traffic state of the port.

If a port is disabled by UDLD, the change also is indicated in the output of the **show interfaces brief** command. Here is an example:

```
BigIron(config)# show interface brief

Port  Link State      Dupl Speed Trunk Tag Priori MAC           Name
1/1   Up    LK-DISABLENone None  None No  level0 00e0.52a9.bb00
1/2   Down None           None None  None No  level0 00e0.52a9.bb01
1/3   Down None           None None  None No  level0 00e0.52a9.bb02
1/4   Down None           None None  None No  level0 00e0.52a9.bb03
```

If the port was already down before you enabled UDLD for the port, the port's state is listed as None.

Syntax: show interface brief

Beginning with Enterprise software release 07.6.06, the **show link-keepalive** command shows the following:

```
BigIron(config)# show link-keepalive ethernet
Current State      : down           Remote MAC Addr   : 0000.0000.0000
Local Port         : 1/1             Remote Port       : n/a
Local System ID    : e0eb8e00      Remote System ID  : 00000000
Packets sent       : 0             Packets received  : 0
Transitions        : 0             Link-vlan       : 100
Port blocking      : No            BM disabled       : Yes
```

The Link-vlan entry shows the ID of the tagged VLAN in the UDLD packet.

Syntax: show link-keepalive ethernet

Displaying Information for a Single Port

To display detailed UDLD information for a specific port, enter a command such as the following:

```
BigIron(config)# show link-keepalive ethernet 4/1

Current State      : up             Remote MAC Addr   : 00e0.52d2.5100
Local Port         : 4/1             Remote Port       : 2/1
Local System ID    : e0927400      Remote System ID  : e0d25100
Packets sent       : 254            Packets received  : 255
Transitions        : 1

Port blocking      : No            BM disabled       : No
```

Table 19.2: CLI Display of Detailed UDLD Information

This Field...	Displays...
Current State	The state of the logical link. This is the link between this Foundry port and the Foundry port on the other end of the link.
Remote MAC Addr	The MAC address of the port or device at the remote end of the logical link.

Table 19.2: CLI Display of Detailed UDLD Information (Continued)

This Field...	Displays...
Local Port	The port number on this Foundry device.
Remote Port	The port number on the Foundry device at the remote end of the link.
Local System ID	A unique value that identifies this Foundry device. The ID can be used by Foundry technical support for troubleshooting.
Remote System ID	A unique value that identifies the Foundry device at the remote end of the link.
Packets sent	The number of UDLD health-check packets sent on this port.
Packets received	The number of UDLD health-check packets received on this port.
Transitions	The number of times the logical link state has changed between up and down.
Port blocking	Information used by Foundry technical support for troubleshooting.
BM disabled	Information used by Foundry technical support for troubleshooting.

The **show interface ethernet** <portnum> command also displays the UDLD state for an individual port. In addition, the line protocol state listed in the first line will say “down” if UDLD has brought the port down. Here is an example:

```
BigIron(config)# show interface ethernet 1/1
FastEthernet1/1 is down, line protocol is down, link keepalive is enabled
  Hardware is FastEthernet, address is 00e0.52a9.bbca (bia 00e0.52a9.bbca)
  Configured speed auto, actual unknown, configured duplex fdx, actual unknown
  Member of L2 VLAN ID 1, port is untagged, port state is DISABLED
  STP configured to ON, priority is level0, flow control enabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 0 giants, DMA received 0 packets
  19 packets output, 1216 bytes, 0 underruns
  Transmitted 0 broadcasts, 19 multicasts, 0 unicasts
  0 output errors, 0 collisions, DMA transmitted 19 packets
```

In this example, the port has been brought down by UDLD. Notice that in addition to the information in the first line, the port state on the fourth line of the display is listed as DISABLED.

Clearing UDLD Statistics

To clear UDLD statistics, enter the following command:

```
BigIron# clear link-keepalive statistics
```

Syntax: clear link-keepalive statistics

This command clears the Packets sent, Packets received, and Transitions counters in the **show link keepalive ethernet <portnum> display**.

Chapter 20

Configuring Metro Features

This chapter describes the following Metro features:

- Topology groups – A topology group enables you to control the Layer 2 protocol configuration and Layer 2 state of a set of ports in multiple VLANs based on the configuration and states of those ports in a single master VLAN. One instance of the Layer 2 protocol controls all the VLANs. See “Topology Groups”.
- Metro Ring Protocol – MRP is an alternative to STP that provides Layer 2 redundancy and sub-second failover in ring topologies. See “Metro Ring Protocol (MRP)” on page 20-5.
- Virtual Switch Redundancy Protocol (VSRP) – VSRP is an alternative to STP that provides Layer 2 and Layer 3 redundancy and sub-second failover in mesh topologies. See “Virtual Switch Redundancy Protocol (VSRP)” on page 20-25.

You can use these features individually or in combination to provide fast, reliable, and easy to configure Layer 2 connectivity in your Metro network.

Topology Groups

A topology group is a named set of VLANs that share a Layer 2 topology. Topology groups simplify configuration and enhance scalability of Layer 2 protocols by allowing you to run a single instance of a Layer 2 protocol on multiple VLANs.

You can use topology groups with the following Layer 2 protocols:

- STP
- MRP
- VSRP
- 802.1W

Topology groups simplify Layer 2 configuration and provide scalability by enabling you to use the same instance of a Layer 2 protocol for multiple VLANs. For example, if a Foundry device is deployed in a Metro network and provides forwarding for two MRP rings that each contain 128 VLANs, you can configure a topology group for each ring. If a link failure in a ring causes a topology change, the change is applied to all the VLANs in the ring's topology group. Without topology groups, you would need to configure a separate ring for each VLAN.

NOTE: Topology groups are supported on FastIron Edge Switches running software release 03.1.00.

NOTE: If you plan to use a configuration saved under an earlier software release and the configuration contains STP groups, the CLI converts the STP groups into topology groups when you save the configuration under software release 07.6.01. For backward compatibility, you can still use the STP group commands. However, the CLI converts the commands into the topology group syntax. Likewise, the **show stp-group** command displays STP topology groups.

Master VLAN and Member VLANs

Each topology group contains a master VLAN and can contain one or more member VLANs and VLAN groups.

- **Master VLAN** – The master VLAN contains the configuration information for the Layer 2 protocol. For example, if you plan to use the topology group for MRP, the topology group's master VLAN contains the ring configuration information.
- **Member VLANs** – The member VLANs are additional VLANs that share ports with the master VLAN. The Layer 2 protocol settings for the ports in the master VLAN apply to the same ports in the member VLANs. A change to the master VLAN's Layer 2 protocol configuration or Layer 2 topology affects all the member VLANs. Member VLANs do not independently run a Layer 2 protocol.
- **Member VLAN groups** – A VLAN group is a named set of VLANs. The VLANs within a VLAN group have the same ports and use the same values for other VLAN parameters.

When a Layer 2 topology change occurs on a port in the master VLAN, the same change is applied to that port in all the member VLANs that contain the port. For example, if you configure a topology group whose master VLAN contains ports 1/1 and 1/2, a Layer 2 state change on port 1/1 applies to port 1/1 in all the member VLANs that contain that port. However, the state change does not affect port 1/1 in VLANs that are not members of the topology group.

Control Ports and Free Ports

A port that is in a topology group can be a control port or a free port.

- **Control port** – A control port is a port in the master VLAN, and is therefore controlled by the Layer 2 protocol configured in the master VLAN. The same port in all the member VLANs is controlled by the master VLAN's Layer 2 protocol. Each member VLAN must contain all of the control ports and can contain additional ports.
- **Free port** – A free port is not controlled by the master VLAN's Layer 2 protocol. The master VLAN can contain free ports. (In this case, the Layer 2 protocol is disabled on those ports.) In addition, any ports in the member VLANs that are not also in the master VLAN are free ports.

NOTE: Since free ports are not controlled by the master port's Layer 2 protocol, they are assumed to always be in the Forwarding state.

Configuration Considerations

- You can configure up to 256 topology groups. Each group can control up to 4096 VLANs. A VLAN cannot be controlled by more than one topology group.
- You must configure the master VLAN and member VLANs or member VLAN groups before you configure the topology group.
- The topology group must contain a master VLAN and can also contain individual member VLANs, VLAN groups, or a combination of individual member VLANs and VLAN groups.
- Once you add a VLAN as a member of a topology group, all the Layer 2 protocol information on the VLAN is deleted.

Configuring a Topology Group

To configure a topology group, enter commands such as the following:

```
BigIron(config)# topology-group 2
BigIron(config-topo-group-2)# master-vlan 2
BigIron(config-topo-group-2)# member-vlan 3
BigIron(config-topo-group-2)# member-vlan 4
BigIron(config-topo-group-2)# member-vlan 5
BigIron(config-topo-group-2)# member-group 2
```

These commands create topology group 2 and add the following:

- Master VLAN 2
- Member VLANs 2, 3, and 4
- Member VLAN group 2

Syntax: [no] topology-group <group-id>

The <group-id> parameter specifies the topology group ID and can be from 1 – 256.

Syntax: [no] master-vlan <vlan-id>

This command adds the master VLAN. The VLAN must already be configured. Make sure all the Layer 2 protocol settings in the VLAN are correct for your configuration before you add the VLAN to the topology group. A topology group can have only one master VLAN.

NOTE: If you remove the master VLAN (by entering **no master-vlan <vlan-id>**), the software selects the next-highest numbered member VLAN as the new master VLAN. For example, if you remove master VLAN 2 from the example above, the CLI converts member VLAN 3 into the new master VLAN. The new master VLAN inherits the Layer 2 protocol settings of the older master VLAN.

NOTE: If you add a new master VLAN to a topology group that already has a master VLAN, the new master VLAN replaces the older master VLAN. All member VLANs and VLAN groups follow the Layer 2 protocol settings of the new master VLAN.

Syntax: [no] member-vlan <vlan-id>

The <vlan-id> parameter specifies a VLAN ID. The VLAN must already be configured.

Syntax: [no] member-group <num>

The <num> specifies a VLAN group ID. The VLAN group must already be configured.

NOTE: Once you add a VLAN or VLAN group as a member of a topology group, all the Layer 2 protocol configuration information for the VLAN or group is deleted. For example, if STP is configured on a VLAN and you add the VLAN to a topology group, the STP configuration is removed from the VLAN. Once you add the VLAN to a topology group, the VLAN uses the Layer 2 protocol settings of the master VLAN.

If you remove a member VLAN or VLAN group from a topology group, you will need to reconfigure the Layer 2 protocol information in the VLAN or VLAN group.

Displaying Topology Group Information

The following sections show how to display STP information and topology group information for VLANs.

Displaying STP Information

To display STP information for a VLAN, enter a command such as the following:

```
BigIron(config)# show span vlan 4
VLAN 4 BPDU cam_index is 14344 and the Master DMA Are(HEX) 18 1A
STP instance owned by VLAN 2
```

This example shows STP information for VLAN 4. The line shown in bold type indicates that the VLAN's STP configuration is controlled by VLAN 2. This information indicates that VLAN 4 is a member of a topology group and VLAN 2 is the master VLAN in that topology group.

Displaying Topology Group Information

To display topology group information, enter the following command:

```
BigIron(config)# show topology-group

Topology Group 3
=====
master-vlan 2
member-vlan none

Common control ports          L2 protocol
ethernet 1/1                  MRP
ethernet 1/2                  MRP
ethernet 1/5                  VSRP
ethernet 2/22                 VSRP
Per vlan free ports
ethernet 2/3                  Vlan 2
ethernet 2/4                  Vlan 2
ethernet 2/11                 Vlan 2
ethernet 2/12                 Vlan 2
```

Syntax: show topology-group [<group-id>]

This display shows the following information.

Table 20.1: CLI Display of Topology Group Information

This Field...	Displays...
master-vlan	The master VLAN for the topology group. The settings for STP, MRP, or VSRP on the control ports in the master VLAN apply to all control ports in the member VLANs within the topology group.
member-vlan	The member VLANs in the topology group.
Common control ports	The master VLAN ports that are configured with Layer 2 protocol information. The Layer 2 protocol configuration and state of these ports in the master VLAN applies to the same port numbers in all the member VLANs.

Table 20.1: CLI Display of Topology Group Information (Continued)

This Field...	Displays...
L2 protocol	The Layer 2 protocol configured on the control ports. The Layer 2 protocol can be one of the following: <ul style="list-style-type: none"> • MRP • STP • VSRP
Per vlan free ports	The ports that are not controlled by the Layer 2 protocol information in the master VLAN.

Metro Ring Protocol (MRP)

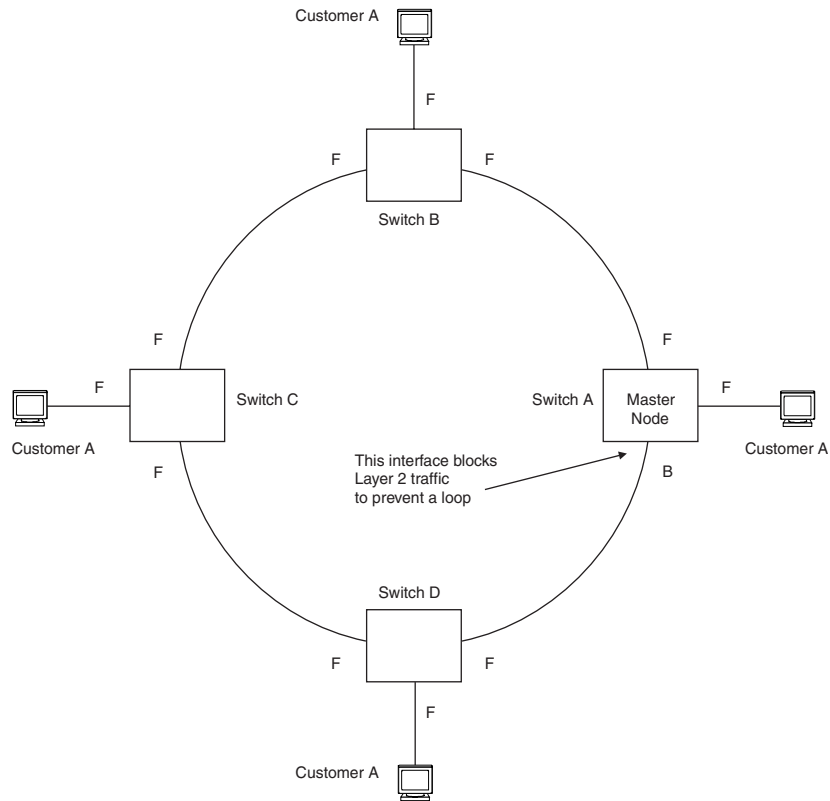
NOTE: This feature is available in software release 07.6.01 and later ; however in release 07.6.01 , this feature is not supported in the B2R Layer 3 Switch image. Beginning with software release 07.7.00, this limitation has been removed. It is now supported on all images of software release 07.7.00 and later. Also, this feature is supported in FES software release 03.1.00 and later.

The Metro Ring Protocol (MRP) is a Foundry proprietary protocol that prevents Layer 2 loops and provides fast reconvergence in Layer 2 ring topologies. It is an alternative to STP and is especially useful in Metropolitan Area Networks (MANs) where using STP has the following drawbacks:

- STP allows a maximum of seven nodes. Metro rings can easily contain more nodes than this.
- STP has a slow reconvergence time, taking many seconds or even minutes. MRP can detect and heal a break in the ring in sub-second time.

Figure 20.1 shows an example of an MRP metro ring.

Figure 20.1 Metro ring – normal state



The ring in this example consists of four MRP nodes (Foundry switches). Each node has two interfaces with the ring. Each node also is connected to a separate customer network. The nodes forward Layer 2 traffic to and from the customer networks through the ring. The ring interfaces are all in one port-based VLAN. Each customer interface can be in the same VLAN as the ring or in a separate VLAN.

One node, is configured as the master node of the MRP ring. One of the two interfaces on the master node is configured as the primary interface; the other is the secondary interface. The primary interface originates Ring Health Packets (RHPs), which are used to monitor the health of the ring. An RHP is forwarded on the ring to the next interface until it reaches the secondary interface of the master node. The secondary interface blocks the packet to prevent a Layer 2 loops.


NOTE: When you configure MRP, Foundry recommends that you disable one of the ring interfaces before beginning the ring configuration. Disabling an interface prevents a Layer 2 loop from occurring while you are configuring MRP on the ring nodes. Once MRP is configured and enabled on all the nodes, you can re-enable the interface.

MRP Rings Without Shared Interfaces (MRP Phase 1)

MRP Phase 1 allows you to configure multiple MRP rings, as shown in Figure 20.2, but the rings cannot share the same link. For example, you cannot configure ring 1 and ring 2 to each have interfaces 1/1 and 1/2.

Also, when you configured an MRP ring, any node on the ring that can be designated as the master node for the ring. A master node can be the master node of more than one ring. (See Figure 20.2.) Each ring is an independent ring and RHP packets are processed within each ring.

Figure 20.2 Metro ring – multiple rings



In this example, two nodes are each configured with two MRP rings. Any node in a ring can be the master for its ring. A node also can be the master for more than one ring.

NOTE: FES devices running software release 03.1.00 or later are capable of being configured as MRP masters or MRP members (for different rings).

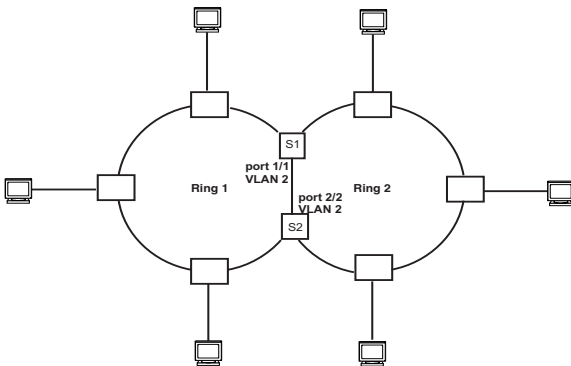
MRP Rings with Shared Interfaces (MRP Phase 2)

NOTE: This feature is supported on devices with IronCore and JetCore modules running Enterprise software running release 07.7.00 and later, devices running Service Provider software release 09.2.00 and later, the BigIron MG8 and NetIron 40G running Terathon IronWare release 02.2.00 and later, and the FES running release 03.4.00 and later.


With MRP Phase 2, MRP rings can be configured to share the same interfaces as long as the interfaces belong to the same VLAN. Figure 20.3 shows examples of multiple MRP rings that share the same interface.

Figure 20.3 Examples of multiple rings sharing the same interface

Example 1




Example 2



On each node that will participate in the ring, you specify the ring's ID and the interfaces that will be used for ring traffic. In a multiple ring configuration, a ring's ID determines its priority. The lower the ring ID, the higher the priority of a ring.

A ring's ID is also used to identify the interfaces that belong to a ring.

Figure 20.4 Interface IDs and Types on Rings with Shared Interfaces



For example, in Figure 20.4, the ID of all interfaces on all nodes on Ring 1 is 1 and all interfaces on all nodes on Ring 2 is 2. Port 1/1 on node S1 and Port 2/2 on S2 have the IDs of 1 and 2 since the interfaces are shared by Rings 1 and 2.

The ring's ID is also used to determine an interface's priority. Generally, a ring's ID is also the ring's priority and the priority of all interfaces on that ring. However, if the interface is shared by two or more rings, then the highest priority (lowest ID) becomes the priority of the interface. For example, in Figure 20.4, all interfaces on Ring 1, except for Port 1/1 on node S1 and Port 2/2 on node S2 have a priority of 1. Likewise, all interfaces on Ring 2, except for Port 1/1 on node S1 and Port 2/2 on node S2 have a priority of 2. Port 1/1 on S1 and Port 2/2 on S2 have a priority of 1 since 1 is the highest priority (lowest ID) of the rings that share the interface.

If a node has interfaces that have different IDs, the interfaces that belong to the ring with the highest priority become regular ports. Those interfaces that do not belong to the ring with the highest priority become tunnel

ports. In Figure 20.4, nodes S1 and S2 have interfaces that belong to Rings 1 and 2. Those interfaces with a priority of 1 are regular ports. The interfaces with a priority of 2 are the tunnel ports since they belong to Ring 2, which has a lower priority than Ring 1.

Selection of Master Node on Shared Interfaces

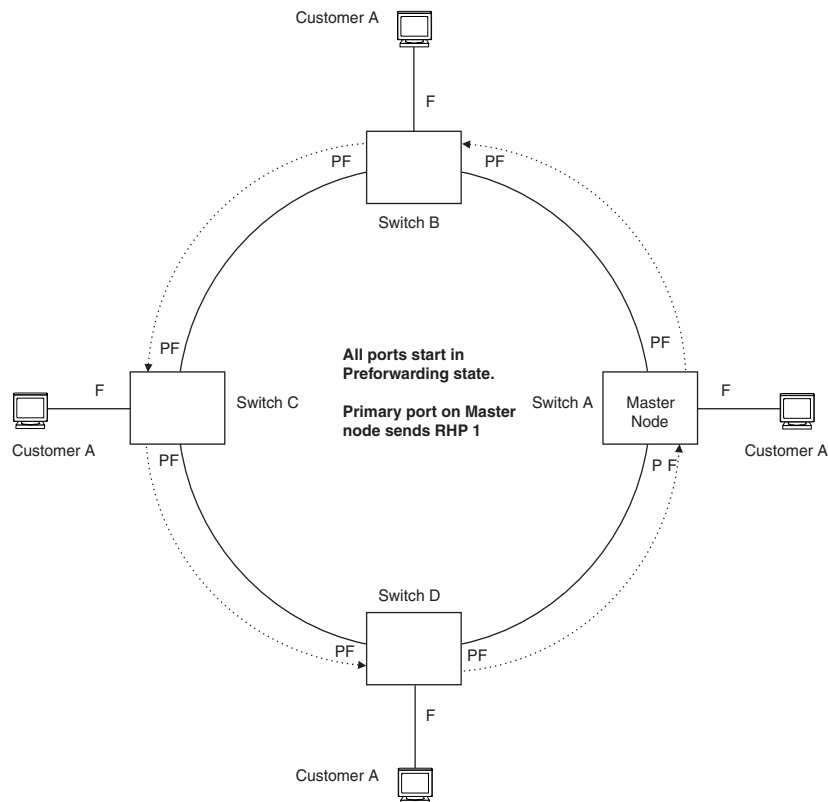
Allowing MRP rings to share interfaces limits the nodes that can be designated as the master node. Any node on an MRP ring that does not have a shared interface can be designated as the ring's master node. However, if all nodes on the ring have shared interfaces, nodes that do not have tunnel ports can be designated as the master node of that ring. If none of the nodes meet these criteria, you must change the rings' priorities by reconfiguring the rings' ID.

In Figure 20.4, any of the nodes on Ring 1, even S1 or S2, can be a master node since none of its interfaces are tunnel ports. However in Ring 2, neither S1 nor S2 can be a master node since these nodes contain tunnel ports.

Ring Initialization

The ring shown in Figure 20.1 shows the port states in a fully initialized ring without any broken links. Figure 20.5 shows the initial state of the ring, when MRP is first enabled on the ring's switches. All ring interfaces on the master node and member nodes begin in the Preforwarding state (PF).

Figure 20.5 Metro ring – initial state



MRP uses Ring Health Packets (RHPs) to monitor the health of the ring. An RHP is an MRP protocol packet. The source address is the MAC address of the master node and the destination MAC address is a protocol address for MRP. The Master node generates RHPs and sends them on the ring. The state of a ring port depends on the RHPs.

A ring interface can have one of the following MRP states:

- Preforwarding (PF) – The interface can forward RHPs but cannot forward data. All ring ports being in this state when you enable MRP.

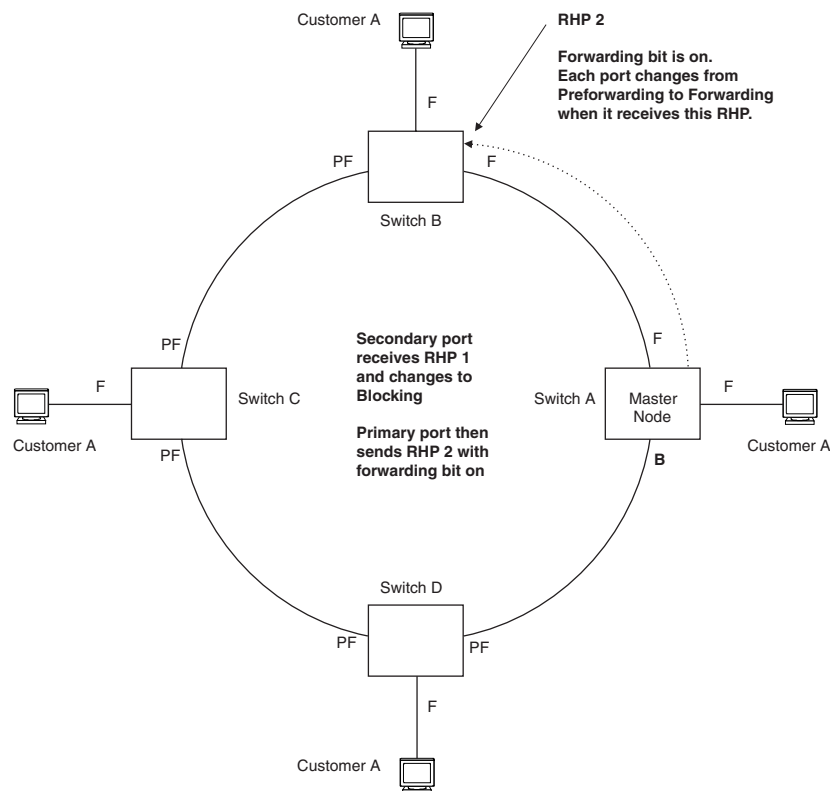
- Forwarding (F) – The interface can forward data as well as RHPs. An interface changes from Preforwarding to Forwarding when the port's preforwarding time expires. This occurs if the port does not receive an RHP from the Master, or if the forwarding bit in the RHPs received by the port is off. This indicates a break in the ring. The port heals the ring by changing its state to Forwarding. The preforwarding time is the number of milliseconds the port will remain in the Preforwarding state before changing to the Forwarding state, even without receiving an RHP.
- Blocking (B) – The interface can process RHPs, but cannot forward data. Only the secondary interface on the Master node can be Blocking.

When MRP is enabled, all ports begin in the Preforwarding state. The primary interface on the Master node, although it is in the Preforwarding state like the other ports, immediately sends an RHP onto the ring. The secondary port on the Master node listens for the RHP.

- If the secondary port receives the RHP, all links in the ring are up and the port changes its state to Blocking. The primary port then sends another MRP with its forwarding bit set on. As each of the member ports receives the RHP, the ports change their state to Forwarding. Typically, this occurs in sub-second time. The ring very quickly enters the fully initialized state.
- If the secondary port does not receive the RHP by the time the preforwarding time expires, a break has occurred in the ring. The port changes its state to Forwarding. The member ports also change their states from Preforwarding to Forwarding as their preforwarding timers expire. The ring is not intact, but data can still travel among the nodes using the links that are up.

Figure 20.6 shows an example.


Figure 20.6 Metro ring – from Preforwarding to Forwarding



Each RHP also has a sequence number. MRP can use the sequence number to determine the round-trip time for RHPs in the ring. See "Using MRP Diagnostics" on page 20-16.

Figure 20.7 shows an example of how RHP packets are processed normally in MRP rings with shared interfaces.

Figure 20.7 Flow of RHP packets on MRP Rings with Shared Interfaces



Port 2/1 on Ring 1's master node is the primary interface of the master node. The primary interface forwards an RHP packet on the ring. Since all the interfaces on Ring 1 are regular ports, the RHP packet is forwarded to all the interfaces until it reaches Port 2/2, the secondary interface of the master node. Port 2/2 then blocks the packet to complete the process.

On Ring 2, Port 3/1, is the primary interface of the master node. It sends an RHP packet on the ring. Since all ports on S4 are regular ports, the RHP packet is forwarded on those interfaces. When the packet reaches S2, the receiving interface is a tunnel port. The port compares the packet's priority to its priority. Since the packet's priority is the same as the tunnel port's priority, the packet is forwarded up the link shared by Rings 1 and 2.


When the RHP packet reaches the interface on node S2 shared by Rings 1 and 2, the packet is forwarded since its priority is less than the interface's priority. The packet continues to be forwarded to node S1 until it reaches the tunnel port on S1. That tunnel port determines that the RHP packet's priority is equal to the port's priority and forwards the packet. The RHP packet is forwarded to the remaining interfaces on Ring 2 until it reaches port 3/2, the secondary interface of the master node. Port 3/2 then blocks the packet to prevent a loop.

When the RHP packet from Ring 2 reached S2, it was also forwarded from S2 to S3 on Ring 1 since the port on S2 has a higher priority than the RHP packet. The packets is forwarded around Ring 1 until it reaches port 2/2, Ring 1's the secondary port. The RHP packet is then blocked by that port.

How Ring Breaks Are Detected and Healed

Figure 20.8 shows ring interface states following a link break. MRP quickly heals the ring and preserves connectivity among the customer networks.

Figure 20.8 Metro ring – ring break



If a break in the ring occurs, MRP heals the ring by changing the states of some of the ring interfaces.

- Blocking interface – The Blocking interface on the Master node has a dead timer. If the dead time expires before the interface receives one of its ring’s RHPs, the interface changes state to Preforwarding. Once the secondary interface changes state to Preforwarding:
 - If the interface receives an RHP, the interface changes back to the Blocking state and resets the dead timer.
 - If the interface does not receive an RHP for its ring before the Preforwarding time expires, the interface changes to the Forwarding state, as shown in Figure 20.8.
- Forwarding interfaces – Each member interface remains in the Forwarding state.


When the broken link is repaired, the link’s interfaces come up in the Preforwarding state, which allows RHPs to travel through the restored interfaces and reach the secondary interface on the Master node.

- If an RHP reaches the Master node’s secondary interface, the ring is intact. The secondary interface changes to Blocking. The Master node sets the forwarding bit on in the next RHP. When the restored interfaces receive this RHP, they immediately change state to Forwarding.
- If an RHP does not reach the Master node’s secondary interface, the ring is still broken. The Master node does not send an RHP with the forwarding bit on. In this case, the restored interfaces remain in the Preforwarding state until the preforwarding timer expires, then change to the Forwarding state.

If the link between **shared interfaces** breaks (Figure 20.9), the secondary interface on Ring 1’s master node changes to a preforwarding state. The RHP packet sent by port 3/1 on Ring 2 is forwarded through the interfaces on S4, then to S2. The packet is then forwarded through S2 to S3, but not from S2 to S1 since the link between the two nodes is not available. When the packet reaches Ring 1’s master node, the packet is forwarded through the secondary interface since it is currently in a preforwarding state. A secondary interface in preforwarding mode ignores any RHP packet that is not from its ring. The secondary interface changes to blocking mode only when the RHP packet forwarded by its primary interface is returned.

The packet then continues around Ring 1, through the interfaces on S1 to Ring 2 until it reaches Ring 2's master node. Port 3/2, the secondary interface on Ring 2 changes to blocking mode since it received its own packet, then blocks the packet to prevent a loop.

Figure 20.9 Flow of RHP packets when a link for shared interfaces brakes

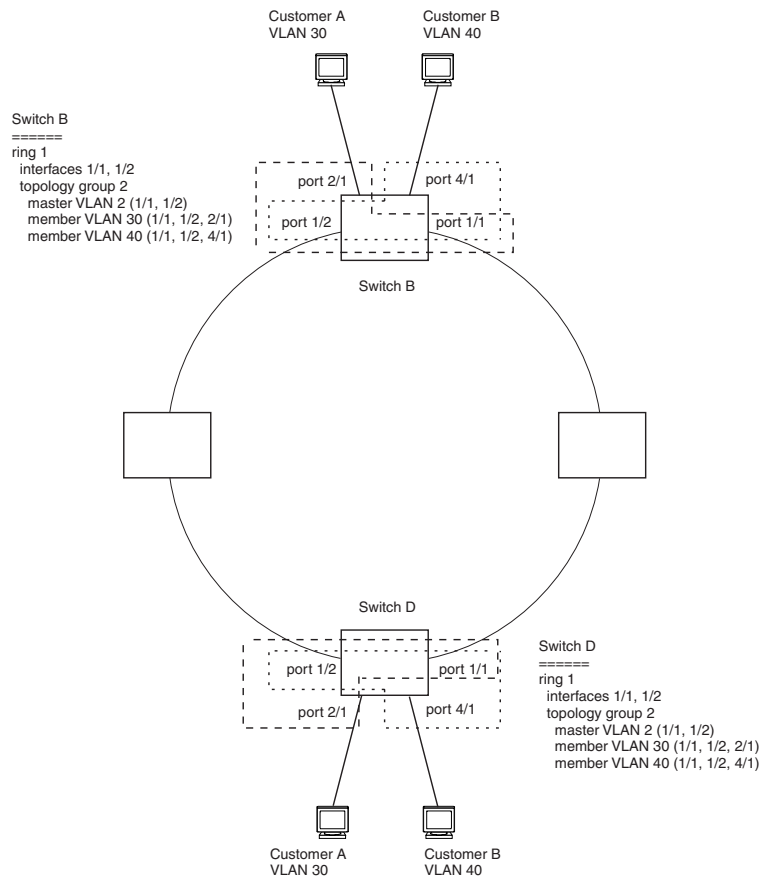


RHP packets follow this flow until the link is restored; then the RHP packet returns to its normal flow as shown in Figure 20.7.

Master VLANs and Customer VLANs

All the ring ports must be in the same VLAN. Placing the ring ports in the same VLAN provides Layer 2 connectivity for a given customer across the ring. Figure 20.10 shows an example.

Figure 20.10 Metro ring – ring VLAN and customer VLANs



Notice that each customer has their own VLAN. Customer A has VLAN 30 and Customer B has VLAN 40. Customer A's host attached to Switch D can reach the Customer A host attached to Switch B at Layer 2 through the ring. Since Customer A and Customer B are on different VLANs, they will not receive each other's traffic.

You can configure MRP separately on each customer VLAN. However, this is impractical if you have many customers. To simplify configuration when you have a lot of customers (and therefore a lot of VLANs), you can use a topology group.

A topology group enables you to control forwarding in multiple VLANs using a single instance of a Layer 2 protocol such as MRP. A topology group contains a master VLAN and member VLANs. The master VLAN contains all the configuration parameters for the Layer 2 protocol (STP, MRP, or VSRP). The member VLANs use the Layer 2 configuration of the master VLAN.

In Figure 20.10, VLAN 2 is the master VLAN and contains the MRP configuration parameters for ring 1. VLAN 30 and VLAN 40, the customer VLANs, are member VLANs in the topology group. Since a topology group is used, a single instance of MRP provides redundancy and loop prevention for both the customer VLANs.

If you use a topology group:

- The master VLAN must contain the ring interfaces. The ports must be tagged, since they will be shared by multiple VLANs.
- The member VLAN for a customer must contain the two ring interfaces and the interfaces for the customer. Since these interfaces are shared with the master VLAN, they must be tagged. Do not add another customer's interfaces to the VLAN.

For more information about topology groups, see "Topology Groups" on page 20-1.

See "MRP CLI Example" on page 20-20 for the configuration commands required to implement the MRP configuration shown in Figure 20.10.

Configuring MRP

To configure MRP, perform the following tasks. You need to perform the first task on only one of the nodes. Perform the remaining tasks on all the nodes.

- Disable one of the ring interfaces. This prevents a Layer 2 loop from occurring while you are configuring the devices for MRP.
- Add an MRP ring to a port-based VLAN. When you add a ring, the CLI changes to the configuration level for the ring, where you can perform the following tasks.
 - Optionally, specify a name for the ring.
 - On the master node only, enable the device to be the master for the ring. Each ring can have only one master node.
 - Specify the MRP interfaces. Each device has two interfaces to an MRP ring.
 - Optionally, change the hello time and the preforwarding time. These parameters control how quickly failover occurs following a change in the state of a link in the ring.
 - Enable the ring.
- Optionally, add the ring's VLAN to a topology group to add more VLANs to the ring. If you use a topology group, make sure you configure MRP on the group's master VLAN. See "Topology Groups" on page 20-1.
- Re-enable the interface you disabled to prevent a Layer 2 loop. Once MRP is enabled, MRP will prevent the Layer 2 loop.

Adding an MRP Ring to a VLAN

To add an MRP ring to a VLAN, enter commands such as the following.

NOTE: If you plan to use a topology group to add VLANs to the ring, make sure you configure MRP on the topology group's master VLAN.

```
BigIron(config)# vlan 2
BigIron(config-vlan-2)# metro-ring 1
BigIron(config-vlan-2-mrp-1)# name CustomerA
BigIron(config-vlan-2-mrp-1)# master
BigIron(config-vlan-2-mrp-1)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
```

These commands configure an MRP ring on VLAN 2. The ring ID is 1, the ring name is CustomerA, and this node (this Foundry device) is the master for the ring. The ring interfaces are 1/1 and 1/2. Interface 1/1 is the primary interface and 1/2 is the secondary interface. The primary interface will initiate RHPs by default. The ring takes effect in VLAN 2.

To configure MRP rings with shared interfaces, enter commands such as the following:

```
BigIron(config)# vlan 2
BigIron(config-vlan-2)# metro-ring 1
BigIron(config-vlan-2-mrp-1)# name CustomerA
BigIron(config-vlan-2-mrp-1)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
BigIron(config-vlan-2-mrp-1)# metro-ring 2
BigIron(config-vlan-2-mrp-2)# name CustomerB
BigIron(config-vlan-2-mrp-2)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
```

Syntax: [no] metro-ring <ring-id>

The <ring-id> parameter specifies the ring ID and can be from 1 – 255. Configure the same ring ID on each of the nodes in the ring.

Syntax: [no] name <string>

The <string> parameter specifies a name for the ring. The name is optional, but it can be up to 20 characters long and can include blank spaces. If you use a name that has blank spaces, enclose the name in double quotation marks (for example: "Customer A").

Syntax: [no] master

Configures this node as the master node for the ring. Enter this command only on one node in the ring. The node is a member (non-master) node by default.

Syntax: [no] ring-interface ethernet | pos <primary-if> ethernet | pos <secondary-if>

The **ethernet | pos <primary-if>** parameter specifies the primary interface. On the master node, the primary interface is the one that originates RHPs. Ring control traffic and Layer 2 data traffic will flow in the outward direction from this interface by default. On member nodes, the direction of traffic flow depends on the traffic direction selected by the master node. Therefore, on a member node, the order in which you enter the interfaces does not matter.

The **ethernet | pos <secondary-if>** parameter specifies the secondary interface.

You can use two Ethernet interfaces, two POS interfaces, or a combination of Ethernet and POS.

NOTE: To take advantage of every interface in a Metro network, you can configure another MRP ring and either configure a different Master node for the ring or reverse the configuration of the primary and secondary interfaces on the Master node. Configuring multiple rings enables you to use all the ports in the ring. The same port can forward traffic one ring while blocking traffic for another ring.

Syntax: [no] enable

The **enable** command enables the ring.

Changing the Hello and PreForwarding Times

You also can change the RHP hello time and preforwarding time. To do so, enter commands such as the following:

```
BigIron(config-vlan-2-mrp-1)# hello-time 200
BigIron(config-vlan-2-mrp-1)# preforwarding-time 400
```

These commands change the hello time to 200 ms and change the preforwarding time to 400 ms.

NOTE: The preforwarding time must be at least twice the value of the hello time and must be a multiple of the hello time.

Syntax: [no] hello-time <ms>

Syntax: [no] preforwarding-time <ms>

The <ms> specifies the number of milliseconds. For the hello time, you can specify from 100 – 1000 (one second). The default hello time is 100 ms. The preforwarding time can be from 200 – 5000 ms, but must be at least twice the value of the hello time and must be a multiple of the hello time. The default preforwarding time is 300 ms. A change to the hello time or preforwarding time takes effect as soon as you enter the command.

NOTE: You can use MRP ring diagnostics to determine whether you need to change the hello time and preforwarding time. See "Using MRP Diagnostics".

Using MRP Diagnostics

The MRP diagnostics feature calculates how long it takes for RHP packets to travel through the ring. When you enable MRP diagnostics, the software tracks RHP packets according to their sequence numbers and calculates how long it takes an RHP packet to travel one time through the entire ring. When you display the diagnostics, the CLI shows the average round-trip time for the RHP packets sent since you enabled diagnostics. The calculated results have a granularity of 1 microsecond.

Enabling MRP Diagnostics

To enable MRP diagnostics for a ring, enter the following command on the Master node, at the configuration level for the ring:

```
BigIron(config-vlan-2-mrp-1)# diagnostics
```

Syntax: [no] diagnostics

NOTE: This command is valid only on the master node.

Displaying MRP Diagnostics

To display MRP diagnostics results, enter the following command on the Master node:

```
BigIron(config)# show metro 2 diag

Metro Ring 2 - CustomerA
=====
diagnostics results

Ring      Diag      RHP average      Recommended      Recommended
id        state     time(microsec)   hello time(ms)   Prefwing time(ms)
2         enabled   125              100              300

Diag frame sent      Diag frame lost
1230                 0
```

Syntax: show metro <ring-id> diag

This display shows the following information.

Table 20.2: CLI Display of MRP Ring Diagnostic Information

This Field...	Displays...
Ring id	The ring ID.
Diag state	The state of ring diagnostics.
RHP average time	The average round-trip time for an RHP packet on the ring. The calculated time has a granularity of 1 microsecond.
Recommended hello time	The hello time recommended by the software based on the RHP average round-trip time.
Recommended Prefwing time	The preforwarding time recommended by the software based on the RHP average round-trip time.
Diag frame sent	The number of diagnostic RHPs sent for the test.
Diag frame lost	The number of diagnostic RHPs lost during the test.

If the recommended hello time and preforwarding time are different from the actual settings and you want to change them, see “Configuring MRP” on page 20-15.

Displaying MRP Information

You can display the following MRP information:

- Topology group configuration information
- Ring configuration information and statistics

Displaying Topology Group Information

To display topology group information, enter the following command:

Syntax: show topology-group [<group-id>]

See “Displaying Topology Group Information” on page 20-3 for more information.

Displaying Ring Information

To display ring information, enter the following command:

```
BigIron(config)# show metro

Metro Ring 2
=====
Ring      State      Ring      Master    Topo      Hello      Prefwing
id        enabled   role      vlan      group     time(ms)   time(ms)
2         enabled   member    2         not conf  100        300

Ring interfaces      Interface role      Forwarding state      Active interface      Interface Type
ethernet 1/1         primary             disabled              none                  Regular
ethernet 1/2         secondary           forwarding            ethernet 2            Tunnel

RHPs sent      RHPs rcvd      TC RHPs rcvd      State changes
3              0              0                 4
```

Syntax: show metro [<ring-id>]

This display shows the following information.

Table 20.3: CLI Display of MRP Ring Information

This Field...	Displays...
Ring id	The ring ID
State	The state of MRP. The state can be one of the following: <ul style="list-style-type: none"> • enabled – MRP is enabled • disabled – MRP is disabled
Ring role	Whether this node is the master for the ring. The role can be one of the following: <ul style="list-style-type: none"> • master • member

Table 20.3: CLI Display of MRP Ring Information (Continued)

This Field...	Displays...
Master vlan	<p>The ID of the master VLAN in the topology group used by this ring. If a topology group is used by MRP, the master VLAN controls the MRP settings for all VLANs in the topology group.</p> <p>Note: The topology group ID is 0 if the MRP VLAN is not the master VLAN in a topology group. Using a topology group for MRP configuration is optional.</p>
Topo group	The topology group ID.
Hello time	The interval, in milliseconds, at which the Forwarding port on the ring's master node sends Ring Hello Packets (RHPs).
Prefwing time	<p>The number of milliseconds an MRP interface that has entered the Preforwarding state will wait before changing to the Forwarding state.</p> <p>If a member port in the Preforwarding state does not receive an RHP within the Preforwarding time (Prefwing time), the port assumes that a topology change has occurred and changes to the Forwarding state.</p> <p>The secondary port on the Master node changes to Blocking if it receives an RHP, but changes to Forwarding if the port does not receive an RHP before the preforwarding time expires.</p> <p>Note: A member node's Preforwarding interface also changes from Preforwarding to Forwarding if it receives an RHP whose forwarding bit is on.</p>
Ring interfaces	<p>The device's two interfaces with the ring.</p> <p>Note: If the interfaces are trunk groups, only the primary ports of the groups are listed.</p>
Interface role	<p>The interface role can be one of the following:</p> <ul style="list-style-type: none"> • primary <ul style="list-style-type: none"> • Master node – The interface generates RHPs. • Member node – The interface forwards RHPs received on the other interface (the secondary interface). • secondary – The interface does not generate RHPs. <ul style="list-style-type: none"> • Master node – The interface listens for RHPs. • Member node – The interface receives RHPs.
Forwarding state	<p>Whether MRP Forwarding is enabled on the interface. The forwarding state can be one of the following:</p> <ul style="list-style-type: none"> • blocking – The interface is blocking Layer 2 data traffic and RHPs • disabled – The interface is down • forwarding – The interface is forwarding Layer 2 data traffic and RHPs • preforwarding – The interface is listening for RHPs but is blocking Layer 2 data traffic

Table 20.3: CLI Display of MRP Ring Information (Continued)

This Field...	Displays...
Active interface	The physical interfaces that are sending and receiving RHPs. Note: If a port is disabled, its state is shown as “disabled”. Note: If an interface is a trunk group, only the primary port of the group is listed.
Interface Type	Shows if the interface is a regular port or a tunnel port. This field is available beginning with software release 07.7.00.
RHPs sent	The number of RHPs sent on the interface. Note: This field applies only to the master node. On non-master nodes, this field contains 0. This is because the RHPs are forwarded in hardware on the non-master nodes.
RHPs rcvd	The number of RHPs received on the interface. Note: On most Foundry devices, this field applies only to the master node. On non-master nodes, this field contains 0. This is because the RHPs are forwarded in hardware on the non-master nodes. However, on the FES and FES X-Series, the RHP received counter on non-master MRP nodes increment. This is because, on the FES and FES X-Series, the CPU receives a copy of the RHPs forwarded in hardware.
TC RHPs rcvd	The number of Topology Change RHPs received on the interface. A Topology Change RHP indicates that the ring topology has changed.
State changes	The number of MRP interface state changes that have occurred. The state can be one of the states listed in the Forwarding state field.

MRP CLI Example

The following examples show the CLI commands required to implement the MRP configuration shown in Figure 20.10 on page 20-14.

NOTE: For simplicity, the figure shows the VLANs on only two switches. The CLI examples implement the ring on all four switches.

Commands on Switch A (Master Node)

The following commands configure a VLAN for the ring. The ring VLAN must contain both of the node’s interfaces with the ring. Add these interfaces as tagged interfaces, since the interfaces also must be in each of the customer VLANs configured on the node.

```
BigIron(config)# vlan 2
BigIron(config-vlan-2)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-2)# metro-ring 1
BigIron(config-vlan-2-mrp-1)# name "Metro A"
BigIron(config-vlan-2-mrp-1)# master
BigIron(config-vlan-2-mrp-1)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
BigIron(config-vlan-2-mrp-1)# exit
BigIron(config-vlan-2)# exit
```

The following commands configure the customer VLANs. The customer VLANs must contain both the ring interfaces as well as the customer interfaces.

```

BigIron(config)# vlan 30
BigIron(config-vlan-30)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-30)# tag ethernet 2/1
BigIron(config-vlan-30)# exit
BigIron(config)# vlan 40
BigIron(config-vlan-40)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-40)# tag ethernet 4/1
BigIron(config-vlan-40)# exit

```

The following commands configure topology group 1 on VLAN 2. The master VLAN is the one that contains the MRP configuration. The member VLANs use the MRP parameters of the master VLAN. The control interfaces (the ones shared by the master VLAN and member VLAN) also share MRP state.

```

BigIron(config)# topology-group 1
BigIron(config-topo-group-1)# master-vlan 2
BigIron(config-topo-group-1)# member-vlan 30
BigIron(config-topo-group-1)# member-vlan 40

```

Commands on Switch B

The commands for configuring switches B, C, and D are similar to the commands for configuring switch A, with two differences: the nodes are not configured to be the ring master. Omitting the **master** command is required for non-master nodes.

```

BigIron(config)# vlan 2
BigIron(config-vlan-2)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-2)# metro-ring 1
BigIron(config-vlan-2-mrp-1)# name "Metro A"
BigIron(config-vlan-2-mrp-1)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
BigIron(config-vlan-2)# exit

```

```

BigIron(config)# vlan 30
BigIron(config-vlan-30)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-30)# tag ethernet 2/1
BigIron(config-vlan-30)# exit
BigIron(config)# vlan 40
BigIron(config-vlan-40)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-40)# tag ethernet 4/1
BigIron(config-vlan-40)# exit

```

```

BigIron(config)# topology-group 1
BigIron(config-topo-group-1)# master-vlan 2
BigIron(config-topo-group-1)# member-vlan 30
BigIron(config-topo-group-1)# member-vlan 40

```

Commands on Switch C

```

BigIron(config)# vlan 2
BigIron(config-vlan-2)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-2)# metro-ring 1
BigIron(config-vlan-2-mrp-1)# name "Metro A"
BigIron(config-vlan-2-mrp-1)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
BigIron(config-vlan-2)# exit

```

```

BigIron(config)# vlan 30
BigIron(config-vlan-30)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-30)# tag ethernet 2/1
BigIron(config-vlan-30)# exit
BigIron(config)# vlan 40
BigIron(config-vlan-40)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-40)# tag ethernet 4/1

```

```
BigIron(config-vlan-40)# exit
BigIron(config)# topology-group 1
BigIron(config-topo-group-1)# master-vlan 2
BigIron(config-topo-group-1)# member-vlan 30
BigIron(config-topo-group-1)# member-vlan 40
```

Commands on Switch D

```
BigIron(config)# vlan 2
BigIron(config-vlan-2)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-2)# metro-ring 1
BigIron(config-vlan-2-mrp-1)# name "Metro A"
BigIron(config-vlan-2-mrp-1)# ring-interface ethernet 1/1 ethernet 1/2
BigIron(config-vlan-2-mrp-1)# enable
BigIron(config-vlan-2)# exit

BigIron(config)# vlan 30
BigIron(config-vlan-30)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-30)# tag ethernet 2/1
BigIron(config-vlan-30)# exit
BigIron(config)# vlan 40
BigIron(config-vlan-40)# tag ethernet 1/1 to 1/2
BigIron(config-vlan-40)# tag ethernet 4/1
BigIron(config-vlan-40)# exit

BigIron(config)# topology-group 1
BigIron(config-topo-group-1)# master-vlan 2
BigIron(config-topo-group-1)# member-vlan 30
BigIron(config-topo-group-1)# member-vlan 40
```

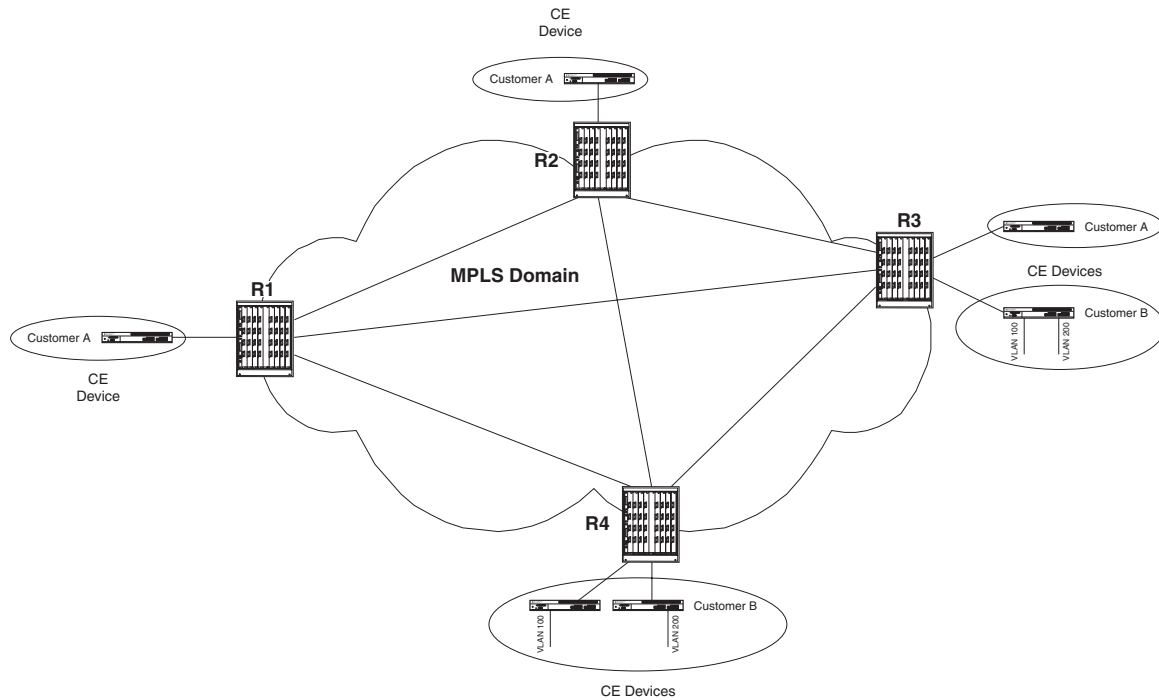
MRP and VPLS Internetworking

NOTE: This feature is supported on Service Provider Release 09.2.00.

MRP (Metro Ring Protocol) is Foundry's proprietary protocol to avoid Layer 2 loops in any type of ring topology and to achieve faster traffic recovery times. VPLS (virtual private LAN service) is an IETF standard that provides Layer 2 VPN over packet (MPLS/IP) backbone. Prior to this release, Layer 2 protocols (for example, MRP or STP) cannot be configured on a VPLS VLAN in a provider edge (PE) device. This release removes this restriction and adds the ability for the MRP protocol to work with VPLS. Ports controlled by MRP can affect VPLS VLAN configurations.

Furthermore:

- VPLS VLANs can now be configured as members of a master VLAN in a topology group.
- When a ring topology change occurs and a PE node needs to flush the MAC table for the master (MRP) VLAN, the MAC tables of the member VPLS VLAN need to be flushed as well.

Figure 21 Sample of VPLS configuration with MRP

Configuration Notes

Topology groups are used to allow MRP and VPLS to be configured on the same VLAN. The topology groups allow the interfaces controlled by the VLAN in an MRP configuration to also be controlled by the VLANs in a VPLS configuration. To allow MRP and VPLS to work together:

1. Configure a port VLAN and its member interfaces. Member interfaces of the VLAN should be tagged interfaces.
2. Configure VPLS and its VLAN. Then assign the interfaces assigned to the port VLAN to be members of the VPLS VLAN. Other interfaces in the PE device can also be members of the VPLS VLAN.
3. Configure a topology group with the port VLAN as the master VLAN. Add the VPLS VLANs as members of the topology group's master VLAN.
4. Configure MRP to run on the master VLAN.

The VPLS VLAN follows the MRP VLAN's state selection. If there is a loop in the MRP topology, traffic for the port VLAN in the MRP topology is blocked. Because of the topology group, traffic for the members of the VPLS VLAN is blocked also. Once the MRP master detects a break in the ring topology, its interface transitions to the forwarding state for the MRP VLAN as well as for the members of the VPLS VLAN.

The following assumptions and restrictions apply to the MRP-VPLS internetworking:

- For MRP-VPLS interworking topology group, only VPLS VLANs can be configured as member VLANs.
- The VPLS VLAN cannot be configured as the master.

When a VPLS VLAN is added as a member VLAN of a topology group that already has a regular VLAN as its member, the configuration is rejected. Similarly, when a regular VLAN is added as a member VLAN to a topology group that already has a VPLS VLAN as its member, the configuration is rejected also.

Adding the VPLS VLAN as a member of a topology group

You add regular a VLAN that is a member of a VPLS VLAN to the master VLAN of a topology group. For example, to add VLAN 4 as a member of VPLS VLAN 10, which is a member of topology group 1:

```
NetIron(config)# topology-group 1
NetIron(config-topo-group-1)# master-vlan 100
NetIron(config-topo-group-1)#member-vlan vpls 10 vlan 4
```

Syntax: member-vlan vpls <vpls-id> vlan <vlan-id>

The <vlan-id> parameter specifies a VLAN ID. The VLAN must already be configured

The following is the output of the **show topology-group** command when a VPLS VLAN is configured as a member VLAN.

Displaying VPLS VLAN membership

To determine which VLANs and VPLS VLANs are members of a topology group, enter the following command:

```
NetIron(config-topo-group-1)#show topology-group

Topology Group 1
=====
  master-vlan 10
  member-vlan VPLS 1 VLAN 100

  Common control ports          L2 protocol
  ethernet 3/1                  MRP
  ethernet 3/2                  MRP
  Per vlan free ports
  ethernet 3/5                   Vlan 10
  ethernet 4/3                   VPLS 1 VLAN 100
```

Configuration Example

The following example shows how to configure MRP and VPLs internetworking on PE4 device (192.168.1.104) in Figure 21.

```
NetIron(config)# vlan 100 by port
NetIron(config-vlan-100)# tagged ethernet 1/6 ethernet 13/1 to 13/8
NetIron(config)# vlan 100 by port
NetIron(config-vlan-100)# metro-ring 1
NetIron(config-vlan-2-mrp-1)# master
NetIron(config-vlan-2-mrp-1)# ring-interfaces ethernet 13/1 ethernet 1/6
NetIron(config-vlan-2-mrp-1)# enable

NetIron(config)#router mpls
NetIron(config-mpls)#policy
NetIron(config-mpls-policy)# traffic-eng ospf
NetIron(config-mpls-policy)# exit

NetIron(config-mpls)#mpls-interface pos 10/1 to 10/2
NetIron(config-mpls-interface)#exit

NetIron(config-mpls)# lsp 1
NetIron(config-mpls-lsp)# to 192.168.2.101
NetIron(config-mpls-lsp)# enable
NetIron(config-mpls-lsp)#exit

NetIron(config-mpls)# lsp 2
NetIron(config-mpls)# to 192.168.2.103
NetIron(config-mpls)# enable
NetIron(config-mpls)# exit

NetIron(config-mpls)#vpls test1 1
NetIron(config-mpls-vpls-test1)#vpls-peer 192.168.2.101 192.168.2.103
NetIron(config-mpls-vpls-test1)#vlan 2
NetIron(config-mpls-vpls-test1-vlan-2)#tag ethernet 1/6 ethernet 13/1
NetIron(config-mpls-vpls-test2-vlan2)#exit
NetIron(config-mpls-vpls-test1)#exit
```



```

NetIron(config-mpls)#vpls test2 2
NetIron(config-mpls-vpls-test2)#vpls-peer 192.168.2.101 192.168.2.103
NetIron(config-mpls-vpls-test2)#vlan 3
NetIron(config-mpls-vpls-test2-vlan-3)#tag ethernet 1/6 ethernet 13/1
NetIron(config-mpls-vpls-test2-vlan-3)#exit
NetIron(config-mpls-vpls-test2)#exit

NetIron(config-mpls)#vpls test3 3
NetIron(config-mpls-vpls-test3)#vpls-peer 192.168.2.101 192.168.2.103
NetIron(config-mpls-vpls-test3)#vlan 4
NetIron(config-mpls-vpls-test3-vlan-4)#tag ethernet 1/6 ethe 13/1
NetIron(config-mpls-vpls-test3-vlan-4)#exit
NetIron(config-mpls-vpls-test3)#exit

NetIron(config)#router ospf
NetIron(config-ospf-router)# area 0

NetIron(config)# interface loopback 1
NetIron(config-lbif-1)# ip address 5.5.5.5 255.255.255.255
NetIron(config-lbif-1)# ip ospf area 0
NetIron(config-lbif-1)# exit

NetIron(config)# interface ethernet 1/6
NetIron(config-if-e100-1/6)# enable
NetIron(config-if-e100-1/6)# exit

NetIron(config)# interface pos 10/1
NetIron(config-poif-10/1)# enable
NetIron(config-poif-10/1)# ip address 10.10.10.1 255.255.255.0
NetIron(config-poif-10/1)# ip ospf area 0
NetIron(config-lbif-1)# exit

NetIron(config)# interface pos 10/2
NetIron(config-poif-10/2)# enable
NetIron(config-poif-10/2)# ip address 30.30.30.1 255.255.255.0
NetIron(config-poif-10/2)# ip ospf area 0
NetIron(config-lbif-1)# exit

NetIron(config)# interface ethernet 13/1
NetIron(config-if-e100-13/1)# enable
NetIron(config-if-e100-13/1)# exit

NetIron(config)# topology-group 1
NetIron(config-topo-group-1)# master-vlan 100
NetIron(config-topo-group-1)# member-vlan vpls 1 vlan 2
NetIron(config-topo-group-1)# member-vlan vpls 2 vlan 3
NetIron(config-topo-group-1)# member-vlan vpls 3 vlan 4

NetIron(config-topo-group-1)# write terminal

```

Virtual Switch Redundancy Protocol (VSRP)

NOTE: This feature cannot be configured in the B2R Layer 3 SwitchH2R image in software release 07.6.01b. However, devices running version 07.6.01b of the B2RH2R image can still be VSRP-aware. (VSRP awareness is described in the following section.) Also, The FES devices running release 03.1.00 support full VSRP, as well as VSRP-awareness.

On the BigIron MG8 and NetIron 40G, VSRP is supported only on Layer 2.

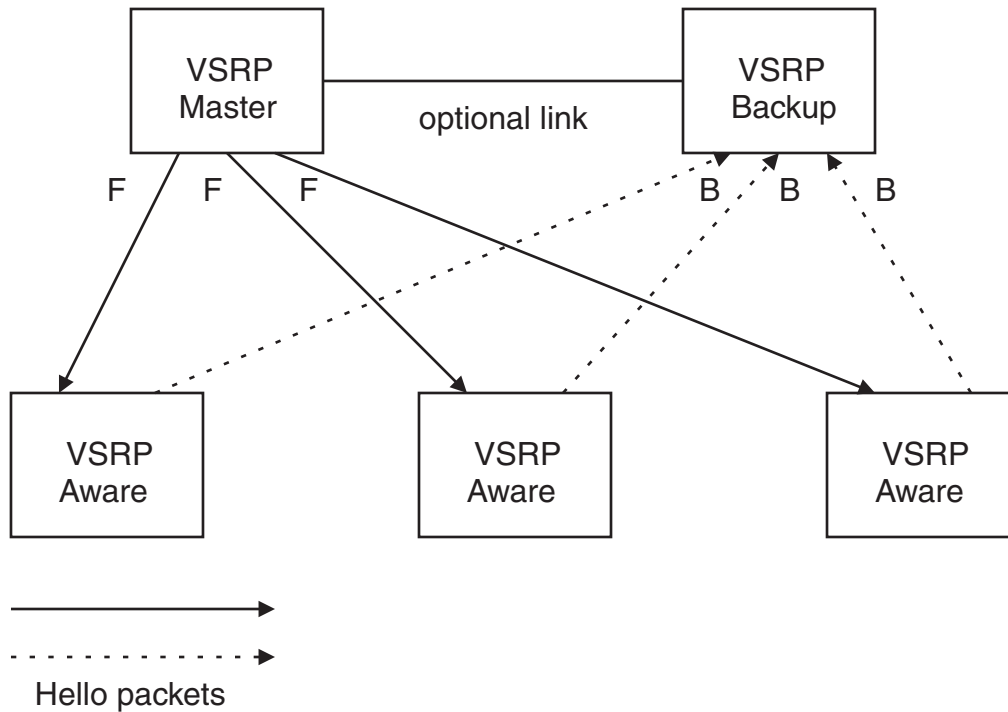
Virtual Switch Redundancy Protocol (VSRP) is a Foundry proprietary protocol that provides redundancy and sub-second failover in Layer 2 and Layer 3 mesh topologies. Based on the Foundry Virtual Router Redundancy Protocol Extended (VRRPE), VSRP provides one or more backups for a Layer 2 Switch or Layer 3 Switch. If the

active Layer 2 Switch or Layer 3 Switch becomes unavailable, one of the backups takes over as the active device and continues forwarding traffic for the network.

You can use VSRP for Layer 2, Layer 3, or for both layers. On Layer 3 Switches, Layer 2 and Layer 3 share the same VSRP configuration information. On Layer 2 Switches, VSRP applies only to Layer 2.

Figure 21.11 shows an example of a VSRP configuration.

Figure 21.11 VSRP mesh – redundant paths for Layer 2 and Layer 3 traffic



In this example, two Foundry devices are configured as redundant paths for VRID 1. On each of the devices, a Virtual Router ID (VRID) is configured on a port-based VLAN. Since VSRP is primarily a Layer 2 redundancy protocol, the VRID applies to the entire VLAN. However, you can selectively remove individual ports from the VRID if needed.

Following Master election (described below), one of the Foundry devices becomes the Master for the VRID and sets the state of all the VLAN's ports to Forwarding. The other device is a Backup and sets all the ports in its VRID VLAN to Blocking.

If a failover occurs, the Backup becomes the new Master and changes all its VRID ports to the Forwarding state.

Other Foundry devices can use the redundant paths provided by the VSRP devices. In this example, three Foundry devices use the redundant paths. A Foundry device that is not itself configured for VSRP but is connected to a Foundry device that is configured for VSRP, is **VSRP aware**. In this example, the three Foundry devices connected to the VSRP devices are VSRP aware. A Foundry device that is VSRP aware can failover its link to the new Master in sub-second time, by changing the MAC address associated with the redundant path.

When you configure VSRP, make sure each of the non-VSRP Foundry devices connected to the VSRP devices has a separate link to each of the VSRP devices.

NOTE: A Foundry device must be running software release 07.6.01 or later to be a VSRP device or a VSRP-aware device.

Layer 2 and Layer 3 Redundancy

You can configure VSRP to provide redundancy for Layer 2 only or also for Layer 3.

- Layer 2 only – The Layer 2 links are backup up but specific IP addresses are not backed up.
- Layer 2 and Layer 3 – The Layer 2 links are backup up and a specific IP address is also backed up. Layer 3 VSRP is the same as VRRPE. However, using VSRP provides redundancy at both layers at the same time.

Layer 2 Switches support Layer 2 VSRP only. Layer 3 Switches support Layer 2 and Layer 3 redundancy. You can configure a Layer 3 Switch for either Layer 2 only or Layer 2 and Layer 3. To configure for Layer 3, specify the IP address you are backing up.

NOTE: If you want to provide Layer 3 redundancy only, disable VSRP and use VRRPE.

Master Election and Failover

Each VSRP device advertises its VSRP priority in Hello messages. During Master election, the VSRP device with the highest priority for a given VRID becomes the Master for that VRID. After Master election, the Master sends Hello messages at regular intervals to inform the Backups that the Master is healthy.

If there is a tie for highest VSRP priority, the tie is resolved as follows:

- Layer 2 Switches – The Layer 2 Switch with the higher management IP address becomes the Master.
 - Switches with management IP addresses are preferred over switches without management IP addresses.
 - If neither of the switches has a management IP address, then the switch with the higher MAC address becomes the Master. (VSRP compares the MAC addresses of the ports configured for the VRID, not the base MAC addresses of the switches.)
- Layer 3 Switches – The Layer 3 Switch whose virtual routing interface has a higher IP address becomes the master.

VSRP Failover

Each Backup listens for Hello messages from the Master. The Hello messages indicate that the Master is still available. If the Backups stop receiving Hello messages from the Master, the election process occurs again and the Backup with the highest priority becomes the new Master.

Each Backup waits for a specific period of time, the Dead Interval, to receive a new Hello message from the Master. If the Backup does not receive a Hello message from the Master by the time the Dead Interval expires, the Backup sends a Hello message of its own, which includes the Backup's VSRP priority, to advertise the Backup's intent to become the Master. If there are multiple Backups for the VRID, each Backup sends a Hello message.

When a Backup sends a Hello message announcing its intent to become the Master, the Backup also starts a hold-down timer. During the hold-down time, the Backup listens for a Hello message with a higher priority than its own.

- If the Backup receives a Hello message with a higher priority than its own, the Backup resets its Dead Interval and returns to normal Backup status.
- If the Backup does not receive a Hello message with a higher priority than its own by the time the hold-down timer expires, the Backup becomes the new Master and starts forwarding Layer 2 traffic on all ports.

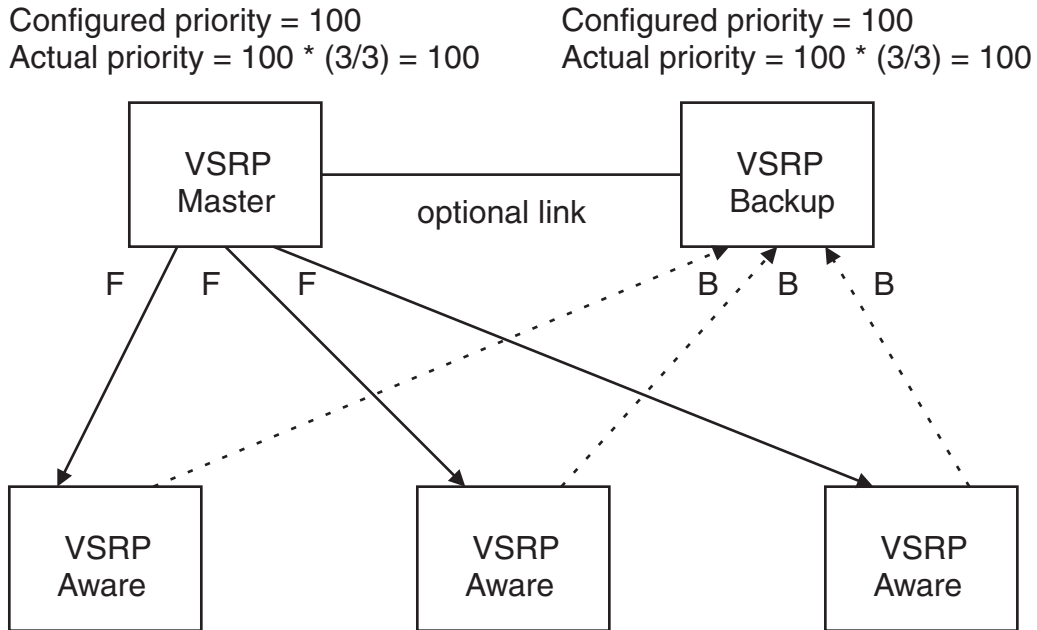
If you increase the timer scale value, each timer's value is divided by the scale value. To achieve sub-second failover times, you can change the scale to a value up to 10. This shortens all the VSRP timers to 10 percent of their configured values.

VSRP Priority Calculation

Each VSRP device has a VSRP priority for each VRID and its VLAN. The VRID is used during Master election for the VRID. By default, a device's VSRP priority is the value configured on the device (which is 100 by default). However, to ensure that a Backup with a high number of up ports for a given VRID is elected, the device reduces

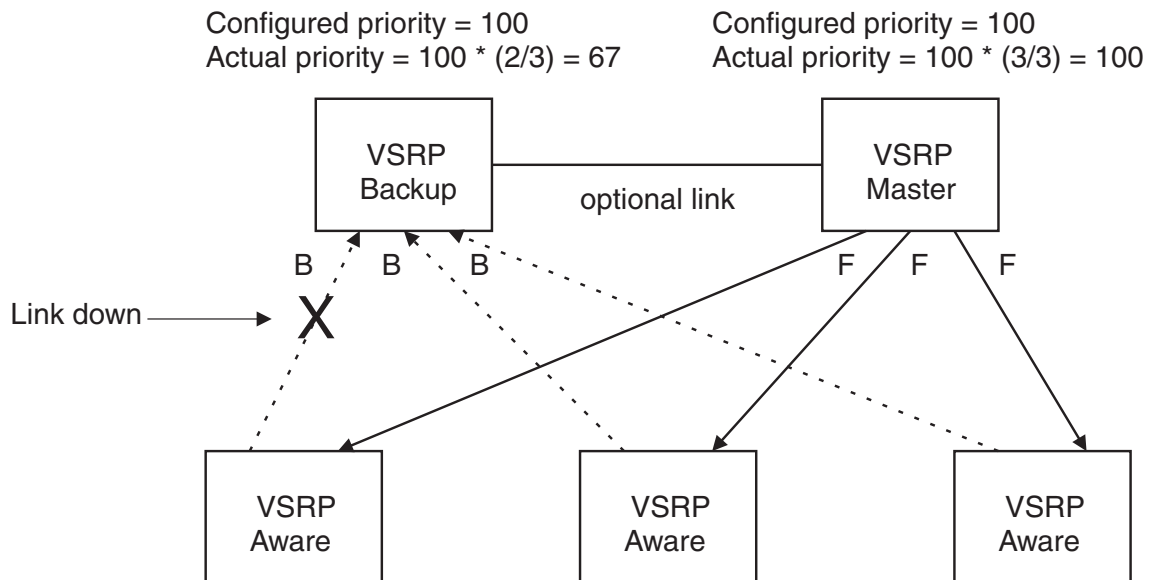
the priority if a port in the VRID's VLAN goes down. For example, if two Backups each have a configured priority of 100, and have three ports in VRID 1 in VLAN 10, each Backup begins with an equal priority, 100. This is shown in Figure 21.12

Figure 21.12 VSRP priority




However, if one of the VRID's ports goes down on one of the Backups, that Backup's priority is reduced. If the Master's priority is reduced enough to make the priority lower than a Backup's priority, the VRID fails over to the Backup. Figure 21.13 shows an example.

Figure 21.13 VSRP priority recalculation



You can reduce the sensitivity of a VSRP device to failover by increasing its configured VSRP priority. For example, you can increase the configured priority of the VSRP device on the left in Figure 21.13 to 150. In this case, failure of a single link does not cause failover. The link failure caused the priority to be reduced to 100, which is still equal to the priority of the other device. This is shown in Figure 21.14.

Figure 21.14 VSRP priority bias




Track Ports

Optionally, you can configure track ports to be included during VSRP priority calculation. In VSRP, a **track port** is a port that is not a member of the VRID's VLAN, but whose state is nonetheless considered when the priority is calculated. Typically, a track port represents the exit side of traffic received on the VRID ports. By default, no track ports are configured.

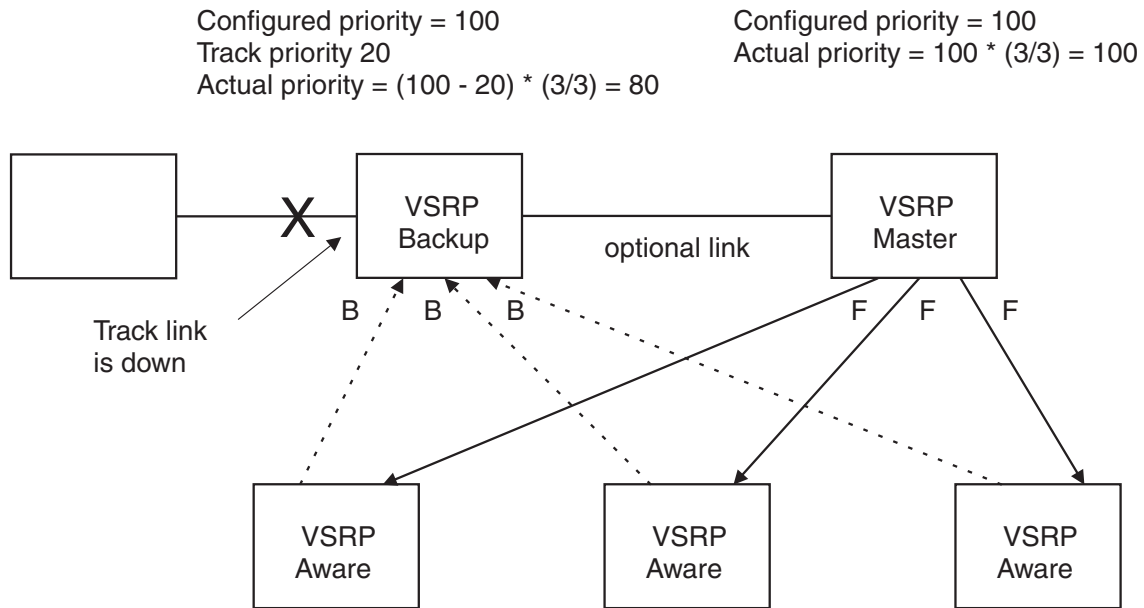
When you configure a track port, you assign a priority value to the port. If the port goes down, VSRP subtracts the track port's priority value from the configured VSRP priority. For example, if you configure a track port with priority 20 and the configured VSRP priority is 100, the software subtracts 20 from 100 if the track port goes down, resulting in a VSRP priority of 80. The new priority value is used when calculating the VSRP priority. Figure 21.15 shows an example.

Figure 21.15 Track port priority



In Figure 21.15, the track port is up. Since the port is up, the track priority does not affect the VSRP priority calculation. If the track port goes down, the track priority does affect VSRP priority calculation, as shown in Figure 21.16.

Figure 21.16 Track port priority subtracted during priority calculation



MAC Address Failover on VSRP-Aware Devices

VSRP-aware devices maintain a record of each VRID and its VLAN. When the device has received a Hello message for a VRID in a given VLAN, the device creates a record for that VRID and VLAN and includes the port number in the record. Each subsequent time the device receives a Hello message for the same VRID and VLAN, the device checks the port number.

- If the port number is the same as the port that previously received a Hello message, the VSRP-aware device assumes that the message came from the same VSRP Master that sent the previous message.
- If the port number does not match, the VSRP-aware device assumes that a VSRP failover has occurred to a new Master, and moves the MAC addresses learned on the previous port to the new port.

The VRID records age out if unused. This can occur if the VSRP-aware device becomes disconnected from the Master. The VSRP-aware device will wait for a Hello message for the period of time equal to the following:

$$\text{VRID Age} = \text{Dead Interval} + \text{Hold-down Interval} + (3 \times \text{Hello Interval})$$

The values for these timers are determined by the VSRP device sending the Hello messages. If the Master uses the default timer values, the age time for VRID records on the VSRP-aware devices is as follows:

$$3 + 2 + (3 \times 1) = 8 \text{ seconds}$$

In this case, if the VSRP-aware device does not receive a new Hello message for a VRID in a given VLAN, on any port, the device assumes the connection to the Master is unavailable and removes the VRID record.

Timer Scale

The VSRP Hello interval, Dead interval, Backup Hello interval, and Hold-down interval timers are individually configurable. You also can easily change all the timers at the same time while preserving the ratios among their values. To do so, change the timer scale. The **timer scale** is a value used by the software to calculate the timers. The software divides a timer's value by the timer scale value. By default, the scale is 1. This means the VSRP timer values are the same as the values in the configuration.

VSRP-Aware Security Features

Software release 07.6.04 and later enhances the security of VSRP-aware switches against unauthorized VSRP hello packets by enabling you to configure VSRP-aware security parameters.

Without VSRP-aware security configured, a VSRP-aware device passively learns the authentication method conveyed by the received VSRP hello packet. The VSRP-aware device then stores the authentication method until it ages out with the aware entry.

With VSRP-aware security, you can:

- Define the specific authentication parameters that a VSRP-aware device will use on a VSRP backup switch. The authentication parameters that you define will not age out.
- Define a list of ports that have authentic VSRP backup switch connections. For ports included in the list, the VSRP-aware switch will process VSRP hello packets using the VSRP-aware security configuration. Conversely, for ports not included in the list, the VSRP-aware switch will not use the VSRP-aware security configuration.

If VSRP hello packets do not meet the acceptance criteria, the VSRP-aware device forwards the packets normally, without any VSRP-aware security processing.

VSRP Parameters

Table 21.4 lists the VSRP parameters.

Table 21.4: VSRP Parameters

Parameter	Description	Default	See page...
Protocol	VSRP state Note: On a Layer 3 Switch, you must disable VSRP to use VRRPE or VRRP.	Enabled	20-35
Virtual Router ID (VRID)	The ID of the virtual switch you are creating by configuring multiple devices as redundant links. You must configure the same VRID on each device that you want to use to back up the links.	None	20-34
Timer scale	The value used by the software to calculate all VSRP timers. Increasing the timer scale value decreases the length of all the VSRP timers equally, without changing the ratio of one timer to another.	1	20-35

Interface Parameters

Table 21.4: VSRP Parameters (Continued)

Parameter	Description	Default	See page...
Authentication type	<p>The type of authentication the VSRP devices use to validate VSRP packets. On Layer 3 Switches, the authentication type must match the authentication type the VRID's port uses with other routing protocols such as OSPF.</p> <ul style="list-style-type: none"> No authentication – The interfaces do not use authentication. This is the VRRP default. Simple – The interface uses a simple text-string as a password in packets sent on the interface. If the interface uses simple password authentication, the VRID configured on the interface must use the same authentication type and the same password. <p>Note: MD5 is not supported.</p>	No authentication	20-36
VSRP-Aware Security Parameters			
VSRP-Aware Authentication type	<p>The type of authentication the VSRP-aware devices will use on a VSRP backup switch.</p> <ul style="list-style-type: none"> No authentication – The device does not accept incoming packets that have authentication strings. Simple – The device uses a simple text-string as the authentication string for accepting incoming packets. 	Not configured	20-36
VRID Parameters			
VSRP device type	<p>Whether the device is a VSRP Backup for the VRID. All VSRP devices for a given VRID are Backups.</p>	Not configured	20-34
VSRP ports	<p>The ports in the VRID's VLAN that you want to use as VRID interfaces. You can selectively exclude individual ports from VSRP while allowing them to remain in the VLAN.</p>	All ports in the VRID's VLAN	20-37
VRID IP address	<p>A gateway address you are backing up. Configuring an IP address provides VRRPE Layer 3 redundancy in addition to VSRP LAYer 2 redundancy.</p> <p>The VRID IP address must be in the same subnet as a real IP address configured on the VSRP interface, but cannot be the same as a real IP address configured on the interface.</p> <p>Note: This parameter is valid only on Layer 3 Switches.</p>	None	20-38

Table 21.4: VSRP Parameters (Continued)

Parameter	Description	Default	See page...
Backup priority	<p>A numeric value that determines a Backup's preferability for becoming the Master for the VRID. During negotiation, the device with the highest priority becomes the Master.</p> <p>In VSRP, all devices are Backups and have the same priority by default.</p> <p>If two or more Backups are tied with the highest priority, the Backup with the highest IP address becomes the Master for the VRID.</p>	100 for all Backups	20-38
Preference of timer source	<p>When you save a Backup's configuration, the software can save the configured VSRP timer values or the VSRP timer values received from the Master.</p> <p>Saving the current timer values instead of the configured ones helps ensure consistent timer usage for all the VRID's devices.</p> <p>Note: The Backup always gets its timer scale value from the Master.</p>	Configured timer values are saved	20-38
Time-to-Live (TTL)	The maximum number of hops a VSRP Hello packet can traverse before being dropped. You can specify from 1 – 255.	2	20-39
Hello interval	<p>The amount of time between Hello messages from the Master to the Backups for a given VRID.</p> <p>The interval can be from 1 – 84 seconds.</p>	One second	20-39
Dead interval	<p>The amount of time a Backup waits for a Hello message from the Master for the VRID before determining that the Master is no longer active.</p> <p>If the Master does not send a Hello message before the dead interval expires, the Backups negotiate (compare priorities) to select a new Master for the VRID.</p>	Three times the Hello Interval	20-39
Backup Hello state and interval	<p>The amount of time between Hello messages from a Backup to the Master.</p> <p>The message interval can be from 60 – 3600 seconds.</p> <p>You must enable the Backup to send the messages. The messages are disabled by default on Backups. The current Master sends Hello messages by default.</p>	<p>Disabled</p> <p>60 seconds when enabled</p>	20-40
Hold-down interval	<p>The amount of time a Backup that has sent a Hello packet announcing its intent to become Master waits before beginning to forward traffic for the VRID. The hold-down interval prevents Layer 2 loops from occurring during VSRP's rapid failover.</p> <p>The interval can from 1 – 84 seconds.</p>	2 seconds	20-40

Table 21.4: VSRP Parameters (Continued)

Parameter	Description	Default	See page...
Track priority	A VSRP priority value assigned to the tracked port(s). If a tracked port's link goes down, the VRID port's VSRP priority is reduced by the amount of the tracked port's priority.	5	20-40
Track port	A track port is a port or virtual routing interface that is outside the VRID but whose link state is tracked by the VRID. Typically, the tracked interface represents the other side of VRID traffic flow through the device. If the link for a tracked interface goes down, the VSRP priority of the VRID interface is changed, causing the devices to renegotiate for Master.	None	20-41
Backup preempt mode	Prevents a Backup with a higher VSRP priority from taking control of the VRID from another Backup that has a lower priority but has already assumed control of the VRID.	Enabled	20-41
VRID active state	The active state of the VSRP VRID.	Disabled	20-34
RIP Parameters			
Suppression of RIP advertisements	A Layer 3 Switch that is running RIP normally advertises routes to a backed up VRID even when the Layer 3 Switch is not currently the active Layer 3 Switch for the VRID. Suppression of these advertisements helps ensure that other Layer 3 Switches do not receive invalid route paths for the VRID. Note: This parameter is valid only on Layer 3 Switches.	Disabled (routes are advertised)	20-41

Configuring Basic VSRP Parameters Using the CLI

To configure VSRP, perform the following required tasks:

- Configure a port-based VLAN containing the ports for which you want to provide VSRP service.

NOTE: If you already have a port-based VLAN but only want to use VSRP on a sub-set of the VLANs ports, you can selectively remove ports from VSRP service in the VLAN. See “Removing a Port from the VRID’s VLAN” on page 20-37.

- Configure a VRID.
 - Specify that the device is a backup. Since VSRP, like VRRPE, does not have an “owner”, all VSRP devices are backups. The active device for a VRID is elected based on the VRID priority, which is configurable.
 - Activate the VRID.

The following example shows a simple VSRP configuration.

```
BigIron(config)# vlan 200
BigIron(config-vlan-200)# tag ethernet 1/1 to 1/8
```

```
BigIron(config-vlan-200)# vsrp vrid 1
BigIron(config-vlan-200-vrid-1)# backup
BigIron(config-vlan-200-vrid-1)# activate
```

Syntax: [no] vsrp vrid <num>

The <num> parameter specifies the VRID and can be from 1 – 255.

Syntax: [no] backup [priority <value>] [track-priority <value>]

This command is required. In VSRP, all devices on which a VRID are configured are Backups. The Master is then elected based on the VSRP priority of each device. There is no “owner” device as there is in VRRP.

For information about the command’s optional parameters, see the following:

- “Changing the Backup Priority” on page 20-38
- “Changing the Default Track Priority” on page 20-40

Syntax: [no] activate

or

Syntax: enable | disable

Configuring Optional VSRP Parameters Using the CLI

The following sections describe how to configure optional VSRP parameters.

Disabling or Re-Enabling VSRP

VSRP is enabled by default on Layer 2 Switches and Layer 3 Switches. On a Layer 3 Switch, if you want to use VRRP or VRRPE for Layer 3 redundancy instead of VSRP, you need to disable VSRP first. To do so, enter the following command at the global CONFIG level:

NOTE: This command is not available on the BigIron MG8 or NetIron 40G.

```
BigIron(config)# no router vsrp
router vsrp is disabled. All vsrp config data will be lost when writing to flash
```

To re-enable the protocol, enter the following command:

```
BigIron(config)# router vsrp
```

Syntax: [no] router vsrp

Since VRRP and VRRPE do not apply to Layer 2 Switches, there is no need to disable VSRP and there is no command to do so. The protocol is always enabled.

Changing the Timer Scale

To achieve sub-second failover times, you can shorten the duration of all VSRP timers by adjusting the timer scale. The **timer scale** is a value used by the software to calculate the timers. By default, the scale value is 1. If you increase the timer scale, each timer’s value is divided by the scale value. Using the timer scale to adjust VSRP timer values enables you to easily change all the timers while preserving the ratios among their values. Here is an example.

Timer	Timer Scale	Timer Value
Hello interval	1	1 second
	2	0.5 seconds

Timer	Timer Scale	Timer Value
Dead interval	1	3 seconds
	2	1.5 seconds
Backup Hello interval	1	60 seconds
	2	30 seconds
Hold-down interval	1	2 seconds
	2	1 second

If you configure the device to receive its timer values from the Master, the Backup also receives the timer scale value from the Master.

NOTE: The Backups always use the value of the timer scale received from the Master, regardless of whether the timer values that are saved in the configuration are the values configured on the Backup or the values received from the Master.

To change the timer scale, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# scale-timer 2
```

This command changes the scale to 2. All VSRP timer values will be divided by 2.

Syntax: [no] scale-timer <num>

The <num> parameter specifies the multiplier. You can specify a timer scale from 1 – 10.

Configuring Authentication

If the interfaces on which you configure the VRID use authentication, the VSRP packets on those interfaces also must use the same authentication. VSRP supports the following authentication types:

- No authentication – The interfaces do not use authentication. This is the default.
- Simple – The interfaces use a simple text-string as a password in packets sent on the interface. If the interfaces use simple password authentication, the VRID configured on the interfaces must use the same authentication type and the same password.

To configure a simple password, enter a command such as the following at the interface configuration level:

```
BigIron(config-if-1/6)# ip vsrp auth-type simple-text-auth ourpword
```

This command configures the simple text password “ourpword”.

Syntax: [no] ip vsrp auth-type no-auth | simple-text-auth <auth-data>

The **auth-type no-auth** parameter indicates that the VRID and the interface it is configured on do not use authentication.

The **auth-type simple-text-auth <auth-data>** parameter indicates that the VRID and the interface it is configured on use a simple text password for authentication. The <auth-data> value is the password. If you use this parameter, make sure all interfaces on all the devices supporting this VRID are configured for simple password authentication and use the same password.

Configuring Security Features on a VSRP-Aware Device

NOTE: On Enterprise software releases, this feature is available in software releases 07.6.04 and later.

The VSRP-aware security feature enables you to:

- Define the specific authentication parameters that a VSRP-aware device will use on a VSRP backup switch. The authentication parameters that you define will not age out.

- Define a list of ports that have authentic VSRP backup switch connections. For ports included in the list, the VSRP-aware switch will process VSRP hello packets using the VSRP-aware security configuration. Conversely, for ports not included in the list, the VSRP-aware switch will not use the VSRP-aware security configuration.

If VSRP hello packets do not meet the acceptance criteria, the VSRP-aware device forwards the packets normally, without any VSRP-aware security processing.

Specifying an Authentication String for VSRP Hello Packets

The following configuration defines **pri-key** as the authentication string for accepting incoming VSRP hello packets. In this example, the VSRP-aware device will accept all incoming packets that have this authorization string.

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# vsrp-aware vrid 3 simple-text-auth pri-key
```

Syntax: vsrp-aware vrid <vrid number> simple text auth <string>

Specifying no Authentication for VSRP Hello Packets

The following configuration specifies no authentication as the preferred VSRP-aware security method. In this case, the VSRP device will not accept incoming packets that have authentication strings.

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# vsrp-aware vrid 2 no-auth
```

Syntax: vsrp-aware vrid <vrid number> no-auth

The following configuration specifies no authentication for VSRP hello packets received on ports 1/1, 1/2, 1/3, and 1/4 in VRID 4. For these ports, the VSRP device will not accept incoming packets that have authentication strings.

```
BigIron(config)# vlan 10
BigIron(config-vlan-10)# vsrp-aware vrid 4 no-auth port-list ethe 1/1 to 1/4
```

Syntax: vsrp-aware vrid <vrid number> no-auth port-list <port range>

<vrid number> is a valid VRID (from 1 to 255).

no-auth specifies no authentication as the preferred VSRP-aware security method. The VSRP device will not accept incoming packets that have authentication strings.

simple-text-auth <string> specifies the authentication string for accepting VSRP hello packets, where <string> can be up to 8 characters.

port-list <port range> specifies the range of ports to include in the configuration.

Removing a Port from the VRID's VLAN

By default, all the ports in the VLAN on which you configure a VRID are interfaces for the VRID. You can remove a port from the VRID while allowing it to remain in the VLAN.

Removing a port is useful in the following cases:

- There is no risk of a loop occurring, such as when the port is attached directly to an end host.
- You plan to use a port in an MRP ring.

To remove a port from a VRID, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# no include-port ethernet 1/2
```

Syntax: [no] include-port ethernet | pos <portnum>

The **ethernet | pos** <portnum> parameter specifies the port you are removing from the VRID. The port remains in the VLAN but its forwarding state is not controlled by VSRP.

Configuring a VRID IP Address

If you are configuring a Layer 3 Switch for VSRP, you can specify an IP address to back up. When you specify an IP address, VSRP provides redundancy for the address. This is useful if you want to back up the gateway address used by hosts attached to the VSRP Backups.

VSRP does not require you to specify an IP address. If you do not specify an address, VSRP provides Layer 2 redundancy. If you do specify an address, VSRP provides Layer 2 and Layer 3 redundancy.

The Layer 3 redundancy support is the same as VRRPE support. For information, see the “Configuring VRRP and VRRPE” chapter in the *Foundry Enterprise Configuration and Management Guide*.

NOTE: The VRID IP address must be in the same subnet as a real IP address configured on the VSRP interface, but cannot be the same as a real IP address configured on the interface.

NOTE: Failover applies to both Layer 2 and Layer 3.

To specify an IP address to back up, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# ip-address 10.10.10.1
```

Syntax: [no] ip-address <ip-addr>

or

Syntax: [no] ip address <ip-addr>

Changing the Backup Priority

When you enter the backup command to configure the device as a VSRP Backup for the VRID, you also can change the backup priority and the track priority.

- The backup priority is used for election of the Master. The VSRP Backup with the highest priority value for the VRID is elected as the Master for that VRID. The default priority is 100. If two or more Backups are tied with the highest priority, the Backup with the highest IP address becomes the Master for the VRID.
- The track priority is used with the track port feature. See “VSRP Priority Calculation” on page 20-27 and “Changing the Default Track Priority” on page 20-40.

To change the backup priority, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# backup priority 75
```

Syntax: [no] backup [priority <value>] [track-priority <value>]

The **priority** <value> parameter specifies the VRRP priority for this interface and VRID. You can specify a value from 3 – 254. The default is 100.

For a description of the **track-priority** <value> parameter, see “Changing the Default Track Priority” on page 20-40.

Saving the Timer Values Received from the Master

The Hello messages sent by a VRID’s master contain the VRID values for the following VSRP timers:

- Hello interval
- Dead interval
- Backup Hello interval
- Hold-down interval

By default, each Backup saves the configured timer values to its startup-config file when you save the device’s configuration.

You can configure a Backup to instead save the current timer values received from the Master when you save the configuration. Saving the current timer values instead of the configured ones helps ensure consistent timer usage for all the VRID's devices.

NOTE: The Backups always use the value of the timer scale received from the Master, regardless of whether the timer values that are saved in the configuration are the values configured on the Backup or the values received from the Master.

To configure a Backup to save the VSRP timer values received from the Master instead of the timer values configured on the Backup, enter the following command:

```
BigIron(config-vlan-200-vrid-1)# save-current-values
```

Syntax: [no] save-current-values

Changing the Time-To-Live (TTL)

A VSRP Hello packet's TTL specifies how many hops the packet can traverse before being dropped. A hop can be a Layer 3 Switch or a Layer 2 Switch. You can specify from 1 – 255. The default TTL is 2. When a VSRP device (Master or Backup) sends a VSRP Hello packet, the device subtracts one from the TTL. Thus, if the TTL is 2, the device that originates the Hello packet sends it out with a TTL of 1. Each subsequent device that receives the packet also subtracts one from the packet's TTL. When the packet has a TTL of 1, the receiving device subtracts 1 and then drops the packet because the TTL is zero.

NOTE: An MRP ring is considered to be a single hop, regardless of the number of nodes in the ring.

To change the TTL for a VRID, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# initial-ttl 5
```

Syntax: [no] initial-ttl <num>

The <num> parameter specifies the TTL and can be from 1 – 255. The default TTL is 2.

Changing the Hello Interval

The Master periodically sends Hello messages to the Backups. To change the Hello interval, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# hello-interval 10
```

Syntax: [no] hello-interval <num>

The <num> parameter specifies the interval and can be from 1 – 84 seconds. The default is 1 second.

NOTE: The default Dead interval is three times the Hello interval plus one-half second. Generally, if you change the Hello interval, you also should change the Dead interval on the Backups.

NOTE: If you change the timer scale, the change affects the actual number of seconds.

Changing the Dead Interval

The Dead interval is the number of seconds a Backup waits for a Hello message from the Master before determining that the Master is dead. The default is 3 seconds. This is three times the default Hello interval.

To change the Dead interval, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# dead-interval 30
```

Syntax: [no] dead-interval <num>

The <num> parameter specifies the interval and can be from 1 – 84 seconds. The default is 3 seconds.

NOTE: If you change the timer scale, the change affects the actual number of seconds.

Changing the Backup Hello State and Interval

By default, Backups do not send Hello messages to advertise themselves to the Master. You can enable these messages if desired and also change the message interval.

To enable a Backup to send Hello messages to the Master, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# advertise backup
```

Syntax: [no] advertise backup

When a Backup is enabled to send Hello messages, the Backup sends a Hello message to the Master every 60 seconds by default. You can change the interval to be up to 3600 seconds.

To change the Backup Hello interval, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# backup-hello-interval 180
```

Syntax: [no] backup-hello-interval <num>

The <num> parameter specifies the message interval and can be from 60 – 3600 seconds. The default is 60 seconds.

NOTE: If you change the timer scale, the change affects the actual number of seconds.

Changing the Hold-Down Interval

The hold-down interval prevents Layer 2 loops from occurring during failover, by delaying the new Master from forwarding traffic long enough to ensure that the failed Master is really unavailable.

To change the Hold-down interval, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# hold-down-interval 4
```

Syntax: [no] hold-down-interval <num>

The <num> parameter specifies the hold-down interval and can be from 1 – 84 seconds. The default is 2 seconds.

NOTE: If you change the timer scale, the change affects the actual number of seconds.

Changing the Default Track Priority

When you configure a VRID to track the link state of other interfaces, if one of the tracked interface goes down, the software changes the VSRP priority of the VRID interface.

The software reduces the VRID priority by the amount of the priority of the tracked interface that went down. For example, if the VSRP interface's priority is 100 and a tracked interface with track priority 60 goes down, the software changes the VSRP interface's priority to 40. If another tracked interface goes down, the software reduces the VRID's priority again, by the amount of the tracked interface's track priority.

The default track priority for all track ports is 1. You can change the default track priority or override the default for an individual track port.

- To change the default track priority, use the **backup track-priority** command, described below.
- To override the default track priority for a specific track port, use the **track-port** command. See "Specifying a Track Port" on page 20-41.

To change the track priority, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# backup track-priority 2
```

Syntax: [no] backup [priority <value>] [track-priority <value>]

Specifying a Track Port

You can configure the VRID on one interface to track the link state of another interface on the device. This capability is useful for tracking the state of the exit interface for the path for which the VRID is providing redundancy. See “VSRP Priority Calculation” on page 20-27.

To configure a VRID to track an interface, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# track-port e 2/4
```

Syntax: [no] track-port ethernet <portnum> | pos <portnum> | ve <num> [priority <num>]

The **priority** <num> parameter changes the VSRP priority of the interface. If this interface goes down, the VRID's VSRP priority is reduced by the amount of the track port priority you specify here.

NOTE: The priority <num> option changes the priority of the specified interface, overriding the default track port priority. To change the default track port priority, use the **backup track-priority** <num> command.

Disabling or Re-Enabling Backup Pre-Emption

By default, a Backup that has a higher priority than another Backup that has become the Master can preempt the Master, and take over the role of Master. If you want to prevent this behavior, disable preemption.

Preemption applies only to Backups and takes effect only when the Master has failed and a Backup has assumed ownership of the VRID. The feature prevents a Backup with a higher priority from taking over as Master from another Backup that has a lower priority but has already become the Master of the VRID.

Preemption is especially useful for preventing flapping in situations where there are multiple Backups and a Backup with a lower priority than another Backup has assumed ownership, because the Backup with the higher priority was unavailable when ownership changed.

If you enable the non-preempt mode (thus disabling the preemption feature) on all the Backups, the Backup that becomes the Master following the disappearance of the Master continues to be the Master. The new Master is not preempted.

To disable preemption on a Backup, enter a command such as the following at the configuration level for the VRID:

```
BigIron(config-vlan-200-vrid-1)# non-preempt-mode
```

Syntax: [no] non-preempt-mode

Suppressing RIP Advertisement from Backups

Normally, for Layer 3 a VSRP Backup includes route information for a backed up IP address in RIP advertisements. As a result, other Layer 3 Switches receive multiple paths for the backed up interface and might sometimes unsuccessfully use the path to the Backup rather than the path to the Master.

You can prevent the Backups from advertising route information for the backed up interface by enabling suppression of the advertisements.

NOTE: This parameter applies only if you specified an IP address to back up and is valid only on Layer 3 Switches.

To suppress RIP advertisements, enter the following commands:

```
Router2(config)# router rip
Router2(config-rip-router)# use-vrrp-path
```

Syntax: [no] use-vrrp-path

Displaying VSRP Information Using the CLI

You can display the following VSRP information:

- Configuration information and current parameter values for a VRID or VLAN
- The interfaces on a VSRP-aware device that are active for the VRID

Displaying VRID Information

To display VSRP information, enter the following command:

On most devices:

```
BigIron(config-vlan-200-vrid-1)# show vsrp vrid 1
Total number of VSRP routers defined: 2
VLAN 200
auth-type no authentication
VRID 1
State      Administrative-status  Advertise-backup  Preempt-mode  save-current
standby    enabled                disabled          true          false

Parameter      Configured Current      Unit
priority        100      80      (100-0)*(4.0/5.0)
hello-interval  1        1        sec/1
dead-interval   3        3        sec/1
hold-interval   3        3        sec/1
initial-ttl     2        2        hops

next hello sent in 00:00:00.8
Member ports:    ethe 1/1 to 1/5
Operational ports: ethe 1/1 to 1/4
Forwarding ports: ethe 1/1 to 1/4
```

On NetIron 40G devices:

```
NetIron 40G# show vsrp vrid 10
VLAN 10
auth-type no authentication
VRID 10
=====
State      Administrative-status  Advertise-backup  Preempt-mode  save-current
standby    enabled                disabled          true          false

Parameter      Configured Current      Unit/Formula
priority        100      50      (100-0)*(1.0/2.0)
hello-interval  1        1        sec/10
dead-interval   3        3        sec/10
hold-interval   3        3        sec/10
initial-ttl     2        2        hops

master router 219.130.154.186 expires in 00:00:00.5
Member ports:    ethe 1/1 to 1/2
Operational ports: ethe 1/1
Forwarding ports:  None
```

On a devices where the VSRP Fast Start feature is enabled:

```
NetIron(config-vlan-100-vrid-100)#show vsrp vrid 100
VLAN 100
  auth-type no authentication
  VRID 100
  =====
  State      Administrative-status Advertise-backup Preempt-mode  save-current
  master     enabled              disabled        true          false
  Parameter  Configured Current      Unit/Formula
  priority   100      50          (100-0)*(2.0/4.0)
  hello-interval 1      1          sec/1
  dead-interval 3      3          sec/1
  hold-interval 3      3          sec/1
  initial-ttl 2      2          hops

  next hello sent in 00:00:00.3
  Member ports:      ethe 2/5 to 2/8
  Operational ports: ethe 2/5 ethe 2/8
  Forwarding ports:  ethe 2/5 ethe 2/8
  Restart ports:     2/5(1) 2/6(1) 2/7(1) 2/8(1)
```

Syntax: show vsrp [vrid <num> | vlan <vlan-id>]

This display shows the following information when you use the **vrid** <num> or **vlan** <vlan-id> parameter. For information about the display when you use the **aware** parameter, see “Displaying the Active Interfaces for a VRID” on page 20-45.

Table 21.5: CLI Display of VSRP VRID or VLAN Information

This Field...	Displays...
Total number of VSRP routers defined	The total number of VRIDs configured on this device.
VLAN	The VLAN on which VSRP is configured.
auth-type	The authentication type in effect on the ports in the VSRP VLAN.
VRID parameters	
VRID	The VRID for which the following information is displayed.
state	<p>This device's VSRP state for the VRID. The state can be one of the following:</p> <ul style="list-style-type: none"> initialize – The VRID is not enabled (activated). If the state remains “initialize” after you activate the VRID, make sure that the VRID is also configured on the other routers and that the routers can communicate with each other. <p>Note: If the state is “initialize” and the mode is incomplete, make sure you have specified the IP address for the VRID.</p> <ul style="list-style-type: none"> standby – This device is a Backup for the VRID. master – This device is the Master for the VRID.

Table 21.5: CLI Display of VSRP VRID or VLAN Information (Continued)

This Field...	Displays...
Administrative-status	<p>The administrative status of the VRID. The administrative status can be one of the following:</p> <ul style="list-style-type: none"> disabled – The VRID is configured on the interface but VSRP or VRRPE has not been activated on the interface. enabled – VSRP has been activated on the interface.
Advertise-backup	<p>Whether the device is enabled to send VSRP Hello messages when it is a Backup. This field can have one of the following values:</p> <ul style="list-style-type: none"> disabled – The device does not send Hello messages when it is a Backup. enabled – The device does send Hello messages when it is a Backup.
Preempt-mode	<p>Whether the device can be pre-empted by a device with a higher VSRP priority after this device becomes the Master. This field can have one of the following values:</p> <ul style="list-style-type: none"> disabled – The device cannot be pre-empted. enabled – The device can be pre-empted.
save-current	<p>The source of VSRP timer values preferred when you save the configuration. This field can have one of the following values:</p> <ul style="list-style-type: none"> false – The timer values configured on this device are saved. true – The timer values most recently received from the Master are saved instead of the locally configured values.
<p>Note: For the following fields:</p>	
<ul style="list-style-type: none"> Configured – indicates the parameter value configured on this device. Current – indicates the parameter value received from the Master. Unit – indicates the formula used for calculating the VSRP priority and the timer scales in effect for the VSRP timers. A timer's true value is the value listed in the Configured or Current field divided by the scale value. 	
priority	<p>The device's preferability for becoming the Master for the VRID. During negotiation, the Backup with the highest priority becomes the Master.</p> <p>If two or more Backups are tied with the highest priority, the Backup interface with the highest IP address becomes the Master for the VRID.</p>
hello-interval	<p>The number of seconds between Hello messages from the Master to the Backups for a given VRID.</p>

Table 21.5: CLI Display of VSRP VRID or VLAN Information (Continued)

This Field...	Displays...
dead-interval	<p>The configured value for the dead interval. The dead interval is the number of seconds a Backup waits for a Hello message from the Master for the VRID before determining that the Master is no longer active.</p> <p>If the Master does not send a Hello message before the dead interval expires, the Backups negotiate (compare priorities) to select a new Master for the VRID.</p> <p>Note: If the value is 0, then you have not configured this parameter.</p>
hold-interval	<p>The number of seconds a Backup that intends to become the Master will wait before actually beginning to forward Layer 2 traffic for the VRID.</p> <p>If the Backup receives a Hello message with a higher priority than its own before the hold-down interval expires, the Backup remains in the Backup state and does not become the new Master.</p>
initial-ttl	<p>The number of hops a Hello message can traverse after leaving the device before the Hello message is dropped.</p> <p>Note: An MRP ring counts as one hop, regardless of the number of nodes in the ring.</p>
next hello sent in	<p>The amount of time until the Master's dead interval expires. If the Backup does not receive a Hello message from the Master by the time the interval expires, either the IP address listed for the Master will change to the IP address of the new Master, or this Layer 3 Switch itself will become the Master.</p> <p>Note: This field applies only when this device is a Backup.</p>
master router (NetTron 40G only)	<p>The IP address of the master router.</p> <p>Note: If VSRP is configured on a virtual routing interface of a NetTron 40G, the device gives its IP address as the decimal expression of the last four bytes of the MAC address it uses for virtual routing interfaces. Thus, in a show interface command, the MAC address of the virtual routing interface is displayed. This address translates to the IP address of the master router on a show vrid display.</p>
Member ports	The ports in the VRID.
Operational ports	The member ports that are currently up.
Forwarding ports	The member ports that are currently in the Forwarding state. Ports that are forwarding on the Master are listed. Ports on the Standby, which are in the Blocking state, are not listed.

Displaying the Active Interfaces for a VRID

On a VSRP-aware device, you can display VLAN and port information for the connections to the VSRP devices (Master and Backups).

To display the active VRID interfaces, enter the following command on the VSRP-aware device:

```
BigIron(config-vlan-200-vrid-1)# show vsrp aware
```

```
Aware port listing
VLAN ID  VRID  Last Port
100      1      3/2
200      2      4/1
```

Syntax: show vsrp aware

This display shows the following information when you use the **aware** parameter. For information about the display when you use the **vrid** <num> or **vlan** <vlan-id> parameter, see “Displaying VRID Information” on page 20-42.

Table 21.6: CLI Display of VSRP-Aware Information

This Field...	Displays...
VLAN ID	The VLAN that contains the VSRP-aware device’s connection with the VSRP Master and Backups.
VRID	The VRID.
Last Port	The most recent active port connection to the VRID. This is the port connected to the current Master. If a failover occurs, the VSRP-aware device changes the port to the port connected to the new Master. The VSRP-aware device uses this port to send and receive data through the backed up node.

Configuring VSRP Using the Web Management Interface

You can use the Web management interface to configure VSRP.

Enabling and Disabling VSRP

To enable VRRP, do the following:

1. Log on to the device using a valid user name and password for read-write access. The General System configuration panel is displayed.
2. Click Enable next to VSRP to enable it, or Disable to disable it.
3. Click the Apply button to apply your changes.
4. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Configuring VSRP Parameters

To configure VSRP parameters using the Web management interface, do the following:

1. Log on to the device using a valid user name and password for read-write access. The General System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of configuration options.
3. Click on the plus sign next to VSRP in the tree view to expand the list of configuration options.
4. Click on the Virtual Switch link.
 - If virtual switches have been configured for the device, you see a list of virtual routers. Click the Modify button if you want to make changes to a virtual switch’s parameters, or click the Add virtual switch link to

add a VSRP interface.

- If virtual switches have not been configured, you see the VSRP configuration panel:

VSRP

VlanId:	<input type="text" value="1"/>
VRId:	<input type="text" value="1"/>
Activate:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Hello Interval:	<input type="text" value="1"/>
Mode:	Backup
Priority:	<input type="text" value="100"/>
Backup mode only	
Backup Hello Interval:	<input type="text" value="60"/>
Dead Interval:	<input type="text" value="0"/>
Advertise Backup:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Preempt:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Hold Down Interval:	<input type="text" value="3"/>
Initial TTL:	<input type="text" value="1"/>
Router Save:	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Track	
Track priority:	<input type="text" value="5"/>

Track Ports

1/1 <input type="checkbox"/>	1/2 <input type="checkbox"/>	1/3 <input type="checkbox"/>	1/4 <input type="checkbox"/>	1/5 <input type="checkbox"/>	1/6 <input type="checkbox"/>	1/7 <input type="checkbox"/>	1/8 <input type="checkbox"/>
3/1 <input type="checkbox"/>	3/2 <input type="checkbox"/>	3/3 <input type="checkbox"/>	3/4 <input type="checkbox"/>	3/5 <input type="checkbox"/>	3/6 <input type="checkbox"/>	3/7 <input type="checkbox"/>	3/8 <input type="checkbox"/>
3/9 <input type="checkbox"/>	3/10 <input type="checkbox"/>	3/11 <input type="checkbox"/>	3/12 <input type="checkbox"/>	3/13 <input type="checkbox"/>	3/14 <input type="checkbox"/>	3/15 <input type="checkbox"/>	3/16 <input type="checkbox"/>
3/17 <input type="checkbox"/>	3/18 <input type="checkbox"/>	3/19 <input type="checkbox"/>	3/20 <input type="checkbox"/>	3/21 <input type="checkbox"/>	3/22 <input type="checkbox"/>	3/23 <input type="checkbox"/>	3/24 <input type="checkbox"/>
7/1 <input type="checkbox"/>	7/2 <input type="checkbox"/>	7/3 <input type="checkbox"/>	7/4 <input type="checkbox"/>	7/5 <input type="checkbox"/>	7/6 <input type="checkbox"/>	7/7 <input type="checkbox"/>	7/8 <input type="checkbox"/>
7/9 <input type="checkbox"/>	7/10 <input type="checkbox"/>	7/11 <input type="checkbox"/>	7/12 <input type="checkbox"/>	7/13 <input type="checkbox"/>	7/14 <input type="checkbox"/>	7/15 <input type="checkbox"/>	7/16 <input type="checkbox"/>
7/17 <input type="checkbox"/>	7/18 <input type="checkbox"/>	7/19 <input type="checkbox"/>	7/20 <input type="checkbox"/>	7/21 <input type="checkbox"/>	7/22 <input type="checkbox"/>	7/23 <input type="checkbox"/>	7/24 <input type="checkbox"/>

[Virtual Switch][Interface]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

- Enter the ID of the VLAN to which the VRID will be assigned in the VlanId field.

NOTE: The VLAN you enter must be configured and must be active. STP must also be disabled on the VLAN.

- Enter the VRID. By default, VRID 1 is assigned to an interface.
- By default, VSRP is disabled. Click Enable next to Activate to enable it on the VRID.
- Enter the amount of time between Hello messages from the Master to the Backups for a given VRID. The interval can be from 1 – 84 seconds. The default is 1 second.
- Backup is always displayed for the Mode field for VSRP.
- In the Hello Interval field, enter a number that determines a Backup's preferability for becoming the Master for the VRID. During negotiation, the device with the highest priority becomes the Master.

11. Backups is the only mode for all VSRP switches are
12. Enter a value for Priority. If two or more Backups are tied with the highest priority, the Backup with the highest IP address becomes the Master for the VRID. The default Backup Priority is 100.
13. Enter a value for Backup Hello Interval. This interval is the number of seconds between Hello messages from the Master to the Backups for a given VRID. The interval can from 60 –3600 seconds, with 60 seconds as the default.

You must enable the Backup to send messages (advertise backup).
14. In the Dead Interval field. enter he amount of time a Backup waits for a Hello message from the Master for the VRID before determining that the Master is no longer active.

If the Master does not send a Hello message before the dead interval expires, the Backups negotiate (compare priorities) to select a new Master for the VRID.
15. Select Enable for Advertise Backup if you want to advertise routes to a backed up VRID even when the Layer 3 Switch is not the current active router for the VRID. Disabling the advertisements helps ensure that other routers do not receive invalid route paths for the VRID. The default is Disabled.
16. Select Enable for the Preempt field to prevent a Backup with a higher VSRP priority from taking control of the VRID from another Backup that has a lower priority but has already assumed control of the VRID. Select Disable if you do not want to disable this feature. The default is enabled.
17. Enter a value for the Hold Down Interval field. This is the amount of time a Backup that has sent a Hello packet announcing its intent to become Master waits before beginning to forward traffic for the VRID. The hold-down interval prevents Layer 2 loops from occurring during VSRP's rapid failover. The interval can from 1 – 84 seconds. The default is 2 seconds.
18. Indicate the maximum time-to-live value, which is the number of hops a VSRP Hello packet can traverse before being dropped. You can specify from 1 – 255. The default is 2.
19. Click Enable for Router Save if you want the Backup to save the VSRP timer values received from the Master instead of the timer values configured on the Backup (above). VSRP timer values that will be saved are:
 - Hello interval
 - Dead interval
 - Backup Hello interval
 - Hold-down interval
20. Enter the Track Priority value or leave it blank to use the default. If a tracked port's link goes down, the VRID port's VSRP priority is reduced by the amount of the tracked port's priority. The default priority is 5.
21. In the Track Ports section, place a check mark in the box for a port whose link state is tracked by the VRID. Typically, the tracked interface represents the other side of VRID traffic flow through the device.

If the link for a tracked interface goes down, the VSRP priority of the VRID interface is changed, causing the devices to renegotiate for Master.
22. Click the Add button to add the VSRP switch.
23. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Modifying Authentication Parameters

You can modify the password that was configured for a VSRP interface on a separate panel of the Web management interface.

1. Log on to the device using a valid user name and password for read-write access. The General System configuration panel is displayed.
2. Click on the plus sign next to Monitor in the tree view to expand the list of configuration options.
3. Click on the plus sign next to VSRP in the tree view to expand the list of configuration options.

- Click on the [Interface](#) link to display the VSRP Interface table, which lists all the VSRP interfaces on the device that have been configured.

VSRP Interface				
Vlan Id	Vrid	Authentication Type	Simple Text Password	
1	1	no authentication		Modify
Vlan Id	Vrid	Authentication Type	Simple Text Password	

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

- Click the Modify button for the interface that you want to configure to display the VSRP Interface configuration panel.

VSRP Interface	
Authentication Type:	<input checked="" type="radio"/> None <input type="radio"/> Simple Text Password
Simple Text Password:	<input type="text"/>
<input type="button" value="Apply"/> <input type="button" value="Reset"/>	
<input type="button" value="[Show]"/>	
[Home] [Site Map] [Logout] [Save] [Frame Enable] [Disable] [TELNET]	

- Select the Authentication Type, either None, Simple Text Password or Ip Auth header.
- Enter a password if the authentication is Simple Text Password. Leave this field blank if other password types are used.
- Click the Apply button to update the information for the VSRP.
- Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Displaying VSRP Statistics Using the Web Management Interface

To display VSRP statistics using the Web management interface:

- Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
- Click on the plus sign next to Monitor in the tree view to expand the list of configuration options.
- Click on the plus sign next to VSRP in the tree view to expand the list of configuration options.

- Click on the [Virtual Switch](#) link to display the VSRP Virtual Switch Statistics Display panel.

VSRP Virtual Switch Statistics												
Vlan Id	VR Id	State	Receive Pkts Drop		Receive Mismatch			Rcv Priority	Rcv Higher	Transition Count		
			Arp	IP	Port	IP	Hello	Zero from Master	Priority	Master	Backup	
1	1	Initialize	0	0	0	0	0	0	0	0	0	
Vlan Id	VR Id	State	Receive Pkts Drop	Receive Mismatch	Rcv Priority	Rcv Higher	Transition Count					
			Arp	IP	Port	IP	Hello	Zero from Master	Priority	Master	Backup	

Clear

[Home](#)
[Site Map](#)
[Logout](#)
[Save](#)
[Frame Enable](#)
[Disable](#)
[TELNET](#)

The panel shows the following information:

Table 21.7: Web Management Interface Display of VSRP Statistics

This Field...	Displays...
VLAN ID	ID of the VLAN used by the virtual switch.
VRId	The VRID for the virtual switch.
State	Current state of the port. It can be: <ul style="list-style-type: none"> Initialize Master Backup
Receive Pkts Drop	Number of packets addressed to the VRID that were dropped. Packets are divided into the following categories: <ul style="list-style-type: none"> ARP packets IP packets
Receive Mismatch	Number of packets that did not match the configured values of the following: <ul style="list-style-type: none"> Port – receiving interface IP – IP addresses Hello – Hello interval
Receive Priority Zero from Master	Number of times the current Master has resigned
Receive Higher Priority	The number of VRRPE packets received by the interface that had a higher backup priority for the VRID than this device’s backup priority for the VRID.
Transition Count	The number of times this device has changed the state of its VRID: <ul style="list-style-type: none"> Master – transition from Backup to Master Backup – transition Master to Backup

To clear the statistics for VSRP, click the Clear button on the display panel.

VSRP Fast Start

Service Provider software release 09.1.00 and later and BigIron MG8 and NetIron 40G software release 02.1.00 and later provide the VSRP fast start feature.

The feature allows non-Foundry or non-VSRP aware devices that are connected to a Foundry device that is the VSRP Master to quickly switch over to the new Master when a VSRP failover occurs

This feature causes the port on a VSRP Master to restart when a VSRP failover occurs. When the port shuts down at the start of the restart, ports on the non-VSRP aware devices that are connected to the VSRP Master flush the MAC address they have learned for the VSRP master. After a specified time, the port on the previous VSRP Master (which now becomes the Backup) returns back online. Ports on the non-VSRP aware devices switch over to the new Master and learn its MAC address.

Special Considerations when Configuring VSRP Fast Start

- VSRP is sensitive to port status. When a port goes down, the VSRP instance lowers its priority based on the port up fraction. (see “VSRP Priority Calculation” on page 20-27 for more information on how priority is changed by port status). Since the VSRP fast start feature toggles port status by bringing ports down and up it can affect VSRP instances because their priorities get reduced when a port goes down. To avoid this, the VSRP fast start implementation keeps track of ports that it brings down and suppresses port down events for these ports (as concerns VSRP).
- Once a VSRP restart port is brought up by a VSRP instance, other VSRP instances (in Master state) that have this port as a member do not go to forwarding immediately. This is a safety measure that is required to prevent transitory loops. This could happen if a peer VSRP node gets completely cut off from this node and assumed Master state. In this case, where there are 2 VSRP instances that are in Master state and forwarding, the port comes up and starts forwarding immediately. This would cause a forwarding loop. To avoid this, the VSRP instance delays forwarding.

Recommendations for Configuring VSRP Fast Start

The following recommendations apply to configurations where multiple VSRP instances are running between peer devices sharing the same set of ports.

- Multiple VSRP instances configured on the same ports can cause VSRP instances to be completely cut off from peer VSRP instances. This can cause VSRP instances to toggle back and forth between master and backup mode. For this reason, we recommend that you configure VSRP fast start on a per port basis rather than for the entire VLAN.
- We recommend that VSRP peers have a directly connected port without VSRP fast start enabled on it. This allows protocol control packets to be received and sent even if other ports between the master and standby are down.
- The VSRP restart time should be configured based on the type of connecting device since some devices can take a long time to bring a port up or down (as long as several seconds). In order to ensure that the port restart is registered by neighboring device, the restart time may need to be changed to a value higher than the default value of 1 second.

Configuring VSRP Fast Start

The VSRP fast start feature can be enabled on a VSRP-configured Foundry device, either on the VLAN to which the VRID of the VSRP-configured device belongs (globally) or on a port that belongs to the VRID.

To globally configure a VSRP-configured device to shut down its ports when a failover occurs, then restart after five seconds, enter the following command:

```
BigIron MG8(configure)# vlan 100
BigIron MG8(configure-vlan-100)# vsrp vrid 1
BigIron MG8(configure-vlan-100-vrid-1)# restart-ports 5
```

Syntax: [no] restart-ports <seconds>

This command shuts down all the ports that belong to the VLAN when a failover occurs. All the ports will have the specified VRID.

To configure a single port on a VSRP-configured device to shut down when a failover occurs, then restart after a period of time, enter the following command:

```
BigIron MG8(configure)# interface ethernet 1/1
BigIron MG8(configure-if-1/1)# vsrp restart-port 5
```

Syntax: [no] vsrp restart-port <seconds>

In both commands, the <seconds> parameter instructs the VSRP Master to shut down its port for the specified number of seconds before it starts back up. Enter a value between 1 – 120 seconds. The default is 1 second.

Displaying Ports that Have VSRP Fast Start Feature Enabled

The **show vsrp vrid** command shows the ports on which the VSRP fast start feature is enabled.

```
BigIron MG8(config-vlan-100-vrid-100)#show vsrp vrid 100

VLAN 100
  auth-type no authentication
  VRID 100
  =====
  State      Administrative-status  Advertise-backup  Preempt-mode  save-current
  master     enabled                disabled          true          false
  Parameter  Configured  Current  Unit/Formula
  priority   100         50      (100-0)*(2.0/4.0)
  hello-interval  1           1       sec/1
  dead-interval  3           3       sec/1
  hold-interval  3           3       sec/1
  initial-ttl   2           2       hops

  next hello sent in 00:00:00.3
  Member ports:      ethe 2/5 to 2/8
  Operational ports: ethe 2/5 ethe 2/8
  Forwarding ports:  ethe 2/5 ethe 2/8
  Restart ports:     2/5(1) 2/6(1) 2/7(1) 2/8(1)
```


The "Restart ports:" line lists the ports that have the VSRP fast start enabled, and the downtime for each port. See Table 21.5 on page 20-43 to interpret the remaining information on the display.

VSRP and MRP Signaling

This feature is available on devices running Service Provider software release 09.1.00 and on BigIron MG8 and NetIron 40G devices running software release 02.1.00.


A device may connect to an MRP ring via VSRP to provide a redundant path between the device and the MRP ring. VSRP and MRP signaling, ensures rapid failover by flushing MAC addresses appropriately. The host on the MRP ring learns the MAC addresses of all devices on the MRP ring and VSRP link. From these MAC addresses, the host creates a MAC database (table), which is used to establish a data path from the host to a VSRP-linked device. Figure 21.17 below shows two possible data paths from the host to Device 1.

Figure 21.17 Two data paths from host on an MRP ring to a VSRP-linked device



If a VSRP failover from master to backup occurs, VSRP needs to inform MRP of the topology change; otherwise, data from the host continues along the obsolete learned path and never reach the VSRP-linked device, as shown in Figure 21.18.

Figure 21.18 VSRP on MRP rings that failed over




To ensure that MRP is informed of the topology change and to achieve convergence rapidly, this release provides a new signaling process for the interaction between VSRP and MRP. When a VSRP node fails, a new VSRP master is selected. The new VSRP master finds all MRP instances impacted by the failover. Then each MRP instance does the following:

- The MRP node sends out an MRP PDU with the mac-flush flag set three times on the MRP ring.
- The MRP node that receives this MRP PDU empties all the MAC entries from its interfaces that participate on the MRP ring.
- The MRP node then forwards the MRP PDU with the mac-flush flag set to the next MRP node that is in forwarding state.

The process continues until the Master MRP node's secondary (blocking) interface blocks the packet. Once the MAC address entries have been flushed, the MAC table can be rebuilt for the new path from the host to the VSRP-linked device (Figure 21.19).

Figure 21.19 New path established



There are no CLI commands used to configure this process.

Chapter 22

Configuring Virtual LANs (VLANs)

This chapter describes how to configure Virtual LANs (VLANs) on Foundry Layer 2 Switches and Layer 3 Switches.

The “Overview” section provides basic information about Foundry’s VLAN options. Following this section, other sections provide configuration procedures and examples.

To display configuration information for VLANs, see “Displaying VLAN Information” on page 22-90.

For complete syntax information for the CLI commands shown in this chapter, see the *Foundry Switch and Router Command Line Interface Reference*.

Most of the configuration examples in this chapter are based on CLI commands. For Web management procedures, see “Configuring VLANs Using the Web Management Interface” on page 22-83.

NOTE: For information about the GARP VLAN Registration Protocol (GVRP), see “Configuring GARP VLAN Registration Protocol (GVRP)” on page 25-1.

Overview

This section describes the Foundry VLAN features. Configuration procedures and examples appear in later sections of this chapter.

Types of VLANs

You can configure the following types of VLANs on Foundry devices.

- Layer 2 port-based VLAN – a set of physical ports that share a common, exclusive Layer 2 broadcast domain
- Layer 3 protocol VLANs – a subset of ports within a port-based VLAN that share a common, exclusive broadcast domain for Layer 3 broadcasts of the specified protocol type
- Protocol-Based VLANs (BigIron MG8 and NetIron 40G software release 02.0.00 and later)
- IP subnet VLANs – a subset of ports in a port-based VLAN that share a common, exclusive subnet broadcast domain for a specified IP subnet
- IPv6 VLANs – a subset of ports in a port-based VLAN that share a common, exclusive network broadcast domain for IPv6 packets
- IPX network VLANs – a subset of ports in a port-based VLAN that share a common, exclusive network broadcast domain for a specified IPX network
- AppleTalk cable VLANs – a subset of ports in a port-based VLAN that share a common, exclusive network

broadcast domain for a specified AppleTalk cable range

When a Foundry device receives a packet on a port that is a member of a VLAN, the device forwards the packet based on the following VLAN hierarchy:

- If the port belongs to an IP subnet VLAN, IPX network VLAN, or AppleTalk cable VLAN and the packet belongs to the corresponding IP subnet, IPX network, or AppleTalk cable range, the device forwards the packet to all the ports within that VLAN.
- If the packet is a Layer 3 packet but cannot be forwarded as described above, but the port is a member of a Layer 3 protocol VLAN for the packet's protocol, the device forwards the packet on all the Layer 3 protocol VLAN's ports.
- If the packet cannot be forwarded based on either of the VLAN membership types listed above, but the packet can be forwarded at Layer 2, the device forwards the packet on all the ports within the receiving port's port-based VLAN.

Protocol VLANs differ from IP subnet, IPX network, and AppleTalk VLANs in an important way. Protocol VLANs accept any broadcast of the specified protocol type. An IP subnet, IPx network, or AppleTalk VLAN accepts only broadcasts for the specified IP subnet, IPX network, or AppleTalk cable range.

NOTE: Protocol VLANs are different from IP subnet, IPX network, and AppleTalk cable VLANs. A port-based VLAN cannot contain both an IP subnet, IPX network, or AppleTalk cable VLAN and a protocol VLAN for the same protocol. For example, a port-based VLAN cannot contain both an IP protocol VLAN and an IP subnet VLAN.

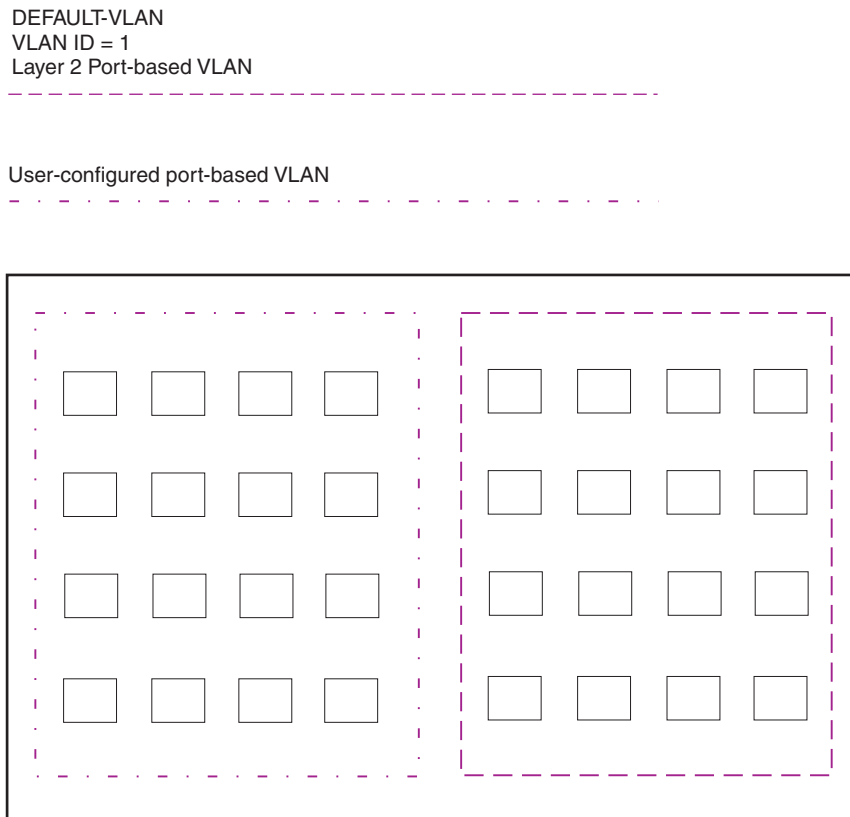
Layer 2 Port-Based VLANs

On all Foundry devices, you can configure port-based VLANs. A port-based VLAN is a subset of ports on a Foundry device that constitutes a Layer 2 broadcast domain.

By default, all the ports on a Foundry device are members of the default VLAN. Thus, all the ports on the device constitute a single Layer 2 broadcast domain. You can configure multiple port-based VLANs. When you configure a port-based VLAN, the device automatically removes the ports you add to the VLAN from the default VLAN.

Figure 22.1 shows an example of a Foundry device on which a Layer 2 port-based VLAN has been configured.

Figure 22.1 Foundry device containing user-defined Layer 2 port-based VLAN



When you add a port-based VLAN, the device removes all the ports in the new VLAN from DEFAULT-VLAN.

A port can belong to only one port-based VLAN, unless you apply 802.1q tagging to the port. **802.1q tagging** allows the port to add a four-byte tag field, which contains the VLAN ID, to each packet sent on the port. You also can configure port-based VLANs that span multiple devices by tagging the ports within the VLAN. The tag enables each device that receives the packet to determine the VLAN the packet belongs to. 802.1q tagging applies only to Layer 2 VLANs, not to Layer 3 VLANs.

Since each port-based VLAN is a separate Layer 2 broadcast domain, by default each VLAN runs a separate instance of the Spanning Tree Protocol (STP).

Layer 2 traffic is bridged within a port-based VLAN and Layer 2 broadcasts are sent to all the ports within the VLAN.

Layer 3 Protocol-Based VLANs

If you want some or all of the ports within a port-based VLAN to be organized according to Layer 3 protocol, you must configure a Layer 3 protocol-based VLAN within the port-based VLAN.

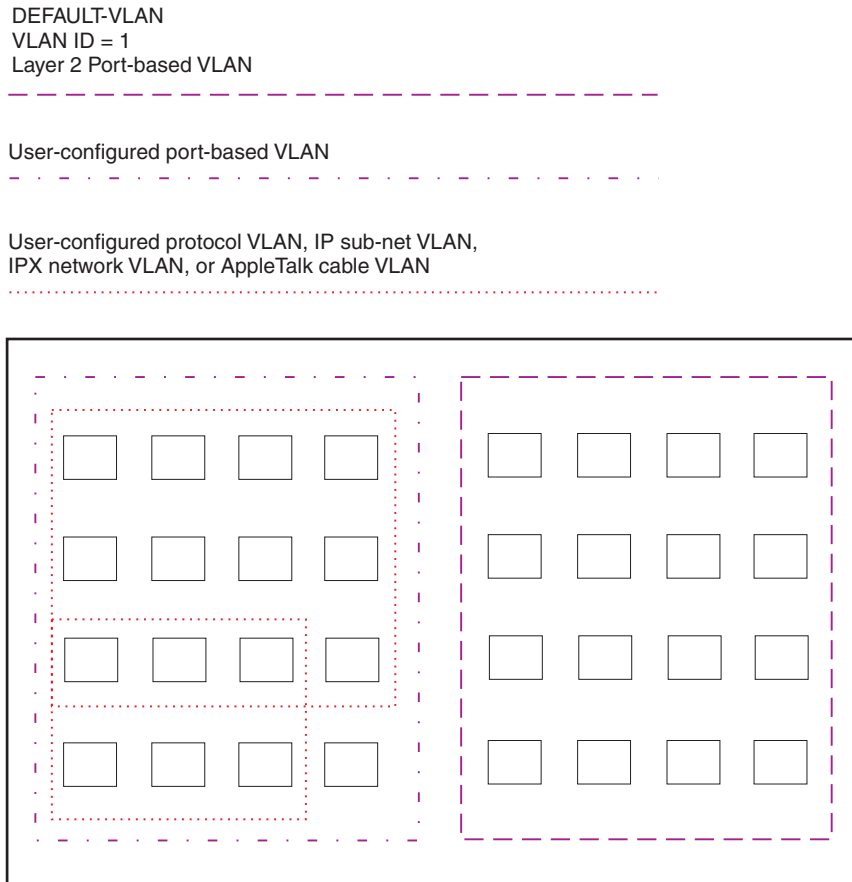
You can configure each of the following types of protocol-based VLAN within a port-based VLAN. All the ports in the Layer 3 VLAN must be in the same Layer 2 VLAN.

- AppleTalk – The device sends AppleTalk broadcasts to all ports within the AppleTalk protocol VLAN.
- IP – The device sends IP broadcasts to all ports within the IP protocol VLAN.

- IPv6 – The device sends IPv6 broadcasts to all ports within the IPv6 protocol VLAN.
- IPX – The device sends IPX broadcasts to all ports within the IPX protocol VLAN.
- DECnet – The device sends DECnet broadcasts to all ports within the DECnet protocol VLAN.
- NetBIOS – The device sends NetBIOS broadcasts to all ports within the NetBIOS protocol VLAN.
- Other – The device sends broadcasts for all protocol types other than those listed above to all ports within the VLAN.

Figure 22.2 shows an example of Layer 3 protocol VLANs configured within a Layer 2 port-based VLAN.

Figure 22.2 Layer 3 protocol VLANs within a Layer 2 port-based VLAN



You can add Layer 3 protocol VLANs or IP sub-net, IPX network, and AppleTalk cable VLANs to port-based VLANs.

Layer 3 VLANs cannot span Layer 2 port-based VLANs.

However, Layer 3 VLANs can overlap within a Layer 2 port-based VLAN.

Integrated Switch Routing (ISR)

Foundry Networks' *Integrated Switch Routing (ISR)* feature enables VLANs configured on Layer 3 Switches to route Layer 3 traffic from one protocol VLAN or IP subnet, IPX network, or AppleTalk cable VLAN to another. Normally, to route traffic from one IP subnet, IPX network, or AppleTalk cable VLAN to another, you would need to

forward the traffic to an external router. The VLANs provide Layer 3 broadcast domains for these protocols but do not in themselves provide routing services for these protocols. This is true even if the source and destination IP subnets, IPX networks, or AppleTalk cable ranges are on the same device.

ISR eliminates the need for an external router by allowing you to route between VLANs using virtual routing interfaces (vris). A **virtual routing interface** is a logical port on which you can configure Layer 3 routing parameters. You configure a separate virtual routing interface on each VLAN that you want to be able to route from or to. For example, if you configure two IP subnet VLANs on a Layer 3 Switch, you can configure a virtual routing interface on each VLAN, then configure IP routing parameters for the subnets. Thus, the Layer 3 Switch forwards IP subnet broadcasts within each VLAN at Layer 2 but routes Layer 3 traffic between the VLANs using the virtual routing interfaces.

NOTE: The Layer 3 Switch uses the lowest MAC address on the device (the MAC address of port 1 or 1/1) as the MAC address for all ports within all virtual routing interfaces you configure on the device.

The routing parameters and the syntax for configuring them are the same as when you configure a physical interface for routing. The logical interface allows the Layer 3 Switch to internally route traffic between the protocol-based VLANs without using physical interfaces.

All the ports within a protocol-based VLAN must be in the same port-based VLAN. The protocol-based VLAN cannot have ports in multiple port-based VLANs, unless the ports in the port-based VLAN to which you add the protocol-based VLAN are 802.1q tagged.

You can configure multiple protocol-based VLANs within the same port-based VLAN. In addition, a port within a port-based VLAN can belong to multiple protocol-based VLANs of the same type or different types. For example, if you have a port-based VLAN that contains ports 1 – 10, you can configure port 5 as a member of an AppleTalk protocol VLAN, an IP protocol VLAN, and an IPX protocol VLAN, and so on.

Protocol-Based VLANs (BigIron MG8 and NetIron 40G Software Release 02.0.00 and later)

Protocol-based VLANs provide the ability to define separate broadcast domains for several unique Layer 3 protocols within a single Layer 2 broadcast domain. Some applications for this feature might include security between departments with unique protocol requirements. This feature enables you to limit the amount of broadcast traffic to end-stations, servers, and routers.

Terathon IronWare software release 02.0.02 and later provides support for the following protocol-based protocols:

- AppleTalk – The device sends AppleTalk broadcasts to all ports within the AppleTalk protocol VLAN.
- IPv4 – The device sends IPv4 broadcasts to all ports within the IP protocol VLAN.
- IPv6 – The device sends IPv6 broadcasts to all ports within the IPv6 protocol VLAN.
- IPX – The device sends IPX broadcasts to all ports within the IPX protocol VLAN.
- Other – For all other protocols that have not been configured as protocol-VLANs under this VLAN.

Protocol-based VLANs can have the following membership types:

- Static ports – Static ports are permanent members of the protocol-based VLAN and remain active members of the VLAN regardless of whether the ports receive traffic for the VLAN's protocol.
- Exclude ports – Prevents a port in a port-based VLAN from ever becoming a member of a protocol-based VLAN.

IP Subnet, IPX Network, and AppleTalk Cable VLANs

The protocol-based VLANs described in the previous section provide separate protocol broadcast domains for specific protocols. For IP, IPX, and AppleTalk, you can provide more granular broadcast control by instead creating the following types of VLAN:

- **IP subnet VLAN** – An IP subnet broadcast domain for a specific IP subnet.

- **IPX network VLAN** – An IPX network broadcast domain for a specific IPX network.
- **AppleTalk cable VLAN** – An AppleTalk broadcast domain for a specific cable range.

You can configure these types of VLANs on Layer 3 Switches only. The Layer 3 Switch sends broadcasts for the IP subnet, IPX network, or AppleTalk cable range to all ports within the IP subnet, IPX network, or AppleTalk cable VLAN at Layer 2.

The Layer 3 Switch routes packets between VLANs at Layer 3. To configure an IP subnet, IPX network, or AppleTalk cable VLAN to route, you must add a virtual routing interface to the VLAN, then configure the appropriate routing parameters on the virtual routing interface.

NOTE: The Layer 3 Switch routes packets between VLANs of the same protocol. The Layer 3 Switch cannot route from one protocol to another.

NOTE: IP subnet VLANs are not the same thing as IP protocol VLANs. An IP protocol VLAN sends all IP broadcasts on the ports within the IP protocol VLAN. An IP subnet VLAN sends only the IP subnet broadcasts for the subnet of the VLAN. You cannot configure an IP protocol VLAN and an IP subnet VLAN within the same port-based VLAN.

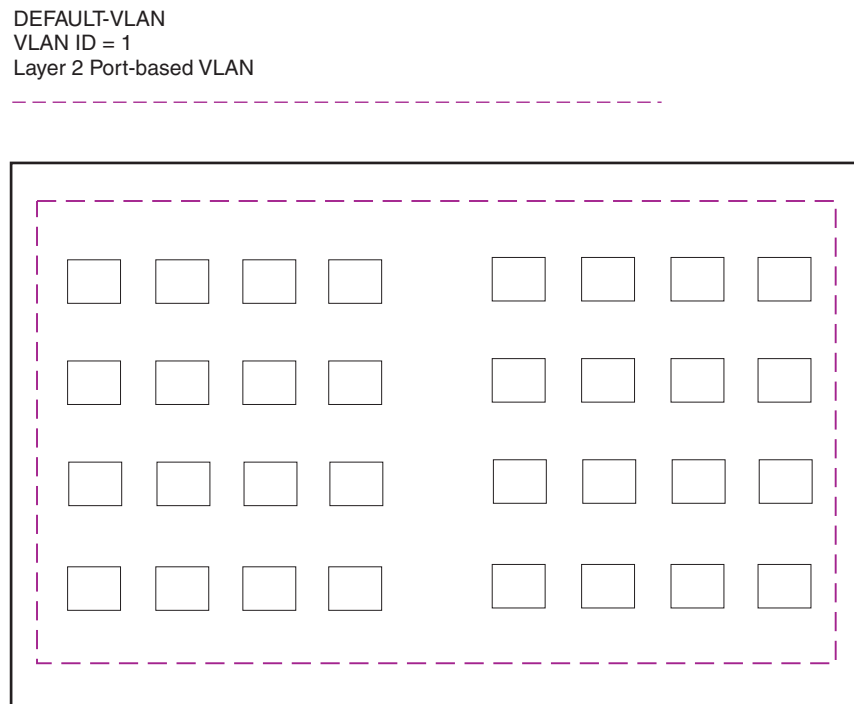
This note also applies to IPX protocol VLANs and IPX network VLANs, and to AppleTalk protocol VLANs and AppleTalk cable VLANs.

Default VLAN

By default, all the ports on a Foundry device are in a single port-based VLAN. This VLAN is called DEFAULT-VLAN and is VLAN number 1. Foundry devices do not contain any protocol VLANs or IP subnet, IPX network, or AppleTalk cable VLANs by default.

Figure 22.3 shows an example of the default Layer 2 port-based VLAN.

Figure 22.3 Default Layer 2 port-based VLAN



By default, all ports belong to a single port-based VLAN, DEFAULT-VLAN. Thus, all ports belong to a single Layer 2 broadcast domain.

When you configure a port-based VLAN, one of the configuration items you provide is the ports that are in the VLAN. When you configure the VLAN, the Foundry device automatically removes the ports that you place in the VLAN from DEFAULT-VLAN. By removing the ports from the default VLAN, the Foundry device ensures that each port resides in only one Layer 2 broadcast domain.

NOTE: Information for the default VLAN is available only after you define another VLAN.

Some network configurations may require that a port be able to reside in two or more Layer 2 broadcast domains (port-based VLANs). In this case, you can enable a port to reside in multiple port-based VLANs by tagging the port. See the following section.

If your network requires that you use VLAN ID 1 for a user-configured VLAN, you can reassign the default VLAN to another valid VLAN ID. See “Assigning a Different VLAN ID to the Default VLAN” on page 22-16.

802.1q Tagging

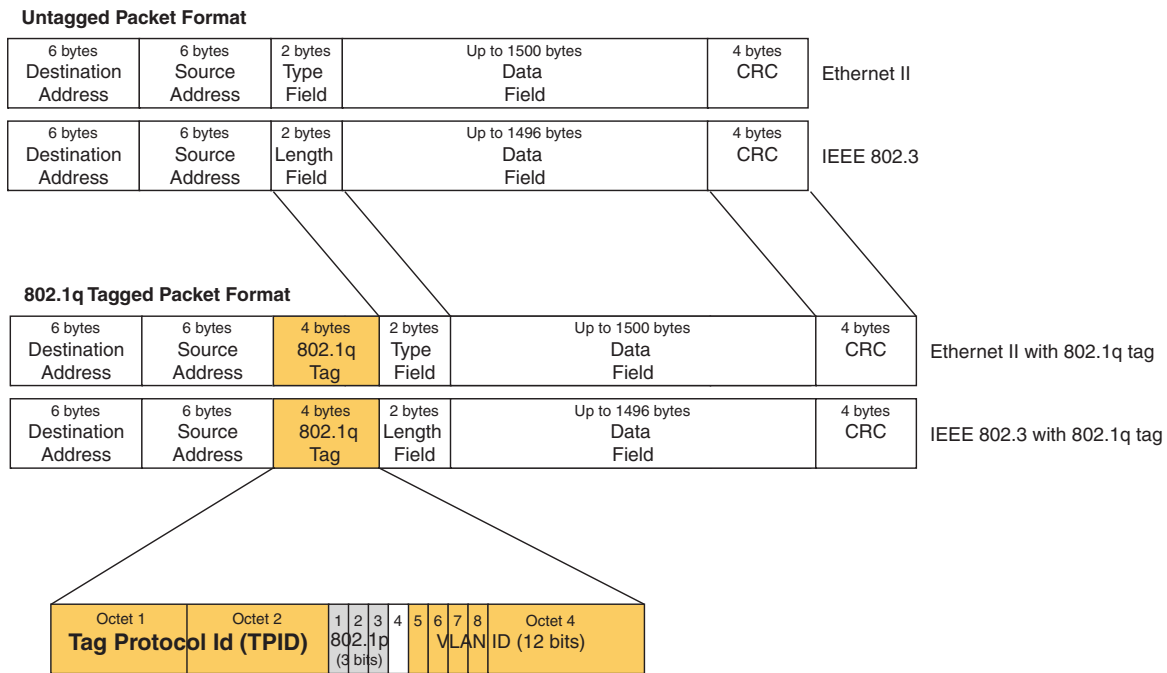
802.1q tagging is an IEEE standard that allows a networking device to add information to a Layer 2 packet in order to identify the VLAN membership of the packet. Foundry devices tag a packet by adding a four-byte tag to the packet. The tag contains the tag value, which identifies the data as a tag, and also contains the VLAN ID of the VLAN from which the packet is sent.

- The default tag value is 8100 (hexadecimal). This value comes from the 802.1q specification. You can change this tag value on a global basis on Foundry devices if needed to be compatible with other vendors' equipment.

- The VLAN ID is determined by the VLAN on which the packet is being forwarded.

Figure 22.4 shows the format of packets with and without the 802.1q tag. The tag format is vendor-specific. To use the tag for VLANs configured across multiple devices, make sure all the devices support the same tag format.

Figure 22.4 Packet containing Foundry's 802.1QVLAN tag




NOTE: You cannot configure a port to be a member of the default port-based VLAN and another port-based VLAN at the same time. Once you add a port to a port-based VLAN, the port is no longer a member of the default VLAN. The port returns to the default VLAN only if you delete the other VLAN(s) that contains the port.

If you configure a VLAN that spans multiple devices, you need to use tagging only if a port connecting one of the devices to the other is a member of more than one port-based VLAN. If a port connecting one device to the other is a member of only a single port-based VLAN, tagging is not required.

If you use tagging on multiple devices, each device must be configured for tagging and must use the same tag value. In addition, the implementation of tagging must be compatible on the devices. The tagging on all Foundry devices is compatible with other Foundry devices.

Figure 22.5 shows an example of two devices that have the same Layer 2 port-based VLANs configured across them. Notice that only one of the VLANs requires tagging.

Figure 22.5 VLANs configured across multiple devices



802.1q-in-q Tagging

NOTE: This feature is supported FastIron Edge Switches running software release 03.1.00 or later.

Software release 03.1.00 and later provide finer granularity for configuring 802.1q tagging, enabling you to configure 802.1Q tag-types on a group of ports, thereby enabling the creation of two identical 802.1Q tags (802.1Q-in-Q tagging) on a single device. This enhancement improves SAV interoperability between Foundry devices and other vendors' devices that support the 802.1Q tag-types, but are not very flexible with the tag-types they accept.

For example applications and configuration details, see "Configuring 802.1q-in-q Tagging" on page 22-57.

802.1q Tag-type Translation

NOTE: This feature is supported on the BigIron MG8 and on Foundry devices running Service Provider release 09.1.00 or later .

The introduction of 802.1q tag-type translation provides finer granularity for configuring multiple 802.1q tag-types on a single device, by enabling you to configure 802.1q tag-types per port group. This enhancement allows for tag-type translation from one port group to the next on tagged interfaces.

For example applications and configuration details, see "Configuring 802.1q Tag-type Translation" on page 22-60

Spanning Tree Protocol (STP)

The default state of STP depends on the device type:

- STP is disabled by default on Foundry Layer 3 Switches.
- STP is enabled by default on Foundry Layer 2 Switches.

Also by default, each port-based VLAN has a separate instance of STP. Thus, when STP is globally enabled, each port-based VLAN on the device runs a separate spanning tree.

You can enable or disable STP on the following levels:

- Globally – Affects all ports on the device.

NOTE: If you configure a port-based VLAN on the device, the VLAN has the same STP state as the default STP state on the device. Thus, on Layer 2 Switches, new VLANs have STP enabled by default. On Layer 3 Switches, new VLANs have STP disabled by default. You can enable or disable STP in each VLAN separately. In addition, you can enable or disable STP on individual ports.

- Port-based VLAN – Affects all ports within the specified port-based VLAN.

STP is a Layer 2 protocol. Thus, you cannot enable or disable STP for individual protocol VLANs or for IP subnet, IPX network, or AppleTalk cable VLANs. The STP state of a port-based VLAN containing these other types of VLANs determines the STP state for all the Layer 2 broadcasts within the port-based VLAN. This is true even though Layer 3 protocol broadcasts are sent on Layer 2 within the VLAN.

It is possible that STP will block one or more ports in a protocol VLAN that uses a virtual routing interface to route to other VLANs. For IP protocol and IP subnet VLANs, even though some of the physical ports of the virtual routing interface are blocked, the virtual routing interface can still route so long as at least one port in the virtual routing interface's protocol VLAN is not blocked by STP.

If you enable Single STP (SSTP) on the device, the ports in all VLANs on which STP is enabled become members of a single spanning tree. The ports in VLANs on which STP is disabled are excluded from the single spanning tree.

For more information, see “Configuring Spanning Tree Protocol (STP) and IronSpan Features” on page 10-1.


Virtual Routing Interfaces

A virtual routing interface is a logical routing interface that Foundry Layer 3 Switches use to route Layer 3 protocol traffic between protocol VLANs.

Foundry devices send Layer 3 traffic at Layer 2 within a protocol VLAN. However, Layer 3 traffic from one protocol VLAN to another must be routed.

If you want the device to be able to send Layer 3 traffic from one protocol VLAN to another, you must configure a virtual routing interface on each protocol VLAN, then configure routing parameters on the virtual routing interfaces. For example, to enable a BigIron Layer 3 Switch to route IP traffic from one IP subnet VLAN to another, you must configure a virtual routing interface on each IP subnet VLAN, then configure the appropriate IP routing parameters on each of the virtual routing interfaces.

Figure 22.6 shows an example of Layer 3 protocol VLANs that use virtual routing interfaces for routing.

Figure 22.6 Use virtual routing interfaces for routing between Layer 3 protocol VLANs

VLAN and Virtual Routing Interface Groups

To simplify configuration, you can configure VLAN groups and virtual routing interface groups. When you create a VLAN group, the VLAN parameters you configure for the group apply to all the VLANs within the group. Additionally, you can easily associate the same IP subnet interface with all the VLANs in a group by configuring a virtual routing interface group with the same ID as the VLAN group.

For configuration information, see "Configuring VLAN Groups and Virtual Routing Interface Groups" on page 22-45.

Dynamic, Static, and Excluded Port Membership

When you add ports to a protocol VLAN, IP subnet VLAN, IPX network VLAN, or AppleTalk cable VLAN, you can add them dynamically or statically:

- Dynamic ports
- Static ports

You also can explicitly exclude ports.

Dynamic Ports

Dynamic ports are added to a VLAN when you create the VLAN. However, if a dynamically added port does not receive any traffic for the VLAN's protocol within ten minutes, the port is removed from the VLAN. However, the port remains a candidate for port membership. Thus, if the port receives traffic for the VLAN's protocol, the device adds the port back to the VLAN.

After the port is added back to the VLAN, the port can remain an active member of the VLAN up to 20 minutes without receiving traffic for the VLAN's protocol. If the port ages out, it remains a candidate for VLAN membership and is added back to the VLAN when the VLAN receives protocol traffic. At this point, the port can remain in the VLAN up to 20 minutes without receiving traffic for the VLAN's protocol, and so on.

Unless you explicitly add a port statically or exclude a port, the port is a dynamic port and thus can be an active member of the VLAN, depending on the traffic it receives.

NOTE: You cannot configure dynamic ports in an AppleTalk cable VLAN. The ports in an AppleTalk cable VLAN must be static. However, ports in an AppleTalk protocol VLAN can be dynamic or static.


Figure 22.7 shows an example of a VLAN with dynamic ports. Dynamic ports not only join and leave the VLAN according to traffic, but also allow some broadcast packets of the specific protocol to "leak" through the VLAN. See "Broadcast Leaks" on page 22-13.

Figure 22.7 VLAN with dynamic ports—all ports are active when you create the VLAN

A = active port

C = candidate port

When you add ports dynamically, all the ports are added when you add the VLAN.




Ports in a new protocol VLAN that do not receive traffic for the VLAN's protocol age out after 10 minutes and become candidate ports. Figure 22.8 shows what happens if a candidate port receives traffic for the VLAN's protocol.

Figure 22.8 VLAN with dynamic ports—candidate ports become active again if they receive protocol traffic

Ports that time out remain candidates for membership in the VLAN and become active again if they receive traffic for the VLAN's protocol, IP sub-net, IPX network, or AppleTalk cable range.

When a candidate port rejoins a VLAN, the timeout for that port becomes 20 minutes. Thus, the port remains an active member of the VLAN even if it does not receive traffic for 20 minutes. After that, the port becomes a candidate port again.



Static Ports

Static ports are permanent members of the protocol VLAN. The ports remain active members of the VLAN regardless of whether the ports receive traffic for the VLAN's protocol. You must explicitly identify the port as a static port when you add it to the VLAN. Otherwise, the port is dynamic and is subject to aging out.

Excluded Ports

If you want to prevent a port in a port-based VLAN from ever becoming a member of a protocol, IP subnet, IPX network, or AppleTalk cable VLAN configured in the port-based VLAN, you can explicitly exclude the port. You exclude the port when you configure the protocol, IP subnet, IPX network, or AppleTalk cable VLAN.

Excluded ports do not leak broadcast packets. See "Broadcast Leaks" on page 22-13.

Broadcast Leaks

A dynamic port becomes a member of a Layer 3 protocol VLAN when traffic from the VLAN's protocol is received on the port. After this point, the port remains an active member of the protocol VLAN, unless the port does not receive traffic from the VLAN's protocol for 20 minutes. If the port does not receive traffic for the VLAN's protocol for 20 minutes, the port ages out and is no longer an active member of the VLAN.

To enable a host that has been silent for awhile to send and receive packets, the dynamic ports that are currently members of the Layer 3 protocol VLAN "leak" Layer 3 broadcast packets to the ports that have aged out. When a host connected to one of the aged out ports responds to a leaked broadcast, the port is added to the protocol VLAN again.

To "leak" Layer 3 broadcast traffic, an active port sends 1/8th of the Layer 3 broadcast traffic to the inactive (aged out) ports.

Static ports do not age out and do not leak broadcast packets.

Super Aggregated VLANs

You can aggregate multiple VLANs within another VLAN. This feature allows you to construct Layer 2 paths and channels. This feature is particularly useful for Virtual Private Network (VPN) applications in which you need to provide a private, dedicated Ethernet connection for an individual client to transparently reach its subnet across multiple networks.

For an application example and configuration information, see “Configuring Super Aggregated VLANs” on page 22-50.

Trunk Group Ports and VLAN Membership

A trunk group is a set of physical ports that are configured to act as a single physical interface. Each trunk group's port configuration is based on the configuration of the lead port, which is the lowest numbered port in the group.

If you add a trunk group's lead port to a VLAN, all of the ports in the trunk group become members of that VLAN.

Summary of VLAN Configuration Rules

A hierarchy of VLANs exists between the Layer 2 and Layer 3 protocol-based VLANs:

- Port-based VLANs are at the lowest level of the hierarchy.
- Layer 3 protocol-based VLANs, IP, IPv6, IPX, AppleTalk, Decnet, and NetBIOS are at the middle level of the hierarchy.
- IP subnet, IPX network, and AppleTalk cable VLANs are at the top of the hierarchy.

NOTE: You cannot have a protocol-based VLAN and a subnet or network VLAN of the same protocol type in the same port-based VLAN. For example, you can have an IPX protocol VLAN and IP subnet VLAN in the same port-based VLAN, but you cannot have an IP protocol VLAN and an IP subnet VLAN in the same port-based VLAN, nor can you have an IPX protocol VLAN and an IPX network VLAN in the same port-based VLAN.

As a Foundry device receives packets, the VLAN classification starts from the highest level VLAN first. Therefore, if an interface is configured as a member of both a port-based VLAN and an IP protocol VLAN, IP packets coming into the interface are classified as members of the IP protocol VLAN because that VLAN is higher in the VLAN hierarchy.

Multiple VLAN Membership Rules

- A port can belong to multiple, unique, overlapping Layer 3 protocol-based VLANs without VLAN tagging.
- A port can belong to multiple, overlapping Layer 2 port-based VLANs only if the port is a tagged port. Packets sent out of a tagged port use an 802.1q-tagged frame.
- When both port and protocol-based VLANs are configured on a given device, all protocol VLANs must be strictly contained within a port-based VLAN. A protocol VLAN cannot include ports from multiple port-based VLANs. This rule is required to ensure that port-based VLANs remain loop-free Layer 2 broadcast domains.
- IP protocol VLANs and IP subnet VLANs cannot operate concurrently on the system or within the same port-based VLAN.
- IPX protocol VLANs and IPX network VLANs cannot operate concurrently on the system or within the same port-based VLAN.
- If you first configure IP and IPX protocol VLANs before deciding to partition the network by IP subnet and IPX network VLANs, then you need to delete those VLANs before creating the IP subnet and IPX network VLANs.
- One of each type of protocol VLAN is configurable within each port-based VLAN on the Layer 2 Switch.
- Multiple IP subnet and IPX network VLANs are configurable within each port-based VLAN on the Layer 2 Switch.

- Removing a configured port-based VLAN from a Foundry Networks Layer 2 Switch or Layer 3 Switch automatically removes any protocol-based VLAN, IP subnet VLAN, AppleTalk cable VLAN, or IPX network VLAN, or any Virtual Ethernet router interfaces defined within the Port-based VLAN.

Routing Between VLANs (Layer 3 Switches Only)

Foundry Layer 3 Switches can locally route IP, IPX, and Appletalk between VLANs defined within a single router. All other routable protocols or protocol VLANs (for example, DecNet) must be routed by another external router capable of routing the protocol.

Virtual Routing Interfaces (Layer 3 Switches Only)

You need to configure virtual routing interfaces if an IP, IPX, or Appletalk protocol VLAN, IP subnet VLAN, AppleTalk cable VLAN, or IPX network VLAN needs to route protocols to another port-based VLAN on the same router. A virtual routing interface can be associated with the ports in only a single port-based VLAN. Virtual router interfaces must be defined at the highest level of the VLAN hierarchy.

If you do not need to further partition the port-based VLAN by defining separate Layer 3 VLANs, you can define a single virtual routing interface at the port-based VLAN level and enable IP, IPX, and Appletalk routing on a single virtual routing interface.

Bridging and Routing the Same Protocol Simultaneously on the Same Device (Layer 3 Switches Only)

Some configurations may require simultaneous switching and routing of the same single protocol across different sets of ports on the same router. When IP, IPX, or Appletalk routing is enabled on a Foundry Layer 3 Switch, you can route these protocols on specific interfaces while bridging them on other interfaces. In this scenario, you can create two separate backbones for the same protocol, one bridged and one routed.

To bridge IP, IPX, or Appletalk at the same time these protocols are being routed, you need to configure an IP protocol, IP subnet, IPX protocol, IPX network, or Appletalk protocol VLAN and not assign a virtual routing interface to the VLAN. Packets for these protocols are bridged or switched at Layer 2 across ports on the router that are included in the Layer 3 VLAN. If these VLANs are built within port-based VLANs, they can be tagged across a single set of backbone fibers to create separate Layer 2 switched and Layer 3 routed backbones for the same protocol on a single physical backbone.

Routing Between VLANs Using Virtual Routing Interfaces (Layer 3 Switches Only)

Foundry calls the ability to route between VLANs with virtual routing interfaces **Integrated Switch Routing (ISR)**. There are some important concepts to understand before designing an ISR backbone.

Virtual router interfaces can be defined on port-based, IP protocol, IP subnet, IPX protocol, IPX network, AppleTalk protocol, and AppleTalk cable VLANs.

To create any type of VLAN on a Foundry Layer 3 Switch, Layer 2 forwarding must be enabled. When Layer 2 forwarding is enabled, the Layer 3 Switch becomes a Switch on all ports for all non-routable protocols.

If the router interfaces for IP, IPX, or AppleTalk are configured on physical ports, then routing occurs independent of the Spanning Tree Protocol (STP). However, if the router interfaces are defined for any type VLAN, they are virtual routing interfaces and are subject to the rules of STP.

If your backbone is consisted of virtual routing interfaces all within the same STP domain, it is a bridged backbone, not a routed one. This means that the set of backbone interfaces that are blocked by STP will be blocked for routed protocols as well. The routed protocols will be able to cross these paths only when the STP state of the link is FORWARDING. This problem is easily avoided by proper network design.

When designing an ISR network, pay attention to your use of virtual routing interfaces and the spanning-tree domain. If Layer 2 switching of your routed protocols (IP, IPX, AppleTalk) is not required across the backbone, then the use of virtual routing interfaces can be limited to edge switch ports within each router. Full backbone routing can be achieved by configuring routing on each physical interface that connects to the backbone. Routing is independent of STP when configured on a physical interface.

If your ISR design requires that you switch IP, IPX, or Appletalk at Layer 2 while simultaneously routing the same protocols over a single backbone, then create multiple port-based VLANs and use VLAN tagging on the backbone links to separate your Layer 2 switched and Layer 3 routed networks.

There is a separate STP domain for each port-based VLAN. Routing occurs independently across port-based VLANs or STP domains. You can define each end of each backbone link as a separate tagged port-based VLAN. Routing will occur independently across the port-based VLANs. Because each port-based VLAN's STP domain is a single point-to-point backbone connection, you are guaranteed to never have an STP loop. STP will never block the virtual router interfaces within the tagged port-based VLAN, and you will have a fully routed backbone.

Dynamic Port Assignment (Layer 2 Switches and Layer 3 Switches)

All switch ports are dynamically assigned to any Layer 3 VLAN on Foundry Layer 2 Switches and any non-routable VLAN on Foundry Layer 3 Switches. To maintain explicit control of the VLAN, you can explicitly exclude ports when configuring any Layer 3 VLAN on a Foundry Layer 2 Switch or any non-routable Layer 3 VLAN on a Foundry Layer 3 Switch.

If you do not want the ports to have dynamic membership, you can add them statically. This eliminates the need to explicitly exclude the ports that you do not want to participate in a particular Layer 3 VLAN.

Assigning a Different VLAN ID to the Default VLAN

When you enable port-based VLANs, all ports in the system are added to the default VLAN. By default, the default VLAN ID is "VLAN 1". The default VLAN is not configurable. If you want to use the VLAN ID "VLAN 1" as a configurable VLAN, you can assign a different VLAN ID to the default VLAN.

To reassign the default VLAN to a different VLAN ID, enter the following command:

```
BigIron(config)# default-vlan-id 4095
```

Syntax: [no] default-vlan-d <vlan-id>

You must specify a valid VLAN ID that is not already in use. For example, if you have already defined VLAN 10, do not try to use "10" as the new VLAN ID for the default VLAN. Valid VLAN IDs are numbers from 1 – 4096.

NOTE: Changing the default VLAN name does not change the properties of the default VLAN. Changing the name allows you to use the VLAN ID "1" as a configurable VLAN.

Assigning Trunk Group Ports

When a "lead" trunk group port is assigned to a VLAN, all other members of the trunk group are automatically added to that VLAN. A lead port is the first port of a trunk group port range; for example, "1" in 1 – 4 or "5" in 5 – 8. See "Trunk Group Rules" on page 17-3 for more information.

Configuring Port-Based VLANs

Port-based VLANs allow you to provide separate spanning tree protocol (STP) domains or broadcast domains on a port-by-port basis.

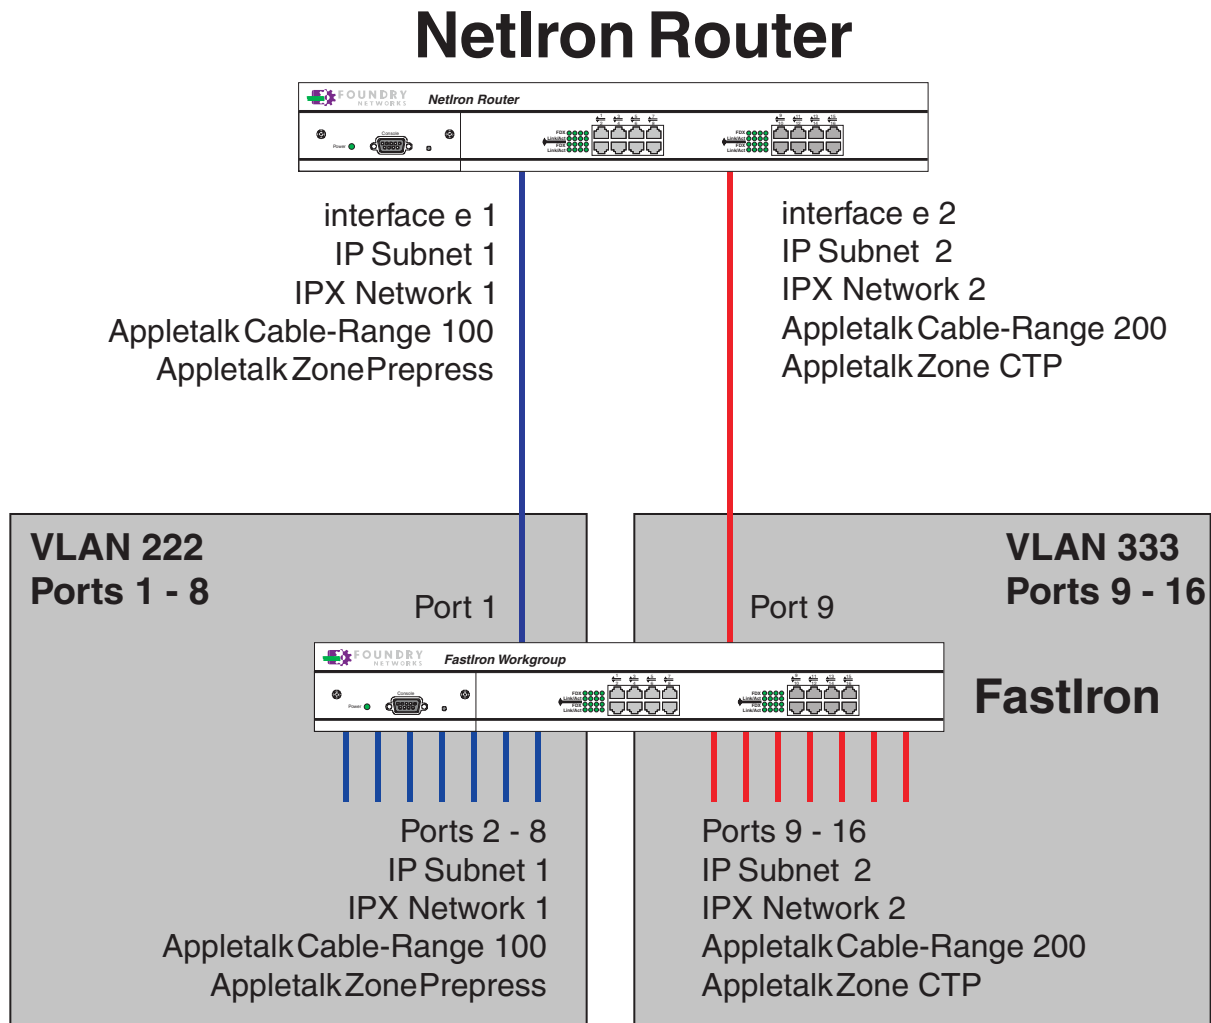
This section describes how to perform the following tasks for port-based VLANs using the CLI:

- Create a VLAN.
- Delete a VLAN.
- Modify a VLAN.
- Assign a higher priority to the VLAN.
- Change a VLAN's priority.
- Enable or disable STP on the VLAN.

EXAMPLE:

Figure 22.9 shows a simple port-based VLAN configuration using a single Foundry Layer 2 Switch. All ports within each VLAN are untagged. One untagged port within each VLAN is used to connect the Layer 2 Switch to a Layer 3 Switch (in this example, a Netron) for Layer 3 connectivity between the two port-based VLANs.

Figure 22.9 Port-based VLANs 222 and 333



To create the two port-based VLANs shown in Figure 22.9, use the following method.

USING THE CLI

```
FastIron(config)# vlan 222 by port
FastIron(config-vlan-222)# untag e 1 to
FastIron(config-vlan-222)# vlan 333 by port
FastIron(config-vlan-333)# untag e 95 to 16
```

Syntax: vlan <vlan-id> by port


Syntax: untagged ethernet <portnum> [to <portnum> | ethernet <portnum>]

EXAMPLE:

Figure 22.10 shows a more complex port-based VLAN configuration using multiple Layer 2 Switches and IEEE 802.1q VLAN tagging. The backbone link connecting the three Layer 2 Switches is tagged. One untagged port

within each port-based VLAN on FastIron-A connects each separate network wide Layer 2 broadcast domain to the router for Layer 3 forwarding between broadcast domains. The STP priority is configured to force FastIron-A to be the root bridge for VLANs RED and BLUE. The STP priority on FastIron-B is configured so that FastIron-B is the root bridge for VLANs GREEN and BROWN.

Figure 22.10 More complex port-based VLAN



To configure the Port-based VLANs on the FastIron Layer 2 Switches in Figure 22.10, use the following method.

USING THE CLI

Configuring FastIron-A

Enter the following commands to configure FastIron-A:

```
FastIron> enable
FastIron# configure terminal
FastIron(config)# hostname FastIron-A
FastIron-A(config)# vlan 2 name BROWN
FastIron-A(config-vlan-2)# untag ethernet 1 to 4 ethernet 17
FastIron-A(config-vlan-2)# tag ethernet 25 to 26
FastIron-A(config-vlan-2)# spanning-tree
FastIron-A(config-vlan-2)# vlan 3 name GREEN
FastIron-A(config-vlan-3)# untag ethernet 5 to 8 ethernet 18
FastIron-A(config-vlan-3)# tag ethernet 25 to 26
FastIron-A(config-vlan-3)# spanning-tree
FastIron-A(config-vlan-3)# vlan 4 name BLUE
FastIron-A(config-vlan-4)# untag ethernet 9 to 12 ethernet 19
FastIron-A(config-vlan-4)# tag ethernet 25 to 26
FastIron-A(config-vlan-4)# spanning-tree
FastIron-A(config-vlan-4)# spanning-tree priority 500
FastIron-A(config-vlan-4)# vlan 5 name RED
FastIron-A(config-vlan-5)# untag ethernet 13 to 16 ethernet 20
```



```
FastIron-A(config-vlan-5)# tag ethernet 25 to 26
FastIron-A(config-vlan-5)# spanning-tree
FastIron-A(config-vlan-5)# spanning-tree priority 500
FastIron-A(config-vlan-5)# end
FastIron-A# write memory
```

Configuring FastIron-B

Enter the following commands to configure FastIron-B:

```
FastIron> en
FastIron# configure terminal
FastIron(config)# hostname FastIron-B
FastIron-B(config)# vlan 2 name BROWN
FastIron-B(config-vlan-2)# untag ethernet 1 to 4
FastIron-B(config-vlan-2)# tag ethernet 25 to 26
FastIron-B(config-vlan-2)# spanning-tree
FastIron-B(config-vlan-2)# spanning-tree priority 500
FastIron-B(config-vlan-2)# vlan 3 name GREEN
FastIron-B(config-vlan-3)# untag ethernet 5 to 8
FastIron-B(config-vlan-3)# tag ethernet 25 to 26
FastIron-B(config-vlan-3)# spanning-tree
FastIron-B(config-vlan-3)# spanning-tree priority 500
FastIron-B(config-vlan-3)# vlan 4 name BLUE
FastIron-B(config-vlan-4)# untag ethernet 9 to 12
FastIron-B(config-vlan-4)# tag ethernet 25 to 26
FastIron-B(config-vlan-4)# vlan 5 name RED
FastIron-B(config-vlan-5)# untag ethernet 13 to 16
FastIron-B(config-vlan-5)# tag ethernet 25 to 26
FastIron-B(config-vlan-5)# end
FastIron-B# write memory
```

Configuring FastIron-C

Enter the following commands to configure FastIron-C:

```
FastIron> en
FastIron# configure terminal
FastIron(config)# hostname FastIron-C
FastIron-C(config)# vlan 2 name BROWN
FastIron-C(config-vlan-2)# untag ethernet 1 to 4
FastIron-C(config-vlan-2)# tag ethernet 25 to 26
FastIron-C(config-vlan-2)# vlan 3 name GREEN
FastIron-C(config-vlan-3)# untag ethernet 5 to 8
FastIron-C(config-vlan-3)# tag ethernet 25 to 26
FastIron-C(config-vlan-3)# vlan 4 name BLUE
FastIron-C(config-vlan-4)# untag ethernet 9 to 12
FastIron-C(config-vlan-4)# tag ethernet 25 to 26
FastIron-C(config-vlan-4)# vlan 5 name RED
FastIron-C(config-vlan-5)# untag ethernet 13 to 16
FastIron-C(config-vlan-5)# tag ethernet 25 to 26
FastIron-C(config-vlan-5)# end
FastIron-C# write memory
```

Syntax: vlan <vlan-id> by port

Syntax: untagged ethernet | pos <portnum> [to <portnum> | ethernet <portnum>]

Syntax: tagged ethernet | pos <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] spanning-tree

Syntax: spanning-tree [ethernet <portnum> path-cost <value> priority <value>] forward-delay <value> hello-time <value> maximum-age <time> priority <value>

Modifying a Port-Based VLAN

You can make the following modifications to a port-based VLAN:

- Add or delete a VLAN port.
- Change its priority.
- Enable or disable STP.

Removing a Port-Based VLAN

Suppose you want to remove VLAN 5 from the example in Figure 22.10. To do so, use the following procedure.

USING THE CLI

1. Access the global CONFIG level of the CLI on FastIron-A by entering the following commands:

```
FastIron-A> enable
No password has been assigned yet...
FastIron-A# configure terminal
FastIron-A(config)#
```

2. Enter the following command:

```
FastIron-A(config)# no vlan 5
FastIron-A(config)#
```

3. Enter the following commands to exit the CONFIG level and save the configuration to the system-config file on flash memory:

```
FastIron-A(config)#
FastIron-A(config)# end
FastIron-A# write memory
FastIron-A#
```

4. Repeat steps 1 – 3 on FastIron-B.

Syntax: no vlan <vlan-id> by port

Removing a Port from a VLAN

Suppose you want to remove port 11 from VLAN 4 on FastIron-A shown in Figure 22.10. To do so, use the following procedure.

USING THE CLI

1. Access the global CONFIG level of the CLI on FastIron-A by entering the following command:

```
FastIron-A> enable
No password has been assigned yet...
FastIron-A# configure terminal
FastIron-A(config)#
```

2. Access the level of the CLI for configuring port-based VLAN 4 by entering the following command:

```
FastIron-A(config)#
FastIron-A(config)# vlan 4
FastIron-A(config-vlan-4)#
```

3. Enter the following commands:

```
FastIron-A(config-vlan-4)#
FastIron-A(config-vlan-4)# no untag ethernet 11
deleted port ethe 11 from port-vlan 4.
FastIron-A(config-vlan-4)#
```

4. Enter the following commands to exit the VLAN CONFIG mode and save the configuration to the system-config file on flash memory:

```
FastIron-A(config-vlan-4)#
FastIron-A(config-vlan-4)# end
FastIron-A# write memory
FastIron-A#
```

NOTE: Beginning in software release 07.5.00, you can remove all the ports from a port-based VLAN without losing the rest of the VLAN's configuration. However, you cannot configure an IP address on a virtual routing interface unless the VLAN contains ports. If the VLAN has a virtual routing interface, the virtual routing interface's IP address is deleted when the ports associated with the interface are deleted. The rest of the VLAN configuration is retained.

In software releases earlier than 07.5.00, if you remove all the ports from a VLAN, the software removes the VLAN configuration entirely.

Assigning a Higher Priority to a VLAN

Suppose you wanted to give all traffic on Purple VLAN 2 in Figure 22.10 higher priority than all the other VLANs. Use the following procedure to do so.

USING THE CLI

1. Access the global CONFIG level of the CLI on FastIron-A by entering the following command:

```
FastIron-A> enable
No password has been assigned yet...
FastIron-A# configure terminal
FastIron-A(config)#
```

2. Access the level of the CLI for configuring port-based VLAN 2 by entering the following command:

```
FastIron-A(config)#
FastIron-A(config)# vlan 2
FastIron-A(config-vlan-2)#
```

3. Enable all packets exiting the Layer 2 Switch on VLAN 2 to transmit from the high priority hardware queue of each transmit interface. Note that possible QoS priority levels for Foundry Stackable devices are normal or high. possible levels are 0 (normal) – 7 (highest).

```
FastIron-A(config-vlan-2)#
FastIron-A(config-vlan-2)# priority high
FastIron-A(config-vlan-2)#
```

4. Enter the following commands to exit the VLAN CONFIG mode and save the configuration to the system-config file on flash memory:

```
FastIron-A(config-vlan-2)#
FastIron-A(config-vlan-2)# end
FastIron-A# write memory
FastIron-A#
```

5. Repeat steps 1 – 4 on FastIron-B.

Syntax: vlan <vlan-id> by port

Syntax: priority normal | high

Enable Spanning Tree on a VLAN

The spanning tree bridge and port parameters are configurable using one CLI command set at the Global Configuration Level of each Port-based VLAN. Suppose you want to enable the IEEE 802.1d STP across VLAN 3. To do so, use the following method.

NOTE: When port-based VLANs are not operating on the system, STP is set on a system-wide level at the global CONFIG level of the CLI.

USING THE CLI

1. Access the global CONFIG level of the CLI on FastIron-A by entering the following commands:

```
FastIron-A> enable
No password has been assigned yet...
FastIron-A# configure terminal
FastIron-A(config)#
```

2. Access the level of the CLI for configuring port-based VLAN 3 by entering the following command:

```
FastIron-A(config)#
FastIron-A(config)# vlan 3
FastIron-A(config-vlan-3)#
```

3. From VLAN 3's configuration level of the CLI, enter the following command to enable STP on all tagged and untagged ports associated with VLAN 3.

```
FastIron-B(config-vlan-3)#
FastIron-B(config-vlan-3)# spanning-tree
FastIron-B(config-vlan-3)#
```

4. Enter the following commands to exit the VLAN CONFIG mode and save the configuration to the system-config file on flash memory:

```
FastIron-B(config-vlan-3)#
FastIron-B(config-vlan-3)# end
FastIron-B# write memory
FastIron-B#
```

5. Repeat steps 1 – 4 on FastIron-B.

NOTE: You do not need to configure values for the STP parameters. All parameters have default values as noted below. Additionally, all values will be globally applied to all ports on the system or on the port-based VLAN for which they are defined.

To configure a specific path-cost or priority value for a given port, enter those values using the key words in the brackets [] shown in the syntax summary below. If you do not want to specify values for any given port, this portion of the command is not required.

Syntax: vlan <vlan-id> by port

Syntax: [no] spanning-tree

Syntax: spanning-tree [ethernet <portnum> path-cost <value> priority <value>] forward-delay <value> hello-time <value> maximum-age <time> priority <value>

Bridge STP Parameters (applied to all ports within a VLAN)

- Forward Delay – the period of time a bridge will wait (the listen and learn period) before forwarding data packets. Possible values: 4 – 30 seconds. Default is 15.
- Maximum Age – the interval a bridge will wait for receipt of a hello packet before initiating a topology change. Possible values: 6 – 40 seconds. Default is 20.
- Hello Time – the interval of time between each configuration BPDU sent by the root bridge. Possible values: 1 – 10 seconds. Default is 2.
- Priority – a parameter used to identify the root bridge in a network. The bridge with the lowest value has the highest priority and is the root. Possible values: 1 – 65,535. Default is 32,678.

Port Parameters (applied to a specified port within a VLAN)

- Path Cost – a parameter used to assign a higher or lower path cost to a port. Possible values: 1 – 65535. Default is (1000/Port Speed) for Half-Duplex ports and is (1000/Port Speed)/2 for Full-Duplex ports.
- Priority – value determines when a port will be rerouted in relation to other ports. Possible values: 0 – 255. Default is 128.

Configuring IP Subnet, IPX Network and Protocol-Based VLANs

Protocol-based VLANs provide the ability to define separate broadcast domains for several unique Layer 3 protocols within a single Layer 2 broadcast domain. Some applications for this feature might include security between departments with unique protocol requirements. This feature enables you to limit the amount of broadcast traffic end-stations, servers, and routers need to accept.

NOTE: See “Configuring AppleTalk Cable VLANs” on page 22-35 for information about configuring an AppleTalk cable VLAN.

Configuration Notes for FES Devices

NOTE: This section applies to the FastIron Edge Switch (FES) running software release 03.1.00 or later.

- When Protocol-based or Subnet VLANs are enabled, the FastIron Edge Switch forwards unknown unicast, unknown multicast, and broadcast packets in software. By default, the FastIron Edge Switch forwards unknown unicast, unknown multicast, and broadcast packets in hardware.
- Dynamic port maintenance occurs based on unknown unicast, unknown multicast, and broadcast traffic only. If a port keeps receiving known unicast traffic only, then the port will age out from the VLAN membership and will not be an active member of the Protocol or Subnet VLAN.

Configuration Considerations for BigIron MG8 and NetIron 40G

NOTE: This section applies to the BigIron MG8 and NetIron 40G running software release 02.0.00 or later.

Note the following configuration limitations:

- The dynamic protocol VLAN option is not supported.
- The other-protocol option defines a protocol-based VLAN for protocols that do not require a singular protocol broadcast domain or are not currently supported on the Foundry device. It is used as a catch-all rule to mean all other protocols in addition to those already assigned.

For example, in the following VLAN configuration, IP protocol is defined and the "other-proto" option is set to become operational when a non-IPv4 packet is received.

```
BigIron MG8 Router(config)#vlan 5
BigIron MG8 Router(config-vlan-5)#ip-proto
BigIron MG8 Router(config-vlan-5)#other-proto
```

Configuration Example

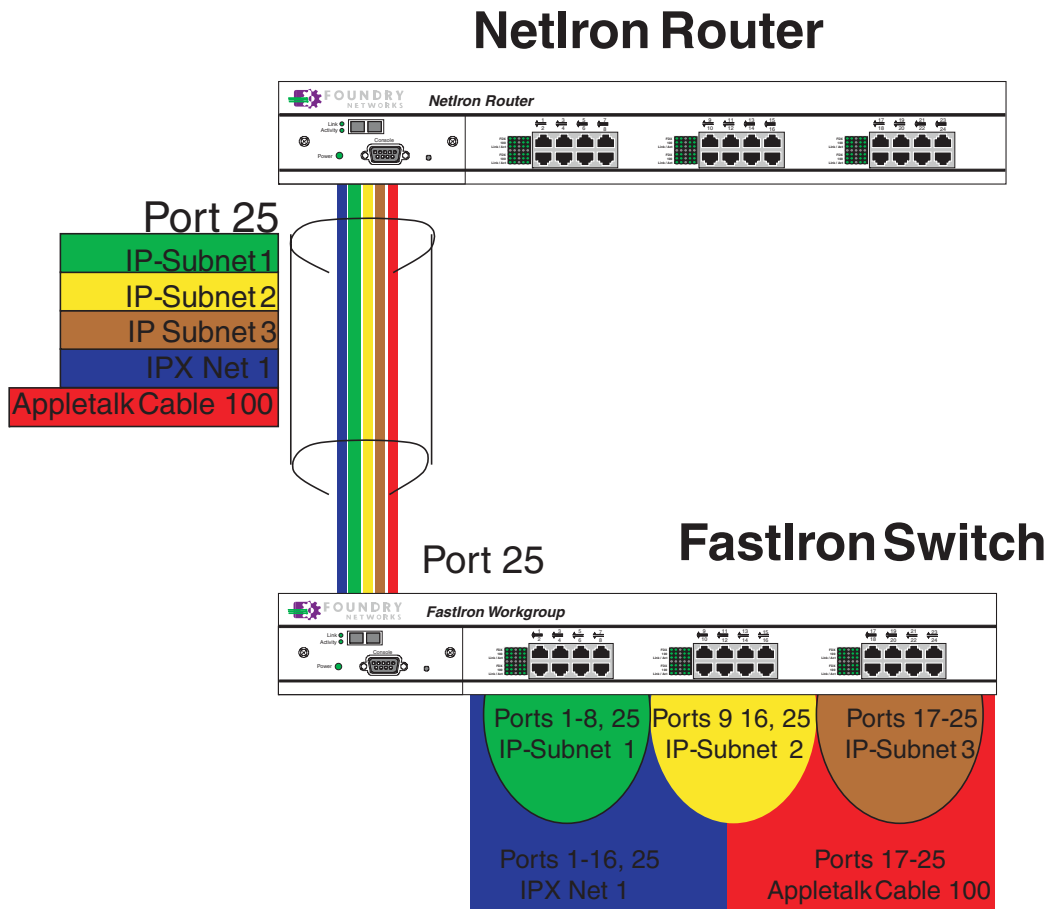
Suppose you want to create five separate Layer 3 broadcast domains within a single Layer 2 STP broadcast domain:

- Three broadcast domains, one for each of three separate IP subnets
- One for IPX Network 1
- One for the Appletalk protocol

Also suppose you want a single router interface to be present within all of these separate broadcast domains, without using IEEE 802.1q VLAN tagging or any proprietary form of VLAN tagging.

Figure 22.11 shows this configuration.

Figure 22.11 Protocol-based (Layer 3) VLANs



To configure the VLANs shown in Figure 22.11, use the following procedure.

USING THE CLI

1. To permanently assign ports 1 – 8 and port 25 to IP subnet VLAN 1.1.1.0, enter the following commands:

```
FastIron> en
No password has been assigned yet...
FastIron# config t
FastIron(config)#
FastIron(config)# ip-subnet 1.1.1.0/24 name Green
FastIron(config-ip-subnet)# no dynamic
FastIron(config-ip-subnet)# static ethernet 1 to 8 ethernet 25
```

2. To permanently assign ports 9 – 16 and port 25 to IP subnet VLAN 1.1.2.0, enter the following commands:

```
FastIron(config-ip-subnet)# ip-subnet 1.1.2.0/24 name Yellow
FastIron(config-ip-subnet)# no dynamic
FastIron(config-ip-subnet)# static ethernet 9 to 16 ethernet 25
```

3. To permanently assign ports 17 – 25 to IP subnet VLAN 1.1.3.0, enter the following commands:

```
FastIron(config-ip-subnet)# ip-subnet 1.1.3.0/24 name Brown
FastIron(config-ip-subnet)# no dynamic
FastIron(config-ip-subnet)# static ethernet 17 to 25
```

4. To permanently assign ports 1 – 12 and port 25 to IPX network 1 VLAN, enter the following commands:

```
FastIron(config-ip-subnet)# ipx-network 1 ethernet_802.3 name Blue
FastIron(config-ipx-network)# no dynamic
FastIron(config-ipx-network)# static ethernet 1 to 12 ethernet 25
FastIron(config-ipx-network)#
```

5. To permanently assign ports 12 – 25 to Appletalk VLAN, enter the following commands:

```
FastIron(config-ipx-PROTO)# atalk-PROTO name Red
FastIron(config-ataalk-PROTO)# no dynamic
FastIron(config-ataalk-PROTO)# static ethernet 13 to 25
FastIron(config-ataalk-PROTO)# end
FastIron# write memory
FastIron#
```

Syntax: ip-subnet <ip-addr> <ip-mask> [name <string>]

Syntax: ipx-network <ipx-network-number> <frame-encapsulation-type> netbios-allow | netbios-disallow [name <string>]

Syntax: ip-PROTO | ipx-PROTO | atalk-PROTO | decnet-PROTO | netbios-PROTO | other-PROTO
static | exclude | dynamic
ethernet <portnum> [to <portnum>] [name <string>]

Configuring IP subnet, IPX Network, and Protocol-Based VLANs Within Port-Based VLANs

If you plan to use port-based VLANs in conjunction with protocol-based VLANs, you must create the port-based VLANs first. Once you create a port-based VLAN, then you can assign Layer 3 protocol VLANs within the boundaries of the port-based VLAN. Generally, you create port-based VLANs to allow multiple separate STP domains.

EXAMPLE:

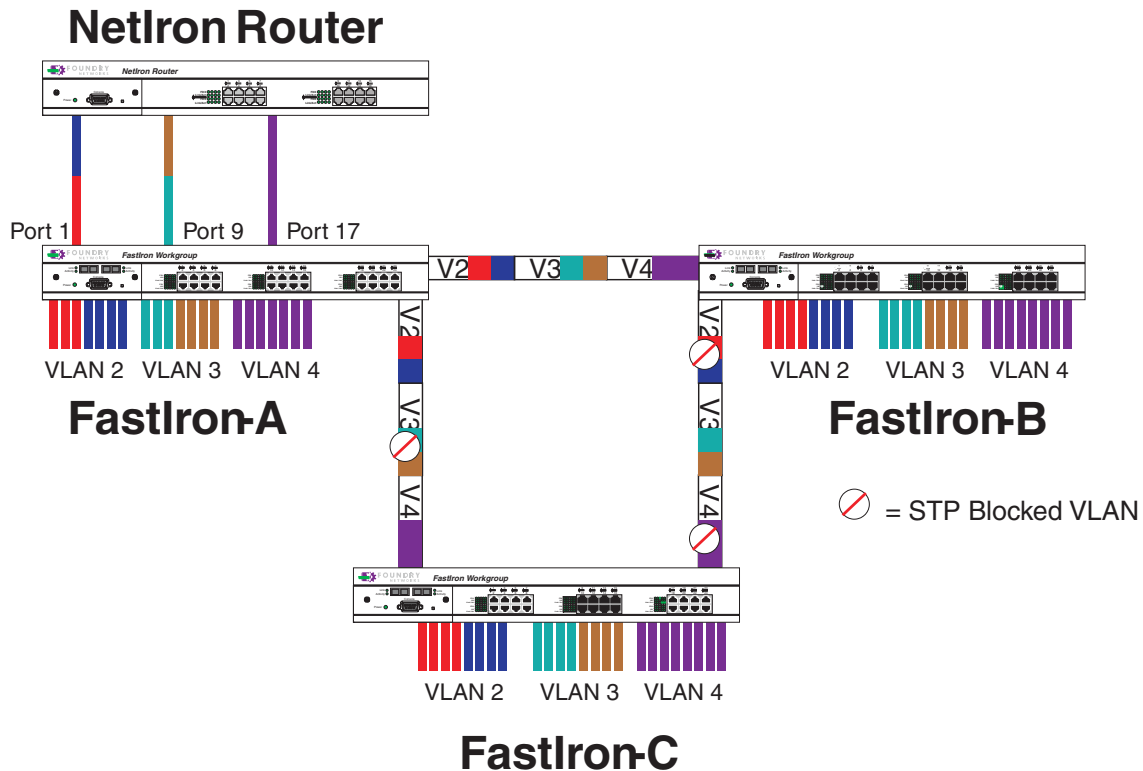
Suppose you need to provide three separate STP domains across an enterprise campus backbone. The first STP domain (VLAN 2) requires a set of ports at each Layer 2 Switch location to be statically mapped to IP only. No other protocols can enter the switches on this set of ports.

A second set of ports within STP domain VLAN 2 will be restricted to only IPX traffic. The IP and IPX protocol VLANs will overlap on Port 1 of FastIron-A to support both protocols on the same router interface. The IP subnets and IPX network that span the two protocol VLANs will be determined by the NetIron router configuration. The IP and IPX Protocol VLANs ensure that only the ports included in the each Layer 3 protocol VLAN will see traffic from the NetIron router.

The second STP domain (VLAN 3) requires that half the ports in the domain are dedicated to IP subnet 1.1.1.0/24 and the other ports are dedicated to IPX network 1. Similar to VLAN 2, Port 9 from VLAN 3 will be used to carry this IP subnet and IPX network to the NetIron router. No other protocols will be allowed to enter the network on VLAN 3. Also, no IP packets with a source address on subnet 1.1.1.0/24 or IPX packets with a source address on network 1 will be allowed to enter the switches on VLAN 3.

There is no need to segment Layer 3 broadcast domains within the STP broadcast domain (VLAN 4). The NetIron router will dictate the IP subnets and IPX network that are on VLAN 4. There are no Layer 3 protocol restrictions on VLAN 4; however, the NetIron router is configured to only forward IP and IPX between STP domains.

Figure 22.12 More protocol-based VLANs



To configure the Layer 3 VLANs on the FastIron Layer 2 Switches in Figure 22.12, use the following procedure.

USING THE CLI

Configuring FastIron-A

Enter the following commands to configure FastIron-A:

1. Create port-based VLAN 2 and assign the untagged and tagged ports that will participate in this VLAN:

```
FastIron-A >en
FastIron-A# config t
FastIron-A(config)# vlan 2 name IP_IPX_Protocol
FastIron-A(config-vlan-2)# untag e1 to 8
FastIron-A(config-vlan-2)# tag e25 to 26
```

2. Enable STP and set the priority to force FastIron-A to be the root bridge for VLAN 2:

```
FastIron-A(config-vlan-2)# spanning-tree
FastIron-A(config-vlan-2)# spanning-tree priority 500
FastIron-A(config-vlan-2)#
```

3. Create the IP and IPX protocol-based VLANs and statically assign the ports within VLAN 2 that will be associated with each protocol-based VLAN:

```
FastIron-A(config-vlan-2)# ip-proto name Red
FastIron-A(config-vlan-ip-proto)# no dynamic
FastIron-A(config-vlan-ip-proto)# static e1 to 4 e25 to 26
FastIron-A(config-vlan-ip-proto)# exclude e5 to 8
FastIron-A(config-vlan-ip-proto)# ipx-proto name Blue
FastIron-A(config-vlan-ipx-proto)# no dynamic
```



```
FastIron-A(config-vlan-ipx-PROTO)# static e1 e5 to 8 e25 to 26
FastIron-A(config-vlan-ipx-PROTO)# exclude e2 to 4
```

- To prevent machines with non-IP protocols from getting into the IP portion of VLAN 2, create another Layer 3 protocol VLAN to exclude all other protocols from the ports that contains the IP-protocol VLAN. To do so, enter the following commands:

```
FastIron-A(config-vlan-ipx-PROTO)# other-PROTO name Block_other_PROTO
FastIron-A(config-vlan-other-PROTO)# no dynamic
FastIron-A(config-vlan-other-PROTO)# exclude e1 to 8
FastIron-A(config-vlan-other-PROTO)#
```

- Create port-based VLAN 3. Note that FastIron-B will be the root for this STP domain, so you do not need to adjust the STP priority.

```
FastIron-A(config-vlan-other-PROTO)# vlan 3 name IP-Sub_IPX-Net_Vlans
FastIron-A(config-vlan-3)# untag e9 to 16
FastIron-A(config-vlan-3)# tag e25 to 26
FastIron-A(config-vlan-3)# spanning-tree
FastIron-A(config-vlan-3)#
```

- Create IP subnet VLAN 1.1.1.0/24, IPX network 1, and other-protocol VLANs

```
FastIron-A(config-vlan-3)# ip-subnet 1.1.1.0/24 name Green
FastIron-A(config-vlan-ip-subnet)# no dynamic
FastIron-A(config-vlan-ip-subnet)# static e9 to 12 e25 to 26
FastIron-A(config-vlan-ip-subnet)# exclude e13 to 16
FastIron-A(config-vlan-ip-subnet)# ipx-net 1 ethernet_802.3 name Brown
FastIron-A(config-vlan-ipx-network)# no dynamic
FastIron-A(config-vlan-ipx-network)# static e9 e13 to 16 e25 to 26
FastIron-A(config-vlan-ipx-network)# exclude e10 to 12
FastIron-A(config-vlan-ipx-network)# other-PROTO name Block_other_PROTO
FastIron-A(config-vlan-other-PROTO)# no dynamic
FastIron-A(config-vlan-other-PROTO)# exclude e9 to 16
FastIron-A(config-vlan-other-PROTO)#
```

- Configure the last port-based VLAN 4. You need to set the STP priority for this VLAN because FastIron-A will be the root bridge for this VLAN. Since you do not need to partition this STP domain into multiple Layer 3 broadcast domains, this is the only configuration required for VLAN 4:

```
FastIron-A(config-vlan-other-PROTO)# vlan 4 name Purple_ALL-Protocols
FastIron-A(config-vlan-4)# untag e17 to 24
FastIron-A(config-vlan-4)# tag e25 to 26
FastIron-A(config-vlan-4)# spanning-tree
FastIron-A(config-vlan-4)# spanning-tree priority 500
FastIron-A(config-vlan-4)#
```

Configuring FastIron-B

Enter the following commands to configure FastIron-B:

```
FastIron# config t
FastIron(config)# host FastIron-B
FastIron-B(config)# vlan 2 name IP_IPX_Protocol
FastIron-B(config-vlan-2)# untag e1 to 8
FastIron-B(config-vlan-2)# tag e25 to 26
FastIron-B(config-vlan-2)# spanning-tree
FastIron-B(config-vlan-2)# ip-PROTO name Red
FastIron-B(config-vlan-ip-PROTO)# no dynamic
FastIron-B(config-vlan-ip-PROTO)# static e1 to 4 e25 to 26
FastIron-B(config-vlan-ip-PROTO)# exclude e5 to 8
FastIron-B(config-vlan-ip-PROTO)# ipx-PROTO name Blue
FastIron-B(config-vlan-ipx-PROTO)# no dynamic
FastIron-B(config-vlan-ipx-PROTO)# static e5 to 8 e25 to 26
```

```

FastIron-B(config-vlan-ipx-proto)# exclude e1 to 4
FastIron-B(config-vlan-other-proto)# vlan 3 name IP-Sub_IPX-Net_VLANs
FastIron-B(config-vlan-3)# untag e9 to 16
FastIron-B(config-vlan-3)# tag e25 to 26
FastIron-B(config-vlan-3)# spanning-tree
FastIron-B(config-vlan-3)# spanning-tree priority 500
FastIron-B(config-vlan-3)# ip-sub 1.1.1.0/24 name Green
FastIron-B(config-vlan-ip-subnet)# no dynamic
FastIron-B(config-vlan-ip-subnet)# static e9 to 12 e25 to 26
FastIron-B(config-vlan-ip-subnet)# exclude e13 to 16
FastIron-B(config-vlan-ip-subnet)# ipx-net 1 ethernet_802.3 name Brown
FastIron-B(config-vlan-ipx-network)# no dynamic
FastIron-B(config-vlan-ipx-network)# static e13 to 16 e25 to 26
FastIron-B(config-vlan-ipx-network)# exclude e9 to 12
FastIron-B(config-vlan-ipx-network)# vlan 4 name Purple_ALL-Protocols
FastIron-B(config-vlan-4)# untag e17 to 24
FastIron-B(config-vlan-4)# tag e25 to 26
FastIron-B(config-vlan-4)# spanning-tree

```

Configuring FastIron-C

Enter the following commands to configure FastIron-C:

```

FastIron# config t
FastIron(config)# host FastIron-C
FastIron-C(config)# vlan 2 name IP_IPX_Protocol
FastIron-C(config-vlan-2)# untag e1 to 8
FastIron-C(config-vlan-2)# tag e25 to 26
FastIron-C(config-vlan-2)# spanning-tree
FastIron-C(config-vlan-2)# ip-proto name Red
FastIron-C(config-vlan-ip-proto)# no dynamic
FastIron-C(config-vlan-ip-proto)# static e1 to 4 e25 to 26
FastIron-C(config-vlan-ip-proto)# exclude e5 to 8
FastIron-C(config-vlan-ip-proto)# ipx-proto name Blue
FastIron-C(config-vlan-ipx-proto)# no dynamic
FastIron-C(config-vlan-ipx-proto)# static e5 to 8 e25 to 26
FastIron-C(config-vlan-ipx-proto)# exclude e1 to 4
FastIron-C(config-vlan-other-proto)# vlan 3 name IP-Sub_IPX-Net_VLANs
FastIron-C(config-vlan-3)# untag e9 to 16
FastIron-C(config-vlan-3)# tag e25 to 26
FastIron-C(config-vlan-3)# spanning-tree
FastIron-C(config-vlan-3)# ip-sub 1.1.1.0/24 name Green
FastIron-C(config-vlan-ip-subnet)# no dynamic
FastIron-C(config-vlan-ip-subnet)# static e9 to 12 e25 to 26
FastIron-C(config-vlan-ip-subnet)# exclude e13 to 16
FastIron-C(config-vlan-ip-subnet)# ipx-net 1 ethernet_802.3 name Brown
FastIron-C(config-vlan-ipx-network)# no dynamic
FastIron-C(config-vlan-ipx-network)# static e13 to 16 e25 to 26
FastIron-C(config-vlan-ipx-network)# exclude e9 to 12
FastIron-C(config-vlan-ipx-network)# vlan 4 name Purple_ALL-Protocols
FastIron-C(config-vlan-4)# untag e17 to 24
FastIron-C(config-vlan-4)# tag e25 to 26
FastIron-C(config-vlan-4)# spanning-tree

```

Configuring an IPv6 Protocol VLAN

You can configure a protocol-based VLAN as a broadcast domain for IPv6 traffic. When the Layer 3 Switch receives an IPv6 multicast packet (a packet with 06 in the version field and 0xFF as the beginning of the destination address), the Layer 3 Switch forwards the packet to all other ports in the VLAN.

NOTE: The Layer 3 Switch forwards all IPv6 multicast packets to all ports in the VLAN except the port that received the packet, and does not distinguish among subnet directed multicasts.

You can add the VLAN ports as static ports or dynamic ports. A static port is always an active member of the VLAN. Dynamic ports within any protocol VLAN age out after 10 minutes, if no member protocol traffic is received on a port within the VLAN. The aged out port, however, remains as a candidate dynamic port for that VLAN. The port becomes active in the VLAN again if member protocol traffic is received on that port.

Once a port is re-activated, the aging out period for the port is reset to 20 minutes. Each time a member protocol packet is received by a candidate dynamic port (aged out port) the port becomes active again and the aging out period is reset for 20 minutes.

To configure an IPv6 VLAN, enter commands such as the following:

```
BigIron(config)# vlan 2
BigIron(config-vlan-2)# untag ethernet 1/1 to 1/8
BigIron(config-vlan-2)# ipv6-proto name V6
BigIron(config-ipv6-subnet)# static ethernet 1/1 to 1/6
BigIron(config-ipv6-subnet)# dynamic
```

The first two commands configure a port-based VLAN and add ports 1/1 – 1/8 to the VLAN. The remaining commands configure an IPv6 VLAN within the port-based VLAN. The **static** command adds ports 1/1 – 1/6 as static ports, which do not age out. The **dynamic** command adds the remaining ports, 1/7 – 1/8, as dynamic ports. These ports are subject to aging as described above.

Syntax: [no] ipv6-proto [name <string>]

Routing Between VLANs Using Virtual Routing Interfaces (Layer 3 Switches Only)


Foundry Layer 3 Switches offer the ability to create a virtual routing interface within a Layer 2 STP port-based VLAN or within each Layer 3 protocol, IP subnet, or IPX network VLAN. This combination of multiple Layer 2 and/or Layer 3 broadcast domains and virtual routing interfaces are the basis for Foundry Networks' very powerful Integrated Switch Routing (ISR) technology. ISR is very flexible and can solve many networking problems. The following example is meant to provide ideas by demonstrating some of the concepts of ISR.

Example: Suppose you want to move routing out to each of three buildings in a network. Remember that the only protocols present on VLAN 2 and VLAN 3 are IP and IPX. Therefore, you can eliminate tagged ports 25 and 26 from both VLAN 2 and VLAN 3 and create new tagged port-based VLANs to support separate IP subnets and IPX networks for each backbone link.

You also need to create unique IP subnets and IPX networks within VLAN 2 and VLAN 3 at each building. This will create a fully routed IP and IPX backbone for VLAN 2 and VLAN 3. However, VLAN 4 has no protocol restrictions across the backbone. In fact there are requirements for NetBIOS and DecNet to be bridged among the three building locations. The IP subnet and IPX network that exists within VLAN 4 must remain a flat Layer 2 switched STP domain. You enable routing for IP and IPX on a virtual routing interface only on NetIron-A. This will provide the flat IP and IPX segment with connectivity to the rest of the network. Within VLAN 4 IP and IPX will follow the STP topology. All other IP subnets and IPX networks will be fully routed and have use of all paths at all times during normal operation.

Figure 22.13 shows the configuration described above.

Figure 22.13 Routing between protocol-based VLANs



To configure the Layer 3 VLANs and virtual routing interfaces on the NetIron Layer 3 Switch in Figure 22.13, use the following procedure.

USING THE CLI

Configuring NetIron-A

Enter the following commands to configure NetIron-A. The following commands enable OSPF or RIP routing and IPX routing.

```

NetIron> en
No password has been assigned yet...
NetIron# configure terminal
NetIron(config)# hostname NetIron-A
NetIron-A(config)# router ospf
NetIron-A(config-ospf-router)# area 0.0.0.0 normal
NetIron-A(config-ospf-router)# router ipx
ipx routing enabled for next power cycle.
Please save configuration to flash and reboot.
NetIron-A(config-ospf-router)#
    
```

The following commands create the port-based VLAN 2. In the previous example, an external NetIron defined the router interfaces for VLAN 2. With ISR, routing for VLAN 2 is done locally within each NetIron. Therefore, there are two ways you can solve this problem. One way is to create a unique IP subnet and IPX network VLAN, each with its own virtual routing interface and unique IP or IPX address within VLAN 2 on each NetIron. In this example, this is the configuration used for VLAN 3. The second way is to split VLAN 2 into two separate port-based VLANs and create a virtual router interface within each port-based VLAN. Later in this example, this second option is used to create a port-based VLAN 8 to show that there are multiple ways to accomplish the same task with ISR.

You also need to create the Other-Protocol VLAN within port-based VLAN 2 and 8 to prevent unwanted protocols from being Layer 2 switched within port-based VLAN 2 or 8. Note that the only port-based VLAN that requires STP in this example is VLAN 4. You will need to configure the rest of the network to prevent the need to run STP.

```
NetIron-A(config-ospf-router)# vlan 2 name IP-Subnet_1.1.2.0/24
NetIron-A(config-vlan-2)# untag e 1 to 4
NetIron-A(config-vlan-2)# no spanning-tree
NetIron-A(config-vlan-2)# router-interface ve1
NetIron-A(config-vlan-2)# other-proto name block_other_protocols
NetIron-A(config-vlan-other-proto)# no dynamic
NetIron-A(config-vlan-other-proto)# exclude e 1 to 4
```

Once you have defined the port-based VLAN and created the virtual routing interface, you need to configure the virtual routing interface just as you would configure a physical interface.

```
NetIron-A(config-vlan-other-proto)# interface ve1
NetIron-A(config-vif-1)# ip address 1.1.2.1/24
NetIron-A(config-vif-1)# ip ospf area 0.0.0.0
```

Do the same thing for VLAN 8.

```
NetIron-A(config-vif-1)# vlan 8 name IPX_Network2
NetIron-A(config-vlan-8)# untag ethernet 5 to 8
NetIron-A(config-vlan-8)# no spanning-tree
NetIron-A(config-vlan-8)# router-interface ve 2
NetIron-A(config-vlan-8)# other-proto name block-other-protocols
NetIron-A(config-vlan-other-proto)# no dynamic
NetIron-A(config-vlan-other-proto)# exclude ethernet 5 to 8
NetIron-A(config-vlan-other-proto)# int ve2
NetIron-A(config-vif-2)# ipx network 2 ethernet_802.3
NetIron-A(config-vif-2)#
```

The next thing you need to do is create VLAN 3. This is very similar to the previous example with the addition of virtual routing interfaces to the IP subnet and IPX network VLANs. Also there is no need to exclude ports from the IP subnet and IPX network VLANs on the router.

```
NetIron-A(config-vif-2)# vlan 3 name IP_Sub_&_IPX_Net_VLAN
NetIron-A(config-vlan-3)# untag e 9 to 16
NetIron-A(config-vlan-3)# no spanning-tree
NetIron-A(config-vlan-3)# ip-subnet 1.1.1.0/24
NetIron-A(config-vlan-ip-subnet)# static e 9 to 12
NetIron-A(config-vlan-ip-subnet)# router-interface ve3
NetIron-A(config-vlan-ip-subnet)# ipx-network 1 ethernet_802.3
NetIron-A(config-vlan-ipx-network)# static e 13 to 16
NetIron-A(config-vlan-ipx-network)# router-interface ve4
NetIron-A(config-vlan-ipx-network)# other-proto name block-other-protocols
NetIron-A(config-vlan-other-proto)# exclude e 9 to 16
NetIron-A(config-vlan-other-proto)# no dynamic
NetIron-A(config-vlan-other-proto)# interface ve 3
NetIron-A(config-vif-3)# ip addr 1.1.1.1/24
NetIron-A(config-vif-3)# ip ospf area 0.0.0.0
NetIron-A(config-vif-3)# int ve4
NetIron-A(config-vif-4)# ipx network 1 ethernet_802.3
NetIron-A(config-vif-4)#
```

Now configure VLAN 4. Remember this is a flat segment that, in the previous example, obtained its IP default gateway and IPX router services from an external NetIron. In this example, NetIron-A will provide the routing services for VLAN 4. You also want to configure the STP priority for VLAN 4 to make NetIron-A the root bridge for this VLAN.

```
NetIron-A(config-vif-4)# vlan 4 name Bridged_ALL_Protocols
NetIron-A(config-vlan-4)# untag ethernet 17 to 24
NetIron-A(config-vlan-4)# tag ethernet 25 to 26
```

```

NetIron-A(config-vlan-4)# spanning-tree
NetIron-A(config-vlan-4)# spanning-tree priority 500
NetIron-A(config-vlan-4)# router-interface ve5
NetIron-A(config-vlan-4)# int ve5
NetIron-A(config-vif-5)# ip address 1.1.3.1/24
NetIron-A(config-vif-5)# ip ospf area 0.0.0.0
NetIron-A(config-vif-5)# ipx network 3 ethernet_802.3
NetIron-A(config-vif-5)#

```

It is time to configure a separate port-based VLAN for each of the routed backbone ports (Ethernet 25 and 26). If you do not create a separate tagged port-based VLAN for each point-to-point backbone link, you need to include tagged interfaces for Ethernet 25 and 26 within VLANs 2, 3, and 8. This type of configuration makes the entire backbone a single STP domain for each VLAN 2, 3, and 8. This is the configuration used in the example in "Configuring IP Subnet, IPX Network and Protocol-Based VLANs" on page 22-23. In this scenario, the virtual routing interfaces within port-based VLANs 2, 3, and 8 will be accessible using only one path through the network. The path that is blocked by STP is not available to the routing protocols until it is in the STP FORWARDING state.

```

NetIron-A(config-vif-5)# vlan 5 name Rtr_BB_to_Bldg.2
NetIron-A(config-vlan-5)# tag e 25
NetIron-A(config-vlan-5)# no spanning-tree
NetIron-A(config-vlan-5)# router-interface ve6
NetIron-A(config-vlan-5)# vlan 6 name Rtr_BB_to_Bldg.3
NetIron-A(config-vlan-6)# tag ethernet 26
NetIron-A(config-vlan-6)# no spanning-tree
NetIron-A(config-vlan-6)# router-interface ve7
NetIron-A(config-vlan-6)# int ve6
NetIron-A(config-vif-6)# ip addr 1.1.4.1/24
NetIron-A(config-vif-6)# ip ospf area 0.0.0.0
NetIron-A(config-vif-6)# ipx network 4 ethernet_802.3
NetIron-A(config-vif-6)# int ve7
NetIron-A(config-vif-7)# ip addr 1.1.5.1/24
NetIron-A(config-vif-7)# ip ospf area 0.0.0.0
NetIron-A(config-vif-7)# ipx network 5 ethernet_802.3
NetIron-A(config-vif-7)#

```

This completes the configuration for NetIron-A. The configuration for NetIron-B and C is very similar except for a few issues.

- IP subnets and IPX networks configured on NetIron-B and NetIron-C must be unique across the entire network, except for the backbone port-based VLANs 5, 6, and 7 where the subnet is the same but the IP address must change.
- There is no need to change the default priority of STP within VLAN 4.
- There is no need to include a virtual router interface within VLAN 4.
- The backbone VLAN between NetIron-B and NetIron-C must be the same at both ends and requires a new VLAN ID. The VLAN ID for this port-based VLAN is VLAN 7.

Configuration for NetIron-B

Enter the following commands to configure NetIron-B.

```

NetIron> en
No password has been assigned yet...
NetIron# config t
NetIron(config)# hostname NetIron-B
NetIron-B(config)# router ospf
NetIron-B(config-ospf-router)# area 0.0.0.0 normal
NetIron-B(config-ospf-router)# router ipx
NetIron-B(config-ospf-router)# vlan 2 name IP-Subnet_1.1.6.0/24
NetIron-B(config-vlan-2)# untag e 1 to 4
NetIron-B(config-vlan-2)# no spanning-tree

```

```
NetIron-B(config-vlan-2)# router-interface ve1
NetIron-B(config-vlan-2)# other-proto name block-other-protocols
NetIron-B(config-vlan-other-proto)# no dynamic
NetIron-B(config-vlan-other-proto)# exclude e 1 to 4
NetIron-B(config-vlan-other-proto)# int ve1
NetIron-B(config-vif-1)# ip addr 1.1.6.1/24
NetIron-B(config-vif-1)# ip ospf area 0.0.0.0
NetIron-B(config-vif-1)# vlan 8 name IPX_Network6
NetIron-B(config-vlan-8)# untag e 5 to 8
NetIron-B(config-vlan-8)# no span
NetIron-B(config-vlan-8)# router-int ve2
NetIron-B(config-vlan-8)# other-proto name block-other-protocols
NetIron-B(config-vlan-other-proto)# no dynamic
NetIron-B(config-vlan-other-proto)# exclude e 5 to 8
NetIron-B(config-vlan-other-proto)# int ve2
NetIron-B(config-vif-2)# ipx net 6 ethernet_802.3
NetIron-B(config-vif-2)# vlan 3 name IP_Sub_&_IPX_Net_VLAN
NetIron-B(config-vlan-3)# untag e 9 to 16
NetIron-B(config-vlan-3)# no spanning-tree
NetIron-B(config-vlan-3)# ip-subnet 1.1.7.0/24
NetIron-B(config-vlan-ip-subnet)# static e 9 to 12
NetIron-B(config-vlan-ip-subnet)# router-interface ve3
NetIron-B(config-vlan-ip-subnet)# ipx-network 7 ethernet_802.3
NetIron-B(config-vlan-ipx-network)# static e 13 to 16
NetIron-B(config-vlan-ipx-network)# router-interface ve4
NetIron-B(config-vlan-ipx-network)# other-proto name block-other-protocols
NetIron-B(config-vlan-other-proto)# exclude e 9 to 16
NetIron-B(config-vlan-other-proto)# no dynamic
NetIron-B(config-vlan-other-proto)# interface ve 3
NetIron-B(config-vif-3)# ip addr 1.1.7.1/24
NetIron-B(config-vif-3)# ip ospf area 0.0.0.0
NetIron-B(config-vif-3)# int ve4
NetIron-B(config-vif-4)# ipx network 7 ethernet_802.3
NetIron-B(config-vif-4)# vlan 4 name Bridged_ALL_Protocols
NetIron-B(config-vlan-4)# untag ethernet 17 to 24
NetIron-B(config-vlan-4)# tag ethernet 25 to 26
NetIron-B(config-vlan-4)# spanning-tree
NetIron-B(config-vlan-4)# vlan 5 name Rtr_BB_to_Bldg.1
NetIron-B(config-vlan-5)# tag e 25
NetIron-B(config-vlan-5)# no spanning-tree
NetIron-B(config-vlan-5)# router-interface ve5
NetIron-B(config-vlan-5)# vlan 7 name Rtr_BB_to_Bldg.3
NetIron-B(config-vlan-7)# tag ethernet 26
NetIron-B(config-vlan-7)# no spanning-tree
NetIron-B(config-vlan-7)# router-interface ve6
NetIron-B(config-vlan-7)# int ve5
NetIron-B(config-vif-5)# ip addr 1.1.4.2/24
NetIron-B(config-vif-5)# ip ospf area 0.0.0.0
NetIron-B(config-vif-5)# ipx network 4 ethernet_802.3
NetIron-B(config-vif-5)# int ve6
NetIron-B(config-vif-6)# ip addr 1.1.8.1/24
NetIron-B(config-vif-6)# ip ospf area 0.0.0.0
NetIron-B(config-vif-6)# ipx network 8 ethernet_802.3
NetIron-B(config-vif-6)#
```

Configuration for NetIron-C

Enter the following commands to configure NetIron-C.

```
NetIron> en
No password has been assigned yet...
NetIron# config t
NetIron-C(config)# hostname NetIron-C
NetIron-C(config)# router ospf
NetIron-C(config-ospf-router)# area 0.0.0.0 normal
NetIron-C(config-ospf-router)# router ipx
NetIron-C(config-ospf-router)# vlan 2 name IP-Subnet_1.1.9.0/24
NetIron-C(config-vlan-2)# untag e 1 to 4
NetIron-C(config-vlan-2)# no spanning-tree
NetIron-C(config-vlan-2)# router-interface ve1
NetIron-C(config-vlan-2)# other-proto name block-other-protocols
NetIron-C(config-vlan-other-proto)# no dynamic
NetIron-C(config-vlan-other-proto)# exclude e 1 to 4
NetIron-C(config-vlan-other-proto)# int ve1
NetIron-C(config-vif-1)# ip addr 1.1.9.1/24
NetIron-C(config-vif-1)# ip ospf area 0.0.0.0
NetIron-C(config-vif-1)# vlan 8 name IPX_Network9
NetIron-C(config-vlan-8)# untag e 5 to 8
NetIron-C(config-vlan-8)# no span
NetIron-C(config-vlan-8)# router-int ve2
NetIron-C(config-vlan-8)# other-proto name block-other-protocols
NetIron-C(config-vlan-other-proto)# no dynamic
NetIron-C(config-vlan-other-proto)# exclude e 5 to 8
NetIron-C(config-vlan-other-proto)# int ve2
NetIron-C(config-vif-2)# ipx net 9 ethernet_802.3
NetIron-C(config-vif-2)# vlan 3 name IP_Sub_&_IPX_Net_VLAN
NetIron-C(config-vlan-3)# untag e 9 to 16
NetIron-C(config-vlan-3)# no spanning-tree
NetIron-C(config-vlan-3)# ip-subnet 1.1.10.0/24
NetIron-C(config-vlan-ip-subnet)# static e 9 to 12
NetIron-C(config-vlan-ip-subnet)# router-interface ve3
NetIron-C(config-vlan-ip-subnet)# ipx-network 10 ethernet_802.3
NetIron-C(config-vlan-ipx-network)# static e 13 to 16
NetIron-C(config-vlan-ipx-network)# router-interface ve4
NetIron-C(config-vlan-ipx-network)# other-proto name block-other-protocols
NetIron-C(config-vlan-other-proto)# exclude e 9 to 16
NetIron-C(config-vlan-other-proto)# no dynamic
NetIron-C(config-vlan-other-proto)# interface ve 3
NetIron-C(config-vif-3)# ip addr 1.1.10.1/24
NetIron-C(config-vif-3)# ip ospf area 0.0.0.0
NetIron-C(config-vif-3)# int ve4
NetIron-C(config-vif-4)# ipx network 10 ethernet_802.3
NetIron-C(config-vif-4)# vlan 4 name Bridged_ALL_Protocols
NetIron-C(config-vlan-4)# untag ethernet 17 to 24
NetIron-C(config-vlan-4)# tag ethernet 25 to 26
NetIron-C(config-vlan-4)# spanning-tree
NetIron-C(config-vlan-4)# vlan 7 name Rtr_BB_to_Bldg.2
NetIron-C(config-vlan-7)# tag e 25
NetIron-C(config-vlan-7)# no spanning-tree
NetIron-C(config-vlan-7)# router-interface ve5
NetIron-C(config-vlan-7)# vlan 6 name Rtr_BB_to_Bldg.1
NetIron-C(config-vlan-6)# tag ethernet 26
NetIron-C(config-vlan-6)# no spanning-tree
NetIron-C(config-vlan-6)# router-interface ve6
```



```
NetIron-C(config-vlan-6)# int ve5
NetIron-C(config-vif-5)# ip addr 1.1.8.2/24
NetIron-C(config-vif-5)# ip ospf area 0.0.0.0
NetIron-C(config-vif-5)# ipx network 8 ethernet_802.3
NetIron-C(config-vif-5)# int ve6
NetIron-C(config-vif-6)# ip addr 1.1.5.2/24
NetIron-C(config-vif-6)# ip ospf area 0.0.0.0
NetIron-C(config-vif-6)# ipx network 5 ethernet_802.3
NetIron-C(config-vif-6)#
```

Configuring AppleTalk Cable VLANs

You can configure up to eight AppleTalk cable VLANs within a port-based VLAN.

NOTE: This feature applies only to Chassis Layer 3 Switches and the Turbolron/8.

To configure an AppleTalk cable VLAN, you create a port-based VLAN, then create up to eight cable VLANs within the port-based VLAN. You create the AppleTalk cable VLAN by assigning a number to the VLAN, optionally naming the cable VLAN, assigning ports from the port-based VLAN, and specifying the router interface (virtual routing interface) on which the Layer 3 Switch will send and receive traffic for the cable VLAN.

All the ports in an AppleTalk cable VLAN are within the same AppleTalk cable range. The device switches traffic within the VLAN and routes traffic between VLANs.

Configuration Guidelines


Use the following guidelines when configuring AppleTalk cable VLANs:

- The number of VLANs you can configure is limited by the system max number configured for protocol VLANs.
- Each AppleTalk cable VLAN can have only one router interface. The router interface must be a virtual routing interface.
- The AppleTalk cable VLANs cannot overlap. Thus, you cannot use the same port in more than one AppleTalk cable VLAN.
- You must add the ports to the AppleTalk cable VLAN using the static option. You cannot use the dynamic or exclude options.
- You cannot have an AppleTalk cable VLAN and an AppleTalk protocol VLAN in the same port-based VLAN. If you already have an AppleTalk protocol VLAN in the port-based VLAN, you must delete the AppleTalk protocol VLAN first, then configure the AppleTalk cable VLAN.

Configuration Example

Figure 22.14 shows an example of a BigIron 8000 Layer 3 Switch with four AppleTalk cable VLANs configured on a single port-based VLAN. In this example, port-based VLAN 10 is configured, then AppleTalk cable VLANs are configured on ports on chassis modules 2 and 3. Each virtual routing interface (ve1, ve2, ve3, and ve4) is then configured with AppleTalk routing information for the cable VLAN.

Figure 22.14 AppleTalk Cable VLANs



Configuring the VLANs

To configure the VLANs shown in Figure 3, enter the following CLI commands:

```
BigIron(config)# vlan 10 by port
BigIron(config-vlan-10)# untag ethe 2/1 to 2/2 ethe 3/1 to 3/8
```

The two commands above add port-based VLAN 10 and add ports 2/1, 2/2, and 3/1 – 3/16 to the VLAN. The **untag** command removes ports from the default VLAN and adds them to port-based VLAN 10. (The default VLAN contains all the ports in the system by default.) The **untag** command also allows the ports to process packets that do not contain 802.1q tagging.

The following commands add four AppleTalk cable VLANs, in groups of three commands each. The **appletalk-cable-vlan** command adds a cable VLAN and, with the optional **name** parameter, names the VLAN. The **static** command adds specific ports within the port-based VLAN to the AppleTalk cable VLAN. The **router-interface** command identifies virtual routing interface that connects to the AppleTalk cable range the VLAN is for.

```
BigIron(config-vlan-10)# appletalk-cable-vlan 1 name cable-one
BigIron(config-vlan-10)# static ethe 2/1 to 2/2 ethe 3/1 to 3/2
BigIron(config-vlan-10)# router-interface ve 1
BigIron(config-vlan-10)# appletalk-cable-vlan 2 name cable-two
BigIron(config-vlan-10)# static ethe 3/3 to 3/4
BigIron(config-vlan-10)# router-interface ve 2
```

```

BigIron(config-vlan-10)# appletalk-cable-vlan 3 name cable-three
BigIron(config-vlan-10)# static ethe 3/5 to 3/6
BigIron(config-vlan-10)# router-interface ve 3
BigIron(config-vlan-10)# appletalk-cable-vlan 4 name cable-four
BigIron(config-vlan-10)# static ethe 3/7 to 3/8
BigIron(config-vlan-10)# router-interface ve 4

```

Syntax: appletalk-cable-vlan <vlan-id> [name <string>]

The <vlan-id> can be from 1 – 8.

The **name** <string> parameter specifies a name and can be a string up to 32 characters long.

Configuring the Router Interfaces

The following commands configure the router interfaces (virtual routing interfaces) associated with the AppleTalk cable VLANs. The **interface ve** commands add the virtual routing interfaces to the system. (The **router-interface** commands above refer to these interfaces but do not add them. You must add the interfaces using the **interface ve** command.)

For each virtual routing interface, additional commands configure the AppleTalk routing parameters for the interface. Notice that each virtual routing interface has a separate set of routing parameters. The routing parameters on each virtual routing interface are independent of the routing parameters on other virtual routing interfaces. Since each AppleTalk cable VLAN is associated with a separate virtual routing interface, each AppleTalk cable VLAN has a distinct set of routing parameters, separate from the routing parameters on other AppleTalk VLANs. In effect, each virtual routing interface contains a separate AppleTalk router.

The **appletalk address** command configures the AppleTalk interface address on the virtual routing interface. The **appletalk cable-range** command specifies the cable range for the network. The **appletalk routing** command enables AppleTalk routing on the virtual routing interface. The **zone-name** commands add zones to the network. For information about the AppleTalk routing commands, see the “Configuring AppleTalk” chapter in the *Foundry Enterprise Configuration and Management Guide*.

The **write memory** command at the end of the example saves the configuration to the startup-config file.

```

BigIron(config-vlan-10)# interface ve 1
BigIron(config-vif-1)# appletalk cable-range 10 - 19
BigIron(config-vif-1)# appletalk address 10.1
BigIron(config-vif-1)# appletalk zone-name AA
BigIron(config-vif-1)# appletalk routing
BigIron(config-vif-1)# interface ve 2
BigIron(config-vif-2)# appletalk cable-range 20 - 29
BigIron(config-vif-2)# appletalk address 20.1
BigIron(config-vif-2)# appletalk zone-name BB
BigIron(config-vif-2)# appletalk routing
BigIron(config-vif-2)# interface ve 3
BigIron(config-vif-3)# appletalk cable-range 30 - 39
BigIron(config-vif-3)# appletalk address 30.1
BigIron(config-vif-3)# appletalk zone-name CC
BigIron(config-vif-3)# appletalk routing
BigIron(config-vif-3)# interface ve 4
BigIron(config-vif-4)# appletalk cable-range 40 - 49
BigIron(config-vif-4)# appletalk address 40.1
BigIron(config-vif-4)# appletalk zone-name DD
BigIron(config-vif-4)# appletalk routing
BigIron(config-vif-4)# write memory

```

Configuring Protocol VLANs With Dynamic Ports

The configuration examples for protocol VLANs in the sections above show how to configure the VLANs using static ports. You also can configure the following types of protocol VLANs with dynamic ports:

- AppleTalk protocol
- IP protocol
- IPX protocol
- IP subnet
- IPX network

NOTE: The software does not support dynamically adding ports to AppleTalk cable VLANs. Conceptually, an AppleTalk cable VLAN consists of a single network cable, connected to a single port. Therefore, dynamic addition and removal of ports is not applicable.

NOTE: This feature applies only to Chassis Layer 3 Switches and the Turbolron/8.

NOTE: You cannot route to or from protocol VLANs with dynamically added ports.

Aging of Dynamic Ports

When you add the ports to the VLAN, the software automatically adds them all to the VLAN. However, dynamically added ports age out. If the age time for a dynamic port expires, the software removes the port from the VLAN. If that port receives traffic for the IP subnet or IPX network, the software adds the port to the VLAN again and starts the aging timer over. Each time the port receives traffic for the VLAN's IP subnet or IPX network, the aging timer starts over.

Dynamic ports within any protocol VLAN age out after 10 minutes, if no member protocol traffic is received on a port within the VLAN. The aged out port, however, remains as a candidate dynamic port for that VLAN. The port becomes active in the VLAN again if member protocol traffic is received on that port.

Once a port is re-activated, the aging out period for the port is reset to 20 minutes. Each time a member protocol packet is received by a candidate dynamic port (aged out port) the port becomes active again and the aging out period is reset for 20 minutes.

Configuration Guidelines

- You cannot dynamically add a port to a protocol VLAN if the port has any routing configuration parameters. For example, the port cannot have a virtual routing interface, IP subnet address, IPX network address, or AppleTalk network address configured on it.
- Once you dynamically add a port to a protocol VLAN, you cannot configure routing parameters on the port.
- Dynamic VLAN ports are not required or supported on AppleTalk cable VLANs.

Configuring an IP, IPX, or AppleTalk Protocol VLAN with Dynamic Ports

To configure an IP, IPX, or AppleTalk protocol VLAN with dynamic ports, use one of the following methods.

USING THE CLI

To configure port-based VLAN 10, then configure an IP protocol VLAN within the port-based VLAN with dynamic ports, enter the following commands such as the following:

```
BigIron(config)# vlan 10 by port
BigIron(config-vlan-10)# untag ethernet 1/1 to 1/6
added untagged port ethe 1/1 to 1/6 to port-vlan 30.
BigIron(config-vlan-10)# ip-proto name IP_Prot_VLAN
```

```
BigIron(config-vlan-10)# dynamic
BigIron(config)# write memory
```

Syntax: vlan <vlan-id> by port [name <string>]

Syntax: untagged ethernet <portnum> to <portnum>

Or

Syntax: untagged ethernet <portnum> ethernet <portnum>

NOTE: Use the first **untagged** command for adding a range of ports. Use the second command for adding separate ports (not in a range).

Syntax: ip-proto [name <string>]

Syntax: ipx-proto [name <string>]

Syntax: appletalk-cable-vlan <num> [name <string>]

Syntax: dynamic

The procedure is similar for IPX and AppleTalk protocol VLANs. Enter **ipx-proto** or **atalk-proto** instead of **ip-proto**.

Configuring an IP Subnet VLAN with Dynamic Ports

To configure an IP subnet VLAN with dynamic ports, use one of the following methods.

USING THE CLI

To configure port-based VLAN 10, then configure an IP subnet VLAN within the port-based VLAN with dynamic ports, enter commands such as the following:

```
BigIron(config)# vlan 10 by port name IP_VLAN
BigIron(config-vlan-10)# untag ethernet 1/1 to 1/6
added untagged port ethe 1/1 to 1/6 to port-vlan 10.
BigIron(config-vlan-10)# ip-subnet 1.1.1.0/24 name Mktg-LAN
BigIron(config-vlan-10)# dynamic
BigIron(config)# write memory
```

These commands create a port-based VLAN on chassis ports 1/1 – 1/6 named “Mktg-LAN”, configure an IP subnet VLAN within the port-based VLAN, and then add ports from the port-based VLAN dynamically.

Syntax: vlan <vlan-id> by port [name <string>]

Syntax: untagged ethernet <portnum> to <portnum>

Or

Syntax: untagged ethernet <portnum> ethernet <portnum>

NOTE: Use the first **untagged** command for adding a range of ports. Use the second command for adding separate ports (not in a range).

Syntax: ip-subnet <ip-addr> <ip-mask> [name <string>]

Or

Syntax: ip-subnet <ip-addr>/<mask-bits> [name <string>]

Syntax: dynamic

Configuring an IPX Network VLAN with Dynamic Ports

To configure an IPX network VLAN with dynamic ports, use one of the following methods.

USING THE CLI

To configure port-based VLAN 20, then configure an IPX network VLAN within the port-based VLAN with dynamic ports, enter commands such as the following:

```
BigIron(config)# vlan 20 by port name IPX_VLAN
BigIron(config-vlan-10)# untag ethernet 2/1 to 2/6
added untagged port ethe 2/1 to 2/6 to port-vlan 20.
BigIron(config-vlan-10)# ipx-network abcd ethernet_ii name Eng-LAN
BigIron(config-vlan-10)# dynamic
BigIron(config)# write memory
```

These commands create a port-based VLAN on chassis ports 2/1 – 2/6 named “Eng-LAN”, configure an IPX network VLAN within the port-based VLAN, and then add ports from the port-based VLAN dynamically.

Syntax: vlan <vlan-id> by port [name <string>]

Syntax: untagged ethernet <portnum> to <portnum>

Or

Syntax: untagged ethernet <portnum> ethernet <portnum>

NOTE: Use the first **untagged** command for adding a range of ports. Use the second command for adding separate ports (not in a range).

Syntax: ipx-network <network-addr> ethernet_ii | ethernet_802.2 | ethernet_802.3 | ethernet_snap [name <string>]

Syntax: dynamic

Configuring Uplink Ports Within a Port-Based VLAN

You can configure a subset of the ports in a port-based VLAN as uplink ports. When you configure uplink ports in a port-based VLAN, the device sends all broadcast and unknown-unicast traffic from a port in the VLAN to the uplink ports, but not to other ports within the VLAN. Thus, the uplink ports provide tighter broadcast control within the VLAN.

For example, if two ports within a port-based VLAN are Gigabit ports attached to the network and the other ports in the VLAN are 10/100 ports attached to clients, you can configure the two ports attached to the network as uplink ports. In this configuration, broadcast and unknown-unicast traffic in the VLAN does not go to all ports in the VLAN. The traffic goes only to the uplink ports. The clients on the network do not receive broadcast and unknown-unicast traffic from other ports, including other clients.

To configure uplink ports in a port-based VLAN, use the following CLI method.

USING THE CLI

To configure a port-based VLAN containing uplink ports, enter commands such as the following:

```
BigIron(config)# vlan 10 by port
BigIron(config-vlan-10)# untag ethernet 1/1 to 1/24
BigIron(config-vlan-10)# untag ethernet 2/1 to 2/2
BigIron(config-vlan-10)# uplink-switch ethernet 2/1 to 2/2
```

Syntax: [no] uplink-switch ethernet <portnum> [to <portnum> | ethernet <portnum>]

In this example, 24 ports on a 10/100 module and two Gigabit ports on a Gigabit module are added to port-based VLAN 10. The two Gigabit ports are then configured as uplink ports.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure uplink ports in a port-based VLAN using the Web management interface.

Configuring the Same IP Subnet Address on Multiple Port-Based VLANs

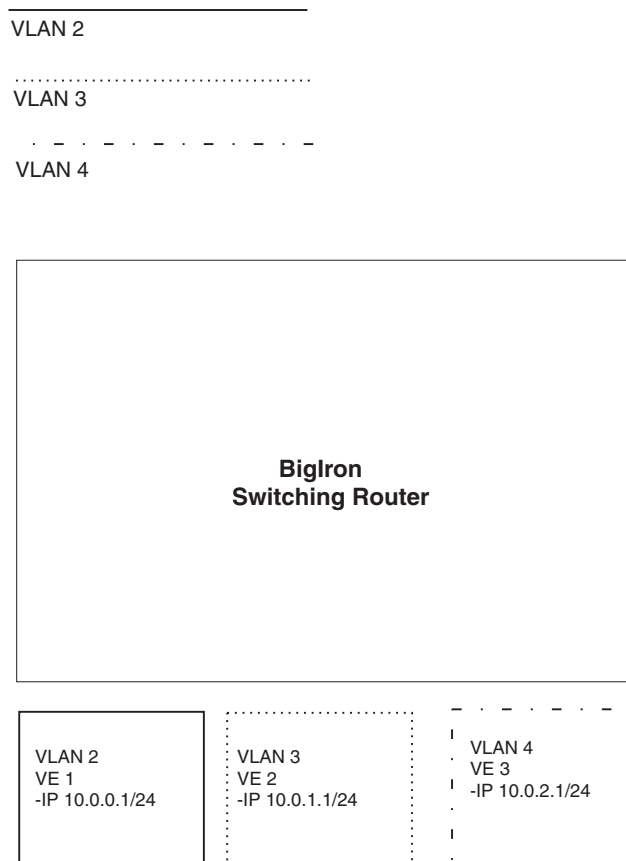
For a Foundry device to route between port-based VLANs, you must add a virtual routing interface to each VLAN. Generally, you also configure a unique IP subnet address on each virtual routing interface. For example, if you have three port-based VLANs, you add a virtual routing interface to each VLAN, then add a separate IP subnet address to each virtual routing interface. The IP address on each of the virtual routing interfaces must be in a separate subnet. The Foundry device routes Layer 3 traffic between the subnets using the subnet addresses.

NOTE: You can create virtual routing interfaces on FES running FEL or FER code, FESX running FEXR or FEXL code, FastIron running full Layer 3 or base Layer 3 code, ServerIronXLs with “ip forward” enabled, and ServerIron chassis running routing code.

NOTE: Before using the method described in this section, see “Configuring VLAN Groups and Virtual Routing Interface Groups” on page 22-45. You might be able to achieve the results you want using the methods in that section instead.

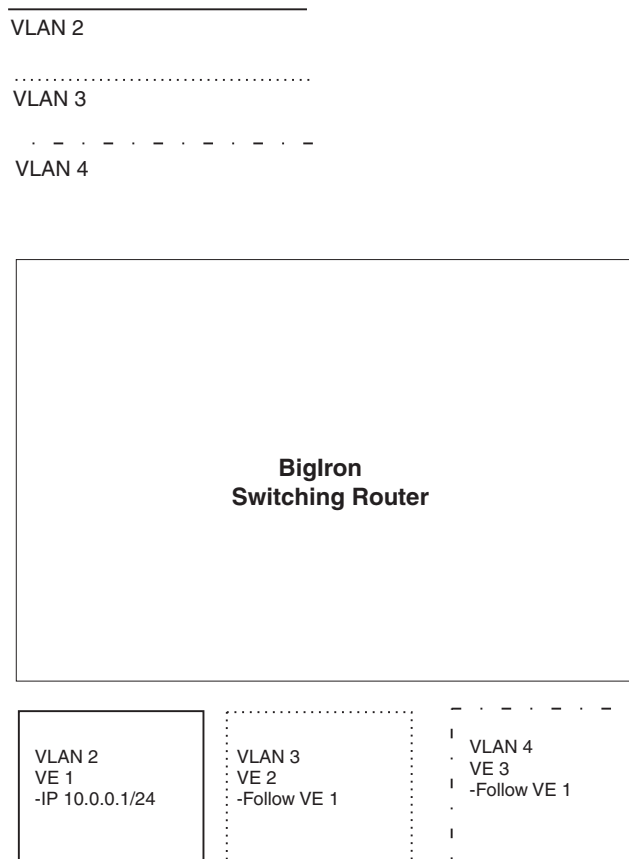
Figure 22.15 shows an example of this type of configuration.

Figure 22.15 Multiple port-based VLANs with separate protocol addresses



As shown in this example, each VLAN has a separate IP subnet address. If you need to conserve IP subnet addresses, you can configure multiple VLANs with the same IP subnet address, as shown in Figure 22.16.

Figure 22.16 Multiple port-based VLANs with the same protocol address



Each VLAN still requires a separate virtual routing interface. However, all three VLANs now use the same IP subnet address.

In addition to conserving IP subnet addresses, this feature allows containment of Layer 2 broadcasts to segments within an IP subnet. For ISP environments where the same IP subnet is allocated to different customers, placing each customer in a separate VLAN allows all customers to share the IP subnet address, while at the same time isolating them from one another's Layer 2 broadcasts.

NOTE: You can provide redundancy to an IP subnet address that contains multiple VLANs using a pair of Foundry Layer 3 Switches configured for Foundry's VRRP (Virtual Router Redundancy Protocol) or FSRP (Foundry Standby Router Protocol).

The Foundry device performs proxy Address Resolution Protocol (ARP) for hosts that want to send IP traffic to hosts in other VLANs that are sharing the same IP subnet address. If the source and destination hosts are in the same VLAN, the Foundry device does not need to use ARP.

- If a host attached to one VLAN sends an ARP message for the MAC address of a host in one of the other VLANs using the same IP subnet address, the Foundry device performs a proxy ARP on behalf of the other host. The Foundry device then replies to the ARP by sending the virtual routing interface MAC address. The Foundry device uses the same MAC address for all virtual routing interfaces.

When the host that sent the ARP then sends a unicast packet addressed to the virtual routing interface's MAC address, the device switches the packet on Layer 3 to the destination host on the VLAN.

NOTE: If the Foundry device's ARP table does not contain the requested host, the Foundry device forwards the ARP request on Layer 2 to the same VLAN as the one that received the ARP request. Then the device sends an ARP for the destination to the other VLANs that are using the same IP subnet address.

- If the destination is in the same VLAN as the source, the Foundry device does not need to perform a proxy ARP.

To configure multiple VLANs to use the same IP subnet address:

- Configure each VLAN, including adding tagged or untagged ports.
- Configure a separate virtual routing interface for each VLAN, but do not add an IP subnet address to more than one of the virtual routing interfaces.
- Configure the virtual routing interfaces that do not have the IP subnet address to “follow” the virtual routing interface that does have the address.

To configure the VLANs shown in Figure 22.16, you could enter the following commands.

```
BigIron(config)# vlan 1 by port
BigIron(config-vlan-1)# untag ethernet 1/1
BigIron(config-vlan-1)# tag ethernet 1/8
BigIron(config-vlan-1)# router-interface ve 1
```

Syntax: ip follow ve <num>

The commands above configure port-based VLAN 1. The VLAN has one untagged port (1/1) and a tagged port (1/8). In this example, all three VLANs contain port 1/8 so the port must be tagged to allow the port to be in multiple VLANs. You can configure VLANs to share a Layer 3 protocol interface regardless of tagging. A combination of tagged and untagged ports is shown in this example to demonstrate that sharing the interface does not change other VLAN features.

Notice that each VLAN still requires a unique virtual routing interface.

The following commands configure port-based VLANs 2 and 3.

```
BigIron(config-vlan-1)# vlan 2 by port
BigIron(config-vlan-2)# untag ethernet 1/2
BigIron(config-vlan-2)# tag ethernet 1/8
BigIron(config-vlan-2)# router-interface ve 2
BigIron(config-vlan-2)# vlan 3 by port
BigIron(config-vlan-3)# untag ethernet 1/5 to 1/6
BigIron(config-vlan-3)# tag ethernet 1/8
BigIron(config-vlan-3)# router-interface ve 3
```

The following commands configure an IP subnet address on virtual routing interface 1.

```
BigIron(config-vlan-3)# interface ve 1
BigIron(config-vif-1)# ip address 10.0.0.1/24
```

The following commands configure virtual routing interfaces 2 and 3 to “follow” the IP subnet address configured on virtual routing interface 1.

```
BigIron(config-vif-1)# interface ve 2
BigIron(config-vif-2)# ip follow ve 1
BigIron(config-vif-2)# interface ve 3
BigIron(config-vif-3)# ip follow ve 1
```

NOTE: Since virtual routing interfaces 2 and 3 do not have their own IP subnet addresses but instead are “following” virtual routing interface 1's IP address, you still can configure an IPX or AppleTalk interface on virtual routing interfaces 2 and 3.

Using Separate ACLs on IP Follower Virtual Routing Interfaces

NOTE: This section applies to flow-based ACLs only.

The IP follower feature allows multiple virtual routing interfaces to share the same IP address. One virtual routing interface has the IP address and the other virtual routing interfaces are configured to follow the virtual routing interface that has the address.

By default, the follower interfaces are secured by the ACLs that are applied to the interface that has the address. In fact, an ACL applied to a follower interface is ignored. For example, if you configure virtual routing interfaces 1, 2, and 3, and configure interfaces 2 and 3 to follow interface 1, then the ACLs applied to interface 1 also apply to interfaces 2 and 3. Any ACLs applied separately to interface 2 or 3 are ignored.

You can enable a follower virtual routing interface to use the ACLs you apply to it instead of using the ACLs applied to the interface that has the address. For example, you can enable virtual routing interface 2 to use its own ACLs instead of using interface 1's ACLs.

To enable a virtual routing interface to use its own ACLs instead of the ACLs of the interface it is following, enter the following command at the configuration level for the interface:

```
BigIron(config-vif-2)# no ip follow acl
```

Syntax: [no] ip follow acl

The following commands show a complete IP follower configuration. Virtual routing interfaces 2 and 3 have been configured to share the IP address of virtual routing interface 1, but also have been configured to use their own ACLs instead of virtual routing interface 1's ACLs.

```
BigIron(config)# vlan 1 name primary_vlan
BigIron(config-vlan-1)# untag ethernet 1/1
BigIron(config-vlan-1)# tag ethernet 1/8
BigIron(config-vlan-1)# router-interface ve 1
BigIron(config-vlan-1)# exit
BigIron(config)# interface ve 1
BigIron(config-ve-1)# ip address 10.0.0.1/24
BigIron(config-ve-1)# ip access-group 1 in
BigIron(config-ve-1)# exit

BigIron(config)# vlan 2 name followerA
BigIron(config-vlan-2)# untag ethernet 1/2
BigIron(config-vlan-2)# tag ethernet 1/8
BigIron(config-vlan-2)# router-interface ve 2
BigIron(config-vlan-2)# exit
BigIron(config)# interface ve 2
BigIron(config-ve-2)# ip follow ve 1
BigIron(config-ve-2)# no ip follow acl
BigIron(config-ve-2)# ip access-group 2 in
BigIron(config-ve-2)# exit

BigIron(config)# vlan 3 name followerB
BigIron(config-vlan-3)# untag ethernet 1/5 to 1/6
BigIron(config-vlan-3)# tag ethernet 1/8
BigIron(config-vlan-3)# router-interface ve 3
BigIron(config-vlan-3)# exit
BigIron(config)# interface ve 3
BigIron(config-ve-3)# ip follow ve 1
BigIron(config-ve-3)# no ip follow acl
BigIron(config-ve-3)# ip access-group 3 out
BigIron(config-ve-3)# exit
```

Configuring VLAN Groups and Virtual Routing Interface Groups

To simplify configuration when you have many VLANs with the same configuration, you can configure VLAN groups and virtual routing interface groups.

NOTE: VLAN groups are supported on the NetIron Internet Backbone router and BigIron Layer 3 Switches and Layer 2 Switches with Management 2 or higher modules. Virtual routing interface groups are supported only on the chassis-based Layer 3 Switches.

When you create a VLAN group, the VLAN parameters you configure for the group apply to all the VLANs within the group. Additionally, you can easily associate the same IP subnet interface with all the VLANs in a group by configuring a virtual routing interface group with the same ID as the VLAN group.

- The VLAN group feature allows you to create multiple port-based VLANs with identical port members. Since the member ports are shared by all the VLANs within the group, you must add the ports as tagged ports. This feature not only simplifies VLAN configuration but also allows you to have a large number of identically configured VLANs in a startup-config file on the device's flash memory module. Normally, a startup-config file with a large number of VLANs might not fit on the flash memory module. By grouping the identically configured VLANs, you can conserve space in the startup-config file so that it fits on the flash memory module.
- The virtual routing interface group feature is useful when you want to configure the same IP subnet address on all the port-based VLANs within a VLAN group. You can configure a virtual routing interface group only after you configure a VLAN group with the same ID. The virtual routing interface group automatically applies to the VLANs in the VLAN group that has the same ID and cannot be applied to other VLAN groups or to individual VLANs.

You can create up to 32 VLAN groups and 32 virtual routing interface groups. A virtual routing interface group always applies only to the VLANs in the VLAN group with the same ID.

NOTE: Depending on the size of the VLAN ID range you want to use for the VLAN group, you might need to allocate additional memory for VLANs. On Layer 3 Switches, if you allocate additional memory for VLANs, you also need to allocate the same amount of memory for virtual routing interfaces. This is true regardless of whether you use the virtual routing interface groups. To allocate additional memory, see "Allocating Memory for More VLANs or Virtual Routing Interfaces" on page 22-48.

Configuring a VLAN Group

To configure a VLAN group, use the following CLI method.

USING THE CLI

To configure a VLAN group, enter commands such as the following:

```
BigIron(config)# vlan-group 1 vlan 2 to 1000
BigIron(config-vlan-group-1)# tagged 1/1 to 1/2
```

The first command in this example begins configuration for VLAN group 1, and assigns VLANs 2 through 1000 to the group. The second command adds ports 1/1 and 1/2 as tagged ports. Since all the VLANs in the group share the ports, you must add the ports as tagged ports.

Syntax: `vlan-group <num> vlan <vlan-id> to <vlan-id>`

Syntax: `tagged ethernet | pos <portnum> [to <portnum> | ethernet <portnum>]`

The `<num>` parameter with the **vlan-group** command specifies the VLAN group ID and can be from 1 – 32. The **vlan <vlan-id> to <vlan-id>** parameters specify a contiguous range (a range with no gaps) of individual VLAN IDs. Specify the low VLAN ID first and the high VLAN ID second. The command adds all the specified VLANs to the VLAN group.

NOTE: The device's memory must be configured to contain at least the number of VLANs you specify for the higher end of the range. For example, if you specify 2048 as the VLAN ID at the high end of the range, you first must increase the memory allocation for VLANs to 2048 or higher. Additionally, on Layer 3 Switches, if you allocate additional memory for VLANs, you also need to allocate the same amount of memory for virtual routing interfaces, before you configure the VLAN groups. This is true regardless of whether you use the virtual routing interface groups. The memory allocation is required because the VLAN groups and virtual routing interface groups have a one-to-one mapping. See "Allocating Memory for More VLANs or Virtual Routing Interfaces" on page 22-48.

If a VLAN within the range you specify is already configured, the CLI does not add the group but instead displays an error message. In this case, create the group by specifying a valid contiguous range. Then add more VLANs to the group after the CLI changes to the configuration level for the group. See the following example.

You can add and remove individual VLANs or VLAN ranges from at the VLAN group configuration level. For example, if you want to add VLANs 1001 and 1002 to VLAN group 1 and remove VLANs 900 through 1000, enter the following commands:

```
BigIron(config-vlan-group-1)# add-vlan 1001 to 1002
BigIron(config-vlan-group-1)# remove-vlan 900 to 1000
```

Syntax: add-vlan <vlan-id> [to <vlan-id>]

Syntax: remove-vlan <vlan-id> [to <vlan-id>]

USING THE WEB MANAGEMENT INTERFACE

You cannot configure this feature using the Web management interface.

Displaying Information about VLAN Groups

To display VLAN group configuration information, enter the following command:

```
BigIron# show vlan-group
vlan-group 1 vlan 2 to 20
  tagged ethe 1/1 to 1/2
!
vlan-group 2 vlan 21 to 40
  tagged ethe 1/1 to 1/2
!
```

Syntax: show vlan-group [<group-id>]

This example shows configuration information for two VLAN groups, group 1 and group 2.

The <group-id> specifies a VLAN group. If you do not use this parameter, the configuration information for all the configured VLAN groups is displayed.

Configuring a Virtual Routing Interface Group

A virtual routing interface group allows you to associate the same IP subnet interface with multiple port-based VLANs. For example, if you associate a virtual routing interface group with a VLAN group, all the VLANs in the group have the IP interface of the virtual routing interface group.

To configure a virtual routing interface group, use the following CLI method.

NOTE: When you configure a virtual routing interface group, all members of the group have the same IP subnet address. This feature is useful in collocation environments where the device has many IP addresses and you want to conserve the IP address space.

USING THE CLI

To configure a virtual routing interface group, enter commands such as the following:

```
BigIron(config)# vlan-group 1
BigIron(config-vlan-group-1)# group-router-interface
```

```
BigIron(config-vlan-group-1)# exit
BigIron(config)# interface group-ve 1
BigIron(config-vif-group-1)# ip address 10.10.10.1/24
```

These commands enable VLAN group 1 to have a group virtual routing interface, then configure virtual routing interface group 1. The software always associates a virtual routing interface group only with the VLAN group that has the same ID. In this example, the VLAN group ID is 1, so the corresponding virtual routing interface group also must have ID 1.

Syntax: group-router-interface

Syntax: interface group-ve <num>

Syntax: [no] ip address <ip-addr> <ip-mask> [secondary]

or

Syntax: [no] ip address <ip-addr>/<mask-bits> [secondary]

The **router-interface-group** command enables a VLAN group to use a virtual routing interface group. Enter this command at the configuration level for the VLAN group. This command configures the VLAN group to use the virtual routing interface group that has the same ID as the VLAN group. You can enter this command when you configure the VLAN group for the first time or later, after you have added tagged ports to the VLAN and so on.

The <num> parameter in the **interface group-ve <num>** command specifies the ID of the VLAN group with which you want to associate this virtual routing interface group. The VLAN group must already be configured and enabled to use a virtual routing interface group. The software automatically associates the virtual routing interface group with the VLAN group that has the same ID. You can associate a virtual routing interface group only with the VLAN group that has the same ID.

The syntax and usage for the **ip address** command is the same as when you use the command at the interface level to add an IP interface.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure this feature using the Web management interface.

Displaying the VLAN Group and Virtual Routing Interface Group Information

To verify configuration of VLAN groups and virtual routing interface groups, display the running-config file. If you have saved the configuration to the startup-config file, you also can verify the configuration by displaying the startup-config file. The following example shows the running-config information for the VLAN group and virtual routing interface group configured in the previous examples. The information appears in the same way in the startup-config file.

```
BigIron(config)# show running-config
```

lines not related to the VLAN group omitted...

```
vlan-group 1 vlan 2 to 900
  add-vlan 1001 to 1002
  tagged ethe 1/1 to 1/2
  router-interface-group
```

lines not related to the virtual routing interface group omitted...

```
interface group-ve 1
  ip address 10.10.10.1 255.255.255.0
```

NOTE: If you have enabled display of subnet masks in CIDR notation, the IP address information is shown as follows: 10.10.10.1/24.

Allocating Memory for More VLANs or Virtual Routing Interfaces

NetIron Internet Backbone routers and BigIron Layer 3 Switches can support up to 4095 VLANs and 4095 virtual routing interfaces.

The number of VLANs and virtual routing interfaces supported on your product depends on the device and, for Chassis devices, the amount of DRAM on the management module. Table 22.1 lists the default and configurable maximum numbers of VLANs and virtual routing interfaces for Layer 3 Switches and Layer 2 Switches. Unless otherwise noted, the values apply to both types of switches.

Table 22.1: VLAN and Virtual Routing Interface Support

Product	VLANs		Virtual Routing Interfaces	
	Default Maximum	Configurable Maximum	Default Maximum	Configurable Maximum
NetIron Internet Backbone router or BigIron Layer 3 Switch M2, M3, or M4 management module with at least 256MB	32	4095	255	4095
BigIron Layer 3 Switch M2 or M3 management module with at least 128MB management module	16	512 (Layer 3 Switch code) 2048 (Layer 2 Switch code)	255	512
BigIron Layer 3 Switch M1 management module	8	512	255	512
FastIron II FastIron II Plus FastIron III	8	512	255	512
FastIron 4802	64	512 (Layer 3 Switch code) 2048 (Layer 2 Switch code)	255	512
TurbolIron/8	8	512	255	512

NOTE: If many of your VLANs will have an identical configuration, you might want to configure VLAN groups and virtual routing interface groups after you increase the system capacity for VLANs and virtual routing interfaces. See “Configuring VLAN Groups and Virtual Routing Interface Groups” on page 22-45.

Increasing the Number of VLANs You Can Configure

To increase the size of the VLAN table, which determines how many VLANs you can configure, use either of the following methods.

NOTE: Although you can specify up to 4095 VLANs, you can configure only 4094 VLANs. VLAN ID 4094 is reserved for use by the Single Spanning Tree feature.

USING THE CLI

To increase the maximum number of VLANs you can configure, enter commands such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# system-max vlan 2048
BigIron(config)# write memory
BigIron(config)# end
BigIron# reload
```

Syntax: system-max vlan <num>

The <num> parameter indicates the maximum number of VLANs. The range of valid values depends on the device you are configuring. See Table 22.1.

USING THE WEB MANAGEMENT INTERFACE

To modify a table size using the Web management interface:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select the [Max-Parameter](#) link to display the Configure System Parameter Maximum Value table. This table lists the settings and valid ranges for all the configurable table sizes on the device.
3. Click the Modify button next to the row for the parameter (in this case, “vlan”).
4. Enter the new value for the table size. The value you enter specifies the maximum number of entries the table can hold.
5. Click Apply to save the changes to the device’s running-config.
6. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.
7. Click on the plus sign next to Command in the tree view to list the command options.
8. Select the [Reload](#) link and select Yes when the Web management interface asks you whether you really want to reload the software. Changes to cache and table sizes do not take effect until you reload the software.

Increasing the Number of Virtual Routing Interfaces You Can Configure

To increase the size of the virtual routing interface table, which determines how many virtual routing interfaces you can configure, use either of the following methods.

USING THE CLI

To increase the maximum number of virtual routing interfaces you can configure, enter commands such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# system-max virtual-interface 4095
BigIron(config)# write memory
BigIron(config)# end
BigIron# reload
```

Syntax: system-max virtual-interface <num>

The <num> parameter indicates the maximum number of virtual routing interfaces. The range of valid values depends on the device you are configuring. See Table 22.1.

USING THE WEB MANAGEMENT INTERFACE

See the Web management procedure for increasing the VLAN table size, in “Increasing the Number of VLANs You Can Configure” on page 22-48.

Configuring Super Aggregated VLANs

You can aggregate multiple VLANs within another VLAN. This feature allows you to construct Layer 2 paths and channels. This feature is particularly useful for Virtual Private Network (VPN) applications in which you need to provide a private, dedicated Ethernet connection for an individual client to transparently reach its subnet across multiple networks.

Conceptually, the paths and channels are similar to Asynchronous Transfer Mode (ATM) paths and channels. A path contains multiple channels, each of which is a dedicated circuit between two end points. The two devices at the end points of the channel appear to each other to be directly attached. The network that connects them is transparent to the two devices.


You can aggregate up to 4094 VLANs within another VLAN. This provides a total VLAN capacity on one Foundry device of 16,760,836 channels ($4094 * 4094$).

The devices connected through the channel are not visible to devices in other channels. Therefore, each client has a private link to the other side of the channel.

The feature allows point-to-point and point-to-multipoint connections.

Figure 22.17 shows a conceptual picture of the service that aggregated VLANs provide. Aggregated VLANs provide a path for multiple client channels. The channels do not receive traffic from other channels. Thus, each channel is a private link.

Figure 22.17 Conceptual Model of the Super Aggregated VLAN Application



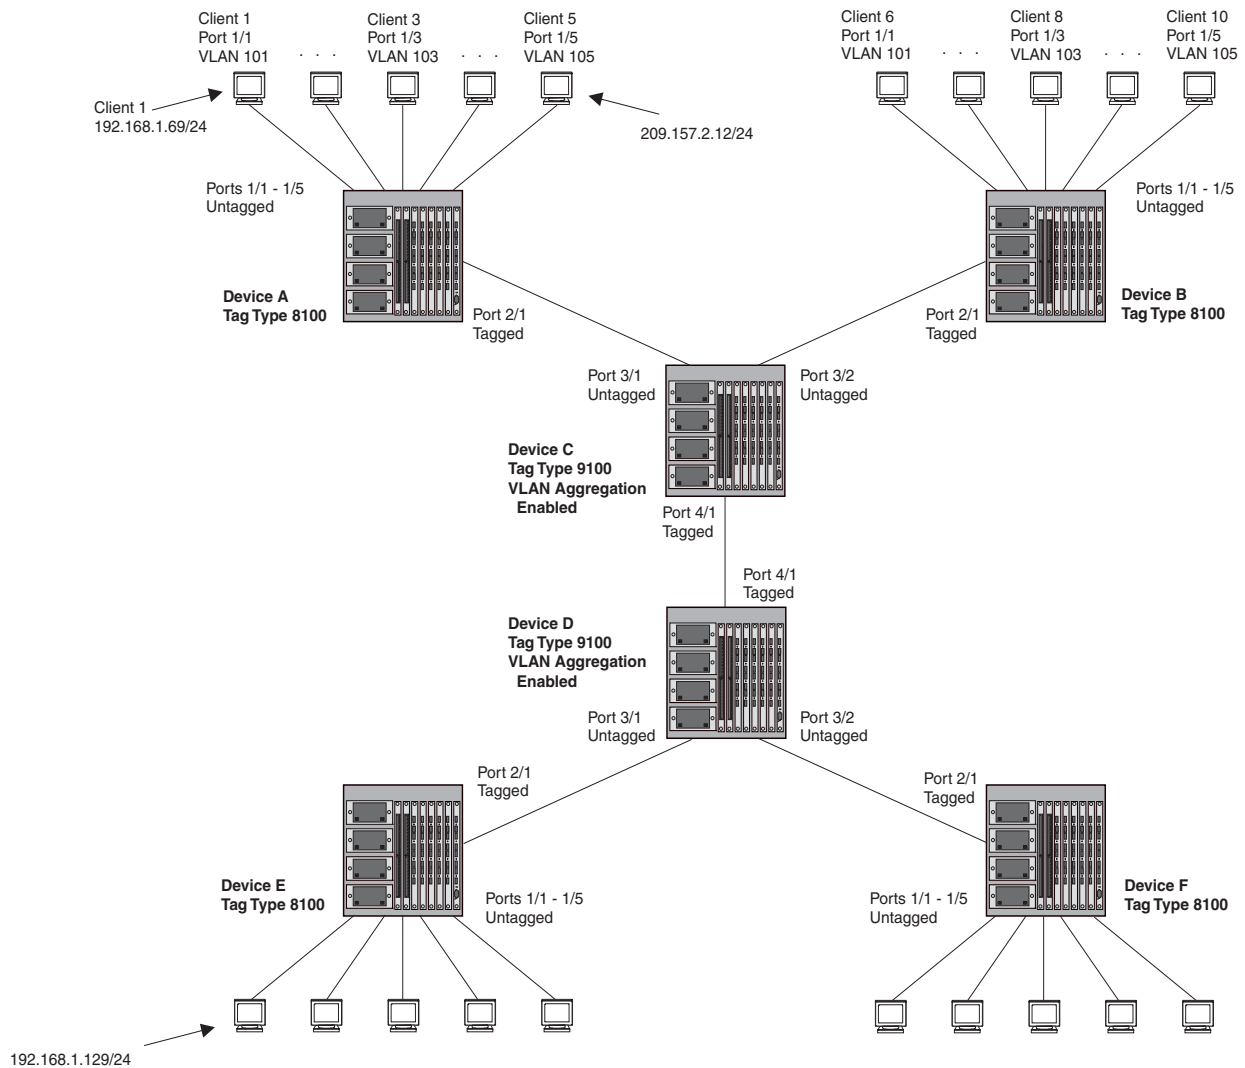
Each client connected to the edge device is in its own port-based VLAN, which is like an ATM channel. All the clients' VLANs are aggregated by the edge device into a single VLAN for connection to the core. The single VLAN that aggregates the clients' VLANs is like an ATM path.

The device that aggregates the VLANs forwards the aggregated VLAN traffic through the core. The core can consist of multiple devices that forward the aggregated VLAN traffic. The edge device at the other end of the core

separates the aggregated VLANs into the individual client VLANs before forwarding the traffic. The edge devices forward the individual client traffic to the clients. For the clients' perspective, the channel is a direct point-to-point link.

Figure 22.18 shows an example application that uses aggregated VLANs. This configuration includes the client connections shown in Figure 22.17.

Figure 22.18 Example Super Aggregated VLAN Application



In this example, a collocation service provides private channels for multiple clients. Although the same devices are used for all the clients, the VLANs ensure that each client receives its own Layer 2 broadcast domain, separate from the broadcast domains of other clients. For example, client 1 cannot ping client 5.

The clients at each end of a channel appear to each other to be directly connected and thus can be on the same subnet and use network services that require connection to the same subnet. In this example, client 1 is in subnet 192.168.1.0/24 and so is the device at the other end of client 1's channel.

Since each VLAN configured on the core devices is an aggregate of multiple client VLANs, the aggregated VLANs greatly increase the number of clients a core device can accommodate.

This example shows a single link between the core devices. However, you can use a trunk group to add link-level redundancy.

Configuration Note for FES Devices

This note applies to the FastIron Edge Switch running software release 03.1.00 or later.

NOTE: With Super Aggregated VLAN (SAV), 802.1Q-in-Q, and server trunk groups configured on the FES, server trunk groups do not function properly and the FES may generate the following message:

```
"Warning: Out of Server Trunk Flow Entries"
```

With multiple 802.1Q tags (802.1Q-in-Q), it is difficult for the trunk server to correctly identify where the Layer 3 header begins. Without proper information with which to forward the packets, server trunks fail. Note that switch trunk groups function properly because the FES examines the Destination and Source MAC addresses, which appear before the Layer 3 header.

When you enable SAV on a port, Foundry recommends that you do not configure any other feature on the port that requires examining the Layer 3 packet header and beyond. This includes features such as access lists, TCP SYN attack protection, and ICMP attack protection.

Configuration Note for BigIron MG8 and NetIron 40G

This note applies to the BigIron MG8 and NetIron 40G running software release 02.1.00 or later.

In earlier releases, the BigIron MG8 and NetIron 40G had a default maximum frame size of 1518 bytes. With release 02.1.00, the maximum frame size supported on a port is modified to dynamically change based upon the port's tagging characteristics as described:

- **Untagged Ports** – The maximum frame size supported on an untagged port is 1518 bytes. This includes 1500 bytes for payload, 14 bytes for the MAC header, and 4 bytes for the CRC. This limit is defined for untagged ports in the IEEE 802.1 specification.
- **Tagged Ports** – The maximum size supported on tagged ports is 1522 bytes. The additional 4 bytes over the untagged port maximum are allowed to support the additional bytes needed to include a VLAN tag.
- **Super-aggregated VLAN Support** – A maximum of 1526 bytes are supported on ports where super-aggregated VLANs are configured. This allows for an additional 8 bytes over the untagged port maximum to allow for support of two VLAN tags.

Configuring Aggregated VLANs

To configure aggregated VLANs, perform the following tasks:

- On each edge device, configure a separate port-based VLAN for each client connected to the edge device. In each client VLAN:
 - Add the port connected to the client as an untagged port.
 - Add the port connected to the core device (the device that will aggregate the VLANs) as a tagged port. This port must be tagged because all the client VLANs share the port as an uplink to the core device.
- On each core device:
 - Enable VLAN aggregation. This support allows the core device to add an additional tag to each Ethernet frame that contains a VLAN packet from the edge device. The additional tag identifies the aggregate VLAN (the path). However, the additional tag can cause the frame to be longer than the maximum supported frame size. The larger frame support allows Ethernet frames up to 1530 bytes long.

NOTE: Enable the VLAN aggregation option only on the core devices.

- Configure a VLAN tag type (tag ID) that is different than the tag type used on the edge devices. If you use the default tag type (8100) on the edge devices, set the tag type on the core devices to another value, such as 9100. The tag type must be the same on all the core devices. The edge devices also must have the same tag type but the type must be different from the tag type on the core devices.

NOTE: You can enable the Spanning Tree Protocol (STP) on the edge devices or the core devices, but not both. If you enable STP on the edge devices and the core devices, STP will prevent client traffic from travelling through the core to the other side.

Configuring Aggregated VLANs on an Edge Device

To configure aggregated VLANs on an edge device, use one of the following methods.

USING THE CLI

To configure the aggregated VLANs on device A in Figure 22.18 on page 22-51, enter the following commands:

```
BigIron(config)# vlan 101 by port
BigIron(config-vlan-101)# tagged ethernet 2/1
BigIron(config-vlan-101)# untagged ethernet 1/1
BigIron(config-vlan-101)# exit
BigIron(config)# vlan 102 by port
BigIron(config-vlan-102)# tagged ethernet 2/1
BigIron(config-vlan-102)# untagged ethernet 1/2
BigIron(config-vlan-102)# exit
BigIron(config)# vlan 103 by port
BigIron(config-vlan-103)# tagged ethernet 2/1
BigIron(config-vlan-103)# untagged ethernet 1/3
BigIron(config-vlan-103)# exit
BigIron(config)# vlan 104 by port
BigIron(config-vlan-104)# tagged ethernet 2/1
BigIron(config-vlan-104)# untagged ethernet 1/4
BigIron(config-vlan-104)# exit
BigIron(config)# vlan 105 by port
BigIron(config-vlan-105)# tagged ethernet 2/1
BigIron(config-vlan-105)# untagged ethernet 1/5
BigIron(config-vlan-105)# exit
BigIron(config)# write memory
```

Syntax: [no] vlan <vlan-id> [by port]

Syntax: [no] tagged ethernet <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] untagged ethernet <portnum> [to <portnum> | ethernet <portnum>]

Use the **tagged** command to add the port that the device uses for the uplink to the core device. Use the **untagged** command to add the ports connected to the individual clients.

USING THE WEB MANAGEMENT INTERFACE

You cannot enable VLAN aggregation using the Web management interface. The other options you need for configuring Aggregated VLANs are present in earlier software releases and are supported in the Web management interface.

Configuring Aggregated VLANs on a Core Device

To configure aggregated VLANs on a core device, use one of the following methods.

USING THE CLI

To configure the aggregated VLANs on device C in Figure 22.18 on page 22-51, enter the following commands:

```
BigIron(config)# tag-type 9100
BigIron(config)# aggregated-vlan
BigIron(config)# vlan 101 by port
BigIron(config-vlan-101)# tagged ethernet 4/1
BigIron(config-vlan-101)# untagged ethernet 3/1
BigIron(config-vlan-101)# exit
BigIron(config)# vlan 102 by port
BigIron(config-vlan-102)# tagged ethernet 4/1
```

```
BigIron(config-vlan-102)# untagged ethernet 3/2
BigIron(config-vlan-102)# exit
BigIron(config)# write memory
```

Syntax: [no] tag-type <num>

Syntax: [no] aggregated-vlan

The <num> parameter specifies the tag type can be a hexadecimal value from 0 – ffff. The default is 8100.

[USING THE WEB MANAGEMENT INTERFACE](#)

You cannot enable VLAN aggregation using the Web management interface.

Applying QoS Markings to SAV Traffic in the Network Core

NOTE: This feature is available in Service Provider software releases 09.1.00 and later.

Foundry's initial implementation of the SAV feature does not prioritize traffic traversing through the core of the network. This is because the core device must add another tag type on top of the 802.1q tag type set by the edge (client) device. When the core device adds the additional tag type, it does not copy the customer QoS priority into the tag value, and it sets the QoS priority to 0 (zero). Hence, packets traveling through the core are not prioritized.

In software release 09.1.00 for the Service Provider, you can optionally configure an untagged interface to copy the QoS bits from the tag value set by the edge device to the tag value set by the core device. The Foundry device copies the QoS bits when it adds the additional tag type as the packets enter the core network. This way, the Foundry device can prioritize SAV traffic transmitted through the network core.

Configuration Notes and Rules

- The device must be running software release 09.1.00 or later.
- The port on which you enable this feature must be an untagged port.
- The tag type for the entire core device must be 0x9100.

Configuring the Core Device to Copy the QoS Priority

NOTE: This feature is also available in Terathon IronWare releases 02.2.00 and later for the BigIron MG8 and NetIron 40G.

To enable this feature, enter commands such as the following on an untagged port of the core device:

```
NetIron (config)# int e 4/1
NetIron (config-if-e1000-4/1)# aggregated-vlan-copy-cos
```

This configuration causes the Foundry device to copy the QoS marking from packets entering interface e 4/1 to the tag value set by the core device when it applies the additional 802.1q tag type. Consequently, the Foundry device will prioritize high-priority customer traffic as it traverses through the core of the network.

Syntax: [no] aggregated-vlan-copy-cos

Use the **no** form of the command to disable this feature.

NOTE: This command is available in Service Provider software releases 09.1.00 and later and on BigIron MG8 and NetIron 40G software release 02.2.00 and later.

Complete CLI Examples

The following sections show all the Aggregated VLAN configuration commands on the devices in Figure 22.18 on page 22-51.

NOTE: In these examples, the configurations of the edge devices (A, B, E, and F) are identical. The configurations of the core devices (C and D) also are identical. The aggregated VLAN configurations of the edge and core devices on one side must be symmetrical (in fact, a mirror image) to the configurations of the devices on the other side. For simplicity, the example in Figure 22.18 on page 22-51 is symmetrical in terms of the port numbers. This allows the configurations for both sides of the link to be the same. If your configuration does not use symmetrically arranged port numbers, the configurations should not be identical but must use the correct port numbers.

Commands for Device A

```
BigIronA(config)# vlan 101 by port
BigIronA(config-vlan-101)# tagged ethernet 2/1
BigIronA(config-vlan-101)# untagged ethernet 1/1
BigIronA(config-vlan-101)# exit
BigIronA(config)# vlan 102 by port
BigIronA(config-vlan-102)# tagged ethernet 2/1
BigIronA(config-vlan-102)# untagged ethernet 1/2
BigIronA(config-vlan-102)# exit
BigIronA(config)# vlan 103 by port
BigIronA(config-vlan-103)# tagged ethernet 2/1
BigIronA(config-vlan-103)# untagged ethernet 1/3
BigIronA(config-vlan-103)# exit
BigIronA(config)# vlan 104 by port
BigIronA(config-vlan-104)# tagged ethernet 2/1
BigIronA(config-vlan-104)# untagged ethernet 1/4
BigIronA(config-vlan-104)# exit
BigIronA(config)# vlan 105 by port
BigIronA(config-vlan-105)# tagged ethernet 2/1
BigIronA(config-vlan-105)# untagged ethernet 1/5
BigIronA(config-vlan-105)# exit
BigIronA(config)# write memory
```

Commands for Device B

The commands for configuring device B are identical to the commands for configuring device A. Notice that you can use the same channel VLAN numbers on each device. The devices that aggregate the VLANs into a path can distinguish between the identically named channel VLANs based on the ID of the path VLAN.

```
BigIronB(config)# vlan 101 by port
BigIronB(config-vlan-101)# tagged ethernet 2/1
BigIronB(config-vlan-101)# untagged ethernet 1/1
BigIronB(config-vlan-101)# exit
BigIronB(config)# vlan 102 by port
BigIronB(config-vlan-102)# tagged ethernet 2/1
BigIronB(config-vlan-102)# untagged ethernet 1/2
BigIronB(config-vlan-102)# exit
BigIronB(config)# vlan 103 by port
BigIronB(config-vlan-103)# tagged ethernet 2/1
BigIronB(config-vlan-103)# untagged ethernet 1/3
BigIronB(config-vlan-103)# exit
BigIronB(config)# vlan 104 by port
BigIronB(config-vlan-104)# tagged ethernet 2/1
BigIronB(config-vlan-104)# untagged ethernet 1/4
BigIronB(config-vlan-104)# exit
BigIronB(config)# vlan 105 by port
BigIronB(config-vlan-105)# tagged ethernet 2/1
BigIronB(config-vlan-105)# untagged ethernet 1/5
BigIronB(config-vlan-105)# exit
BigIronB(config)# write memory
```

Commands for Device C

Since device C is aggregating channel VLANs from devices A and B into a single path, you need to change the tag type and enable VLAN aggregation.

```
BigIronC(config)# tag-type 9100
BigIronC(config)# aggregated-vlan
BigIronC(config)# vlan 101 by port
BigIronC(config-vlan-101)# tagged ethernet 4/1
BigIronC(config-vlan-101)# untagged ethernet 3/1
BigIronC(config-vlan-101)# exit
BigIronC(config)# vlan 102 by port
BigIronC(config-vlan-102)# tagged ethernet 4/1
BigIronC(config-vlan-102)# untagged ethernet 3/2
BigIronC(config-vlan-102)# exit
BigIronC(config)# write memory
```

Commands for Device D

Device D is at the other end of path and separates the channels back into individual VLANs. The tag type must be the same as tag type configured on the other core device (Device C). In addition, VLAN aggregation also must be enabled.

```
BigIronD(config)# tag-type 9100
BigIronD(config)# aggregated-vlan
BigIronD(config)# vlan 101 by port
BigIronD(config-vlan-101)# tagged ethernet 4/1
BigIronD(config-vlan-101)# untagged ethernet 3/1
BigIronD(config-vlan-101)# exit
BigIronD(config)# vlan 102 by port
BigIronD(config-vlan-102)# tagged ethernet 4/1
BigIronD(config-vlan-102)# untagged ethernet 3/2
BigIronD(config-vlan-102)# exit
BigIronD(config)# write memory
```

Commands for Device E

Since the configuration in Figure 22.18 on page 22-51 is symmetrical, the commands for configuring device E are identical to the commands for configuring device A.

```
BigIronE(config)# vlan 101 by port
BigIronE(config-vlan-101)# tagged ethernet 2/1
BigIronE(config-vlan-101)# untagged ethernet 1/1
BigIronE(config-vlan-101)# exit
BigIronE(config)# vlan 102 by port
BigIronE(config-vlan-102)# tagged ethernet 2/1
BigIronE(config-vlan-102)# untagged ethernet 1/2
BigIronE(config-vlan-102)# exit
BigIronE(config)# vlan 103 by port
BigIronE(config-vlan-103)# tagged ethernet 2/1
BigIronE(config-vlan-103)# untagged ethernet 1/3
BigIronE(config-vlan-103)# exit
BigIronE(config)# vlan 104 by port
BigIronE(config-vlan-104)# tagged ethernet 2/1
BigIronE(config-vlan-104)# untagged ethernet 1/4
BigIronE(config-vlan-104)# exit
BigIronE(config)# vlan 105 by port
BigIronE(config-vlan-105)# tagged ethernet 2/1
BigIronE(config-vlan-105)# untagged ethernet 1/5
BigIronE(config-vlan-105)# exit
BigIronE(config)# write memory
```

Commands for Device F

The commands for configuring device F are identical to the commands for configuring device E. In this example, since the port numbers on each side of the configuration in Figure 22.18 on page 22-51 are symmetrical, the configuration of device F is also identical to the configuration of device A and device B.

```
BigIronF(config)# vlan 101 by port
BigIronF(config-vlan-101)# tagged ethernet 2/1
BigIronF(config-vlan-101)# untagged ethernet 1/1
BigIronF(config-vlan-101)# exit
BigIronF(config)# vlan 102 by port
BigIronF(config-vlan-102)# tagged ethernet 2/1
BigIronF(config-vlan-102)# untagged ethernet 1/2
BigIronF(config-vlan-102)# exit
BigIronF(config)# vlan 103 by port
BigIronF(config-vlan-103)# tagged ethernet 2/1
BigIronF(config-vlan-103)# untagged ethernet 1/3
BigIronF(config-vlan-103)# exit
BigIronF(config)# vlan 104 by port
BigIronF(config-vlan-104)# tagged ethernet 2/1
BigIronF(config-vlan-104)# untagged ethernet 1/4
BigIronF(config-vlan-104)# exit
BigIronF(config)# vlan 105 by port
BigIronF(config-vlan-105)# tagged ethernet 2/1
BigIronF(config-vlan-105)# untagged ethernet 1/5
BigIronF(config-vlan-105)# exit
BigIronF(config)# write memory
```

Configuring 802.1q-in-q Tagging


NOTE: This feature is supported FastIron Edge Switches running software release 03.1.00 or later.

Software release 03.1.00 and later provide finer granularity for configuring 802.1q tagging, enabling you to configure 802.1Q tag-types on a group of ports, thereby enabling the creation of two identical 802.1Q tags (802.1Q-in-Q tagging) on a single device. This enhancement improves SAV interoperability between Foundry devices and other vendors' devices that support the 802.1Q tag-types, but are not very flexible with the tag-types they accept.

- In releases prior to 03.1.00, you can configure a single 802.1Q tag type on all ports of an FES device. The default 802.1Q tag on a Foundry device is 8100 (hexadecimal), compliant with the 802.1Q specification.

Figure 22.19 shows an 802.1Q configuration example.

Figure 22.19 802.1Q Configuration Example



As shown in Figure 22.19, the ports to customer interfaces are untagged, whereas the uplink ports to the


provider cloud are tagged, because multiple client VLANs share the uplink to the provider cloud. In this example, the Foundry device treats the customer's private VLAN ID and 8100 tag type as normal payload, and adds the 9100 tag type to the packet when the packet is sent to the uplink and forwarded along the provider cloud.

As long as the switches in the provider's network are Foundry devices or devices that can use the 9100 tag type, the data gets switched along the network. However, devices along the provider's cloud that do not support the 9100 tag type may not properly handle the packets.

- Release 03.1.00 and the introduction of 802.1Q-in-Q tagging, provide finer granularity for configuring 802.1Q tagging, enabling you to configure 802.1Q tag-types on a group of ports, thereby enabling the creation of two identical 802.1Q tags (802.1Q-in-Q tagging) on a single device. This enhancement improves SAV interoperability between Foundry devices and other vendors' devices that support the 802.1Q tag-types, but are not very flexible with the tag-types they accept.

Figure 22.22 shows an example application of the 802.1Q-in-Q enhancement.

Figure 22.20 802.1Q-in-Q Configuration Example



In Figure 22.22, the untagged ports (to customer interfaces) accept frames that have any 802.1Q tag other than the configured tag-type 9100. These packets are considered untagged on this incoming port and are re-tagged when they are sent out of the uplink towards the provider. The 802.1Q tag-type on the uplink port is 8100, so the Foundry device will switch the frames to the uplink device with an additional 8100 tag, thereby supporting devices that only support this method of VLAN tagging.

Configuration Rules

- On the FastIron Edge Switches, you can configure 802.1Q tag-types per port region.
- Since the uplink (to the provider cloud) and the edge link (to the customer port) must have different 802.1Q tags, make sure the uplink and edge link are in different port regions.
- If you configure a port with an 802.1Q tag-type, the Foundry device automatically applies the 802.1Q tag-type to all ports within the same port region.
- If you remove the 802.1Q tag-type from a port, the Foundry device automatically removes the 802.1Q tag-type from all ports within the same port region.
- The FastIron Edge Switches support one configured tag-type per device along with the default tag-type of 8100. For example, if you configure an 802.1Q tag of 9100 on ports 1 – 8, then later configure an 802.1Q tag of 5100 on port 9, the device automatically applies the 5100 tag to all ports in the same port region as port 9, and also changes the 802.1Q tag-type on ports 1 – 8 to 5100.

Enabling 802.1Q-in-Q Tagging

To enable the 802.1Q-in-Q feature, configure an 802.1Q tag on the untagged edge links (the customer ports) to any value other than the 802.1Q tag for incoming traffic. For example, in Figure 22.21, the 802.1Q tag on the untagged edge links (ports 11 and 12) is 9100, whereas, the 802.1Q tag for incoming traffic is 8100.

To configure 802.1 Q-in-Q tagging as shown in Figure 22.21, enter commands such as the following on the untagged edge links of devices C and D:

```
FES4802(config)# tag-type 9100 e 11 to 12
FES4802(config)# aggregated-vlan
```

Note that since ports 11 and 12 belong to the port region 9 – 16, the 802.1Q tag actually applies to ports 9 – 16.

Syntax: [no] tag-type <num> [e <port number> [to <port number>]]

The <num> parameter specifies the tag-type number and can be a hexadecimal value from 0 - ffff. The default is 8100.

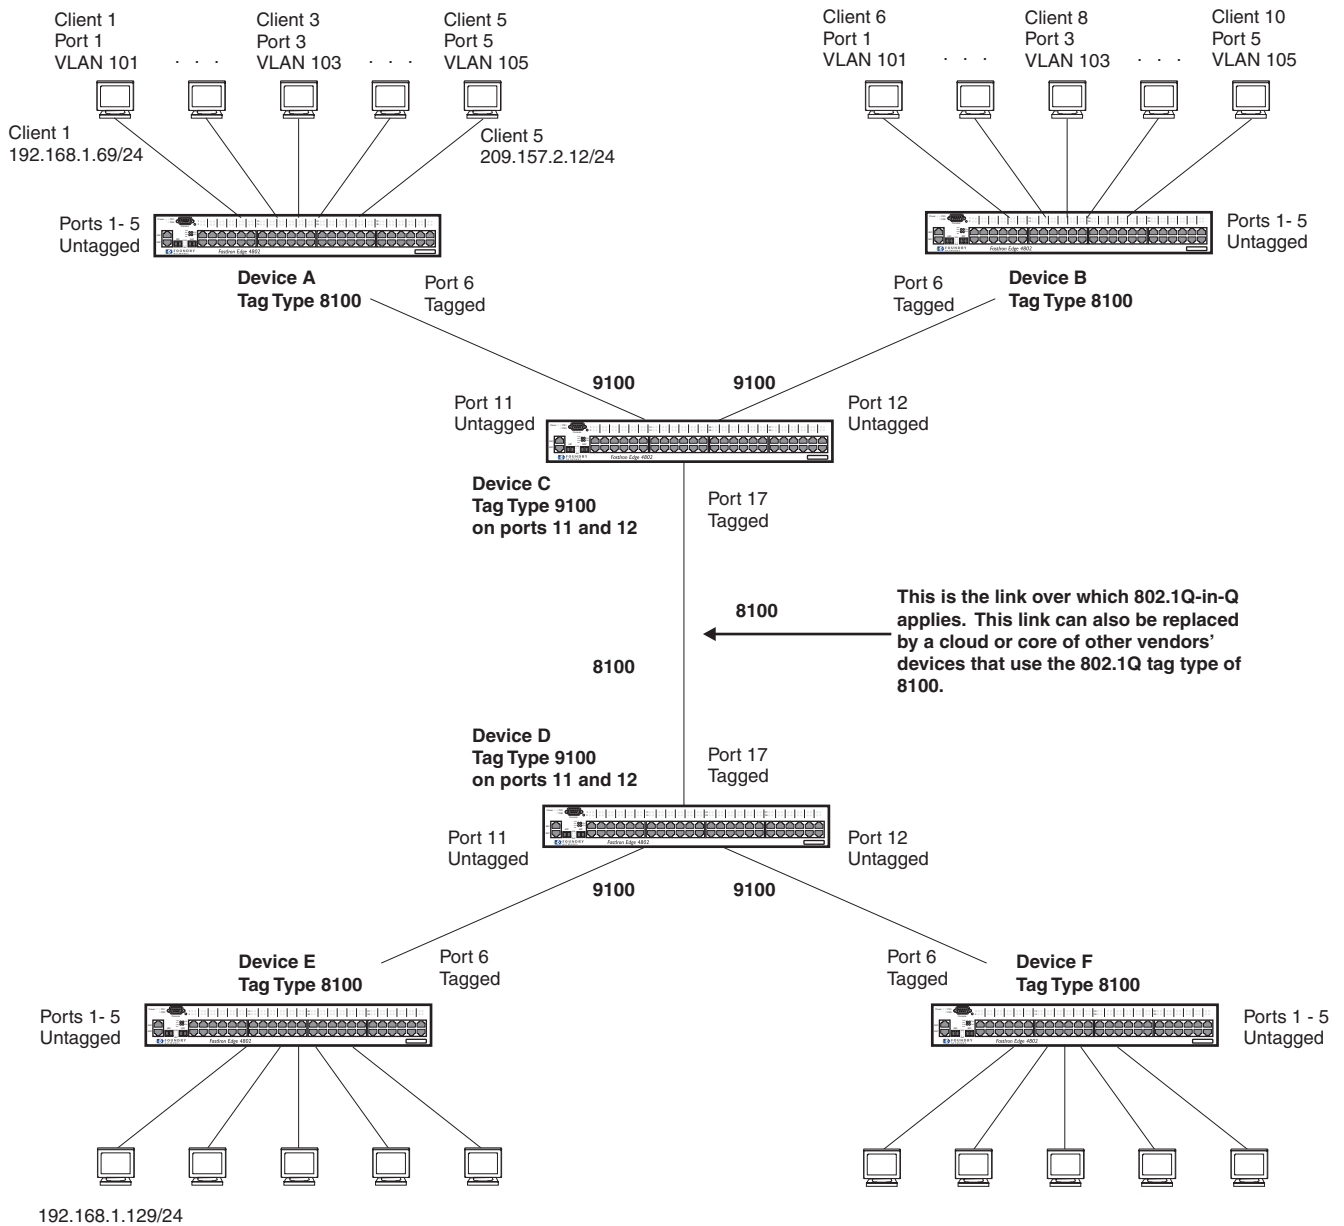
The **e <port number> to <port number>** parameter specifies the port(s) that will use the defined 802.1Q tag. This parameter operates with the following rules:

- If you specify a single port number, the 802.1Q tag applies to all ports within the port region. For example, if you enter the command **tag-type 9100 e 1**, the Foundry device automatically applies the 802.1Q tag to ports 1 – 8 since all of these ports are in the same port region. You can use the **show running-config** command to view how the command has been applied.
- If you do not specify a port or range of ports, the 802.1Q tag applies to all Ethernet ports on the device.

Example Configuration

Figure 22.21 shows an example 802.1Q-in-Q configuration.

Figure 22.21 Example 802.1Q-in-Q Configuration



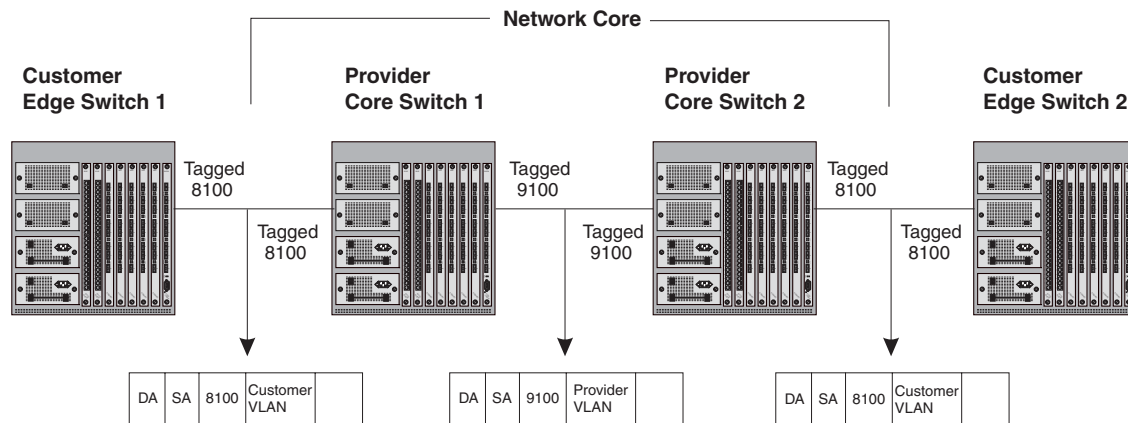
Configuring 802.1q Tag-type Translation

NOTE: This feature is supported on the BigIron MG8 and on Foundry devices running Service Provider software release 09.1.00 or later.

Service Provider release 09.1.00 and the introduction of 802.1Q tag-type translation provides finer granularity for configuring multiple 802.1q tag-types on a single device, by enabling you to configure 802.1q tag-types per port group. This enhancement allows for tag-type translation from one port group to the next on tagged interfaces.

Figure 22.22 shows a basic example application of the 802.1q tag-type translation feature.

Figure 22.22 802.1q Tag-type Translation Configuration Example 1




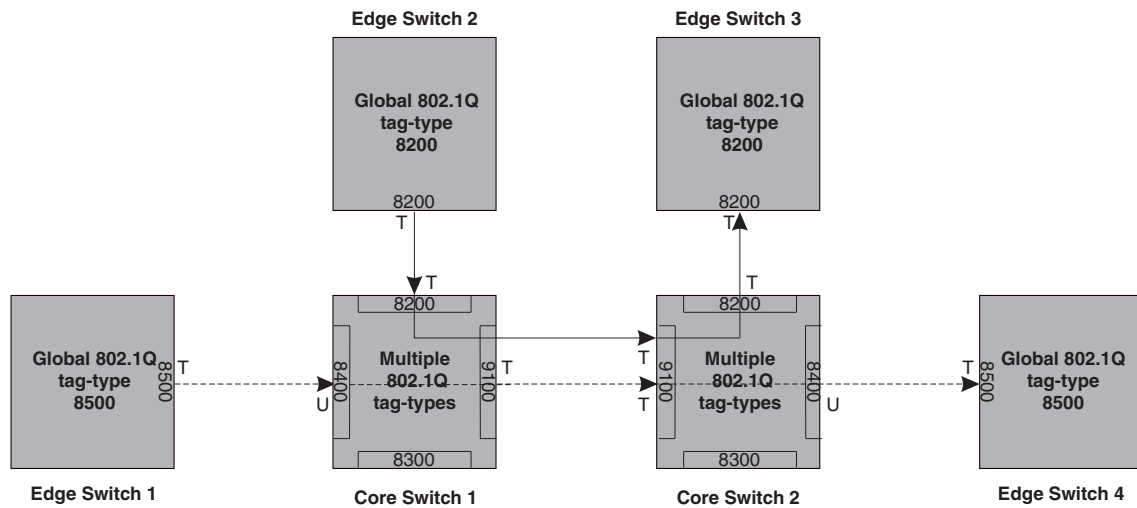
As illustrated in Figure 22.22, the devices process the packet as follows:

- Customer Edge Switch 1 sends a packet with an 802.1q tag-type of 8100 to Provider Core Switch 1.
- Since the customer-facing interface on Provider Core Switch 1 has the same 802.1q tag-type as the incoming packet, it removes the 8100 tag-type and replaces (translates) it with the 9100 tag-type as it sends the packet to the uplink (Provider Core Switch 2).
- The same process occurs between Provider Core Switch 2 and Customer Edge Switch 2.

Figure 22.22 shows a simple application of the 802.1q tag-type translation in which all of the ports are tagged and the tag-types between devices match. In this example, each device performs the 802.1q tag-type translation as the packet traverses the network.

Figure 22.23 shows a more complex example application in which some ports are untagged, not all tag-types between devices match, and the core devices have multiple tag-types. In this example, the tag-type translation feature integrates packets that have single and double tag-types.

Figure 22.23 802.1q Tag-type Translation Configuration Example 2



Legend:

- T - Tagged port
- U - Untagged port
- XXXX - DMA or port group
- > Path A
- > Path B

As illustrated in Figure 22.23, the devices process the packets as follows:

- Path A: When Core Switch 1 receives the tagged packet from Edge Switch 1, it *keeps* the 8500 tag-type in the frame header (because the incoming port on Core Switch 1 is untagged) and *adds* the 9100 tag-type as it sends the packet to the uplink (Core Switch 2). In this case, the packet is double-tagged as it travels between the core devices.
- Path B: When Core Switch 1 receives the tagged packet from Edge Switch 2, it *removes* the 8200 tag-type and replaces (translates) it with the 9100 tag-type as it sends the packet to the uplink (Core Switch 2).

For more information, see “Configuring 802.1q Tag-type Translation” on page 22-60.

Configuration Rules

- On the supported devices, you configure 802.1q tag-types per port region. Use the **show running-config** command at any level of the CLI to view port regions. Note that on Gigabit Ethernet modules, ports 1 and 2 belong to the same port region.
- Since the uplink (to the provider cloud) and the edge link (to the customer port) must have different 802.1q tag-types, make sure the uplink and edge link are in different port regions.

- If you configure a port with an 802.1q tag-type, the Foundry device automatically applies the 802.1q tag-type to all ports within the same port region.
- If you remove the 802.1q tag-type from a port, the Foundry device automatically removes the 802.1q tag-type from all ports within the same port region.
- Foundry does not recommend configuring different 802.1q tag-types on ports that are part of a multi-slot trunk. Use the same 802.1q tag-type for all ports in a multi-slot trunk.
- On the BigIron MG8 software release 02.2.02 and later and on NetIron 40G software release 02.2.04 and later, multiple 802.1Q tag types can be assigned to an interface module. Depending on the module, an 802.1Q tag can be assigned to an individual port or to a group of ports. Table 22.2 describes the granularity at which each of the BigIron MG8 and NetIron 40G interface modules can have 802.1Q tag-types assigned.

Table 22.2: 802.1Q tag-type assignments by module

module type	802.1Q tag-type assignment
4 x 10G	per port
2 x 10G	per port
1 x 10G	per port
40 x 1G	per 10 ports: 1 - 10, 11 - 20, 21 - 30, 31 - 40
20 x 1G	per 10 ports: 1 - 10, 11 - 20
10 x 1G	per 10 ports: 1 - 10
60 x 1G	per 20 ports: 1 - 20, 21 - 40, 41 - 60

Enabling 802.1q Tag-type Translation

To enable 802.1q tag-type translation, configure an 802.1q tag-type on the provider core link, between the provider core switches (see Figure 22.22). Enter commands such as the following:

```
NetIron(config)# tag-type 9100 e 11 to 12
NetIron(config)# aggregated-vlan
```

Note that since ports 11 and 12 belong to the port region 9 – 16, the 802.1q tag-type actually applies to ports 9 – 16.

NOTE: Do not configure 802.1q tag-type translation on the edge link (to the customer edge switch).

Syntax: [no] tag-type <num> [e <slot number>/<port number> [to <port number>]]

The <num> parameter specifies the tag-type number and can be a hexadecimal value from 0 - ffff. The default is 8100. Note that you must specify a value other than 8100.

The <slot number> parameter specifies the slot number on a Foundry Chassis device. This parameter is not applicable to stackable devices.

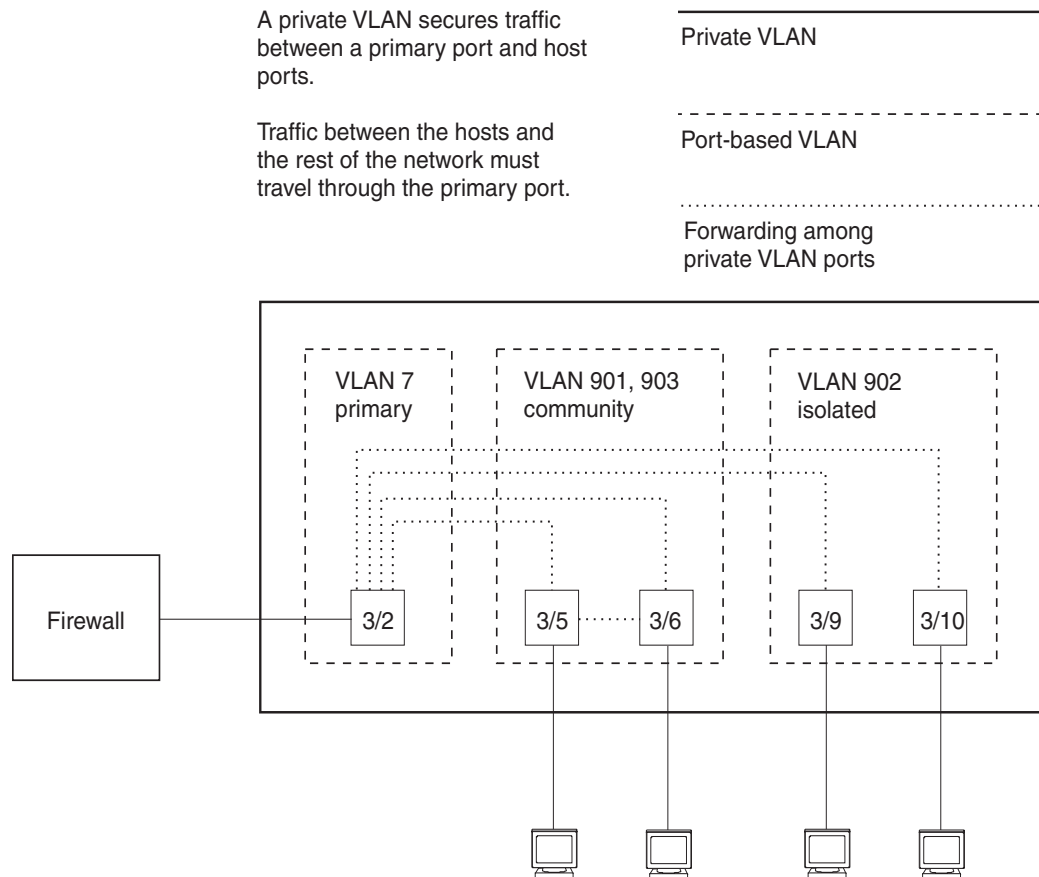
The <port number> [**to** <port number>] parameter specifies the port(s) that will use the defined 802.1q tag-type. This parameter operates with the following rules:

- If you specify a single port number, the 802.1q tag-type applies to all ports within the port region. For example, if you enter the command **tag-type 9100 e 1**, the Foundry device automatically applies the 802.1q tag to ports 1 – 8 since all of these ports are in the same port region (controlled by the same DMA). Use the **show running-config** command at any level of the CLI to view port regions. Note that on Gigabit Ethernet modules, ports 1 and 2 belong to the same port region.
- If the port that you specify is part of a multi-slot trunk, the device automatically applies the 802.1q tag-type to all of the ports that are part of the multi-slot trunk.
- If you do not specify a port or range of ports, the 802.1q tag-type applies to all Ethernet ports on the device.

Configuring Private VLANs

A private VLAN is a VLAN that has the properties of standard Layer 2 port-based VLANs but also provides additional control over flooding packets on a VLAN. Figure 22.24 shows an example of an application using a private VLAN.

Figure 22.24 Private VLAN used to secure communication between a workstation and servers



This example uses a private VLAN to secure traffic between hosts and the rest of the network through a firewall. Five ports in this example are members of a private VLAN. The first port (port 3/2) is attached to a firewall. The next four ports (ports 3/5, 3/6, 3/9, and 3/10) are attached to hosts that rely on the firewall to secure traffic between

the hosts and the rest of the network. In this example, two of the hosts (on ports 3/5 and 3/6) are in a community private VLAN, and thus can communicate with one another as well as through the firewall. The other two hosts (on ports 3/9 and 3/10), are in an isolated VLAN and thus can communicate only through the firewall. The two hosts are secured from communicating with one another even though they are in the same VLAN.

By default, the private VLAN does not forward broadcast or unknown-unicast packets from outside sources into the private VLAN. If needed, you can override this behavior for broadcast packets, unknown-unicast packets, or both. (See “Enabling Broadcast or Unknown Unicast Traffic to the Private VLAN” on page 22-68.)

You can configure a combination of the following types of private VLANs:

- **Primary** – The primary private VLAN ports are “promiscuous”. They can communicate with all the isolated private VLAN ports and community private VLAN ports in the isolated and community VLANs that are mapped to the promiscuous port.
- **Isolated** – Broadcasts and unknown unicasts received on isolated ports are sent only to the primary port. They are not flooded to other ports in the isolated VLAN.
- **Community** – Broadcasts and unknown unicasts received on community ports are sent to the primary port and also are flooded to the other ports in the community VLAN.

Each private VLAN must have a primary VLAN. The primary VLAN is the interface between the secured ports and the rest of the network. The private VLAN can have any combination of community and isolated VLANs. (See “Configuration Rules” on page 22-66.)

Table 22.3 list the differences between private VLANs and standard VLANs.

Table 22.3: Comparison of Private VLANs and Standard Port-Based VLANs

Forwarding Behavior	Private VLANs	Standard VLANs
All ports within a VLAN constitute a common Layer broadcast domain	No	Yes
Broadcasts and unknown unicasts are forwarded to all the VLAN's ports by default	No (isolated VLAN) Yes (community VLAN)	Yes
Known unicasts	Yes	Yes

Implementation Notes

- The private VLAN implementation in the current release uses the CPU for forwarding packets on the primary VLAN's “promiscuous” port. Other forwarding is performed in the hardware. Support for the hardware forwarding in this feature sometimes results in multiple MAC address entries for the same MAC address in the device's MAC address table. In this case, each of the entries is associated with a different VLAN. The multiple entries are a normal aspect of the implementation of this feature and do not indicate a software problem.
- By default, the primary VLAN does not forward broadcast or unknown unicast packets into the private VLAN. You also can use MAC address filters to control traffic forwarded into and out of the private VLAN. If you are implementing the private VLAN on a Layer 2 Switch, you also can use ACLs to control the traffic into and out of the private VLAN.

Configuration Notes for FES Devices

NOTE: This section applies to the FastIron Edge Switch (FES) running software release 03.1.00 or later.

- When Private VLAN mappings are enabled, the FastIron Edge Switch forwards unknown unicast, unknown multicast, and broadcast packets in software. By default, the FastIron Edge Switch forwards unknown unicast,

unknown multicast, and broadcast packets in hardware.

- Release 03.1.00 supports private VLANs on untagged ports only. You cannot configure isolated, community, or primary VLANs on 802.1Q tagged ports.
- The FastIron Edge Switch forwards all known unicast and multicast traffic in hardware. This differs from the way the BigIron implements private VLANs, in that the BigIron uses the CPU to forward packets on the primary VLAN's "promiscuous" port. In addition, on the BigIron, support for the hardware forwarding in this feature sometimes results in multiple MAC address entries for the same MAC address in the device's MAC address table. On the FastIron Edge Switch, multiple MAC entries do not appear in the MAC address table because the FastIron Edge Switch transparently manages multiple MAC entries in hardware.
- You can configure private VLANs and dual-mode VLAN ports on the same device. However, the dual-mode VLAN ports cannot be members of Private VLANs.
- A primary VLAN can have multiple ports. All these ports are active, but the ports that will be used depends on the private VLAN mappings. Also, secondary VLANs (isolated and community VLANs) can be mapped to multiple primary VLAN ports. For example:

```
pvlan mapping 901 ethernet 1
pvlan mapping 901 ethernet 2
pvlan mapping 901 ethernet 3
```

- Switch and server trunks are not supported on the FastIron Edge Switches when the ports are part of a private VLAN.

Configuring a Private VLAN

To configure a private VLAN, configure each of the component VLANs (isolated, community, and public) as a separate port-based VLAN.

- Use standard VLAN configuration commands to create the VLAN and add ports.
- Identify the type private VLAN type (isolated, community, or public)
- For the primary VLAN, map the other private VLANs to the port(s) in the primary VLAN

Configuration Rules

- You can use 10/100 and Gigabit Ethernet ports in a private VLAN, but you cannot use POS ports.
- You cannot configure any of the ports in a private VLAN to be members of a trunk group.
- You cannot share a port between a private VLAN and a standard port-based VLAN or protocol VLAN. You can configure private VLANs and standard port-based VLANs and protocol VLANs on the same device, but a port cannot be a member of both a private VLAN and a port-based VLAN or protocol VLAN.

NOTE: Although a private VLAN resides within a port-based VLAN, the VLAN is considered to be exclusively a private VLAN, not a port-based VLAN.

- You cannot use the private VLAN feature and the dual-mode VLAN port feature on the same device.
- The Spanning Tree Protocol (STP) is independent of this feature, and can be enabled or disabled in the individual port-based VLANs. However, private VLANs are not supported with single-instance STP ("single span").
- You can configure only one private VLAN within a given port-based VLAN. Thus, you must configure a separate port-based VLAN for each private VLAN.
- Each private VLAN can have only one primary VLAN.
- Each private VLAN can have multiple isolated or community VLANs. You can use any combination of isolated or community VLANs with the primary VLAN. You do not need to use both isolated and community VLANs in the private VLAN.
- You can configure the primary VLAN before or after you configure the community or isolated VLANs. You are not required to configure a specific type of private VLAN before you can configure the other types.

- The ports in all three types of private VLANs can be tagged or untagged.

NOTE: If the port in the primary VLAN is tagged, you must add the port as a tagged port to each of the isolated and community VLANs. If the port in the primary VLAN is untagged, you do not need to add the port to the isolated and community VLANs.

- The primary VLAN has only one active port. The primary VLAN can have more than one port, but only the lowest-numbered available port is active. The other ports provide redundancy.
- You cannot configure the default VLAN (VLAN 1) as a private VLAN.

Configuring an Isolated or Community Private VLAN

To configure an isolated or a community private VLAN, use the following CLI methods.

USING THE CLI

To configure a community private VLAN, enter commands such as the following:

```
BigIron(config)# vlan 901
BigIron(config-vlan-901)# tagged ethernet 3/5 to 3/6
BigIron(config-vlan-901)# pvlan type community
```

These commands create port-based VLAN 901, add ports 3/5 and 3/6 to the VLAN as tagged ports, then specify that the VLAN is a community private VLAN.

Syntax: tagged ethernet | pos <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] pvlan type community | isolated | primary

The **tagged** or **untagged** command adds the ports to the VLAN.

The **pvlan type** command specifies that this port-based VLAN is a private VLAN.

- **community** – Broadcasts and unknown unicasts received on community ports are sent to the primary port and also are flooded to the other ports in the community VLAN.
- **isolated** – Broadcasts and unknown unicasts received on isolated ports are sent only to the primary port. They are not flooded to other ports in the isolated VLAN.
- **primary** – The primary private VLAN ports are “promiscuous”. They can communicate with all the isolated private VLAN ports and community private VLAN ports in the isolated and community VLANs that are mapped to the promiscuous port.

Configuring the Primary VLAN

Use the following CLI method to configure the primary VLAN.

NOTE: The primary private VLAN has only one active port. If you configure the VLAN to have more than one port, the lowest-numbered port is the active one. The additional ports provide redundancy. If the active port becomes unavailable, the lowest-numbered available port becomes the active port for the VLAN.

USING THE CLI

To configure a primary private VLAN, enter commands such as the following:

```
BigIron(config)# vlan 7
BigIron(config-vlan-7)# untagged ethernet 3/2
BigIron(config-vlan-7)# pvlan type primary
BigIron(config-vlan-7)# pvlan mapping 901 ethernet 3/2
```

These commands create port-based VLAN 7, add port 3/2 as an untagged port, identify the VLAN as the primary VLAN in a private VLAN, and map the other private VLANs to the port(s) in this VLAN.

Syntax: untagged ethernet <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] pvlan type community | isolated | primary

Syntax: [no] pvlan mapping <vlan-id> ethernet <portnum>

The **tagged** or **untagged** command adds the port(s) to the VLAN.

NOTE: You can add the port as a tagged port if needed. If you add the port as a tagged port, you must also add the port as a tagged port to the isolated and community VLANs. See “CLI Example for Figure 22.24” on page 22-68.

The **pvlan type** command specifies that this port-based VLAN is a private VLAN. Specify **primary** as the type.

The **pvlan mapping** command identifies the other private VLANs for which this VLAN is the primary. The command also specifies the primary VLAN ports to which you are mapping the other private VLANs.

- The <vlan-id> parameter specifies another private VLAN. The other private VLAN you want to specify must already be configured.
- The **ethernet** <portnum> parameter specifies the primary VLAN port to which you are mapping all the ports in the other private VLAN (the one specified by <vlan-id>).

Enabling Broadcast or Unknown Unicast Traffic to the Private VLAN

To enhance private VLAN security, the primary private VLAN does not forward broadcast or unknown unicast packets to its community and isolated VLANs. For example, if port 3/2 in Figure 22.24 on page 22-64 receives a broadcast packet from the firewall, the port does not forward the packet to the other private VLAN ports (3/5, 3/6, 3/9, and 3/10).

This forwarding restriction does not apply to traffic from the private VLAN. The primary port does forward broadcast and unknown unicast packets that are received from the isolated and community VLANs. For example, if the host on port 3/9 sends an unknown unicast packet, port 3/2 forwards the packet to the firewall.

If you want to remove the forwarding restriction, you can enable the primary port to forward broadcast or unknown unicast traffic, if desired, using the following CLI method. You can enable or disable forwarding of broadcast or unknown unicast packets separately.

NOTE: On Layer 2 Switches and Layer 3 Switches, you also can use MAC address filters to control the traffic forwarded into and out of the private VLAN. In addition, if you are using a Layer 2 Switch, you also can use ACLs.

USING THE CLI

To configure the ports in the primary VLAN to forward broadcast or unknown unicast traffic received from sources outside the private VLAN, enter the following commands at the global CONFIG level of the CLI:

```
BigIron(config)# pvlan-preference broadcast flood
BigIron(config)# pvlan-preference unknown-unicast flood
```

These commands enable forwarding of broadcast and unknown-unicast packets to ports within the private VLAN. To again disable forwarding, enter a command such as the following:

```
BigIron(config)# no pvlan-preference broadcast flood
```

This command disables forwarding of broadcast packets within the private VLAN.

Syntax: [no] pvlan-preference broadcast | unknown-unicast flood

CLI Example for Figure 22.24

To configure the private VLANs shown in Figure 22.24 on page 22-64, enter the following commands:

```
BigIron(config)# vlan 901
BigIron(config-vlan-901)# tagged ethernet 3/5 to 3/6
BigIron(config-vlan-901)# pvlan type community
BigIron(config-vlan-901)# exit
BigIron(config)# vlan 902
BigIron(config-vlan-902)# tagged ethernet 3/9 to 3/10
BigIron(config-vlan-902)# pvlan type isolated
```

```
BigIron(config-vlan-902)# exit
BigIron(config)# vlan 903
BigIron(config-vlan-903)# tagged ethernet 3/5 to 3/6
BigIron(config-vlan-903)# pvlan type community
BigIron(config-vlan-903)# exit
BigIron(config)# vlan 7
BigIron(config-vlan-7)# untagged ethernet 3/2
BigIron(config-vlan-7)# pvlan type primary
BigIron(config-vlan-7)# pvlan mapping 901 ethernet 3/2
BigIron(config-vlan-7)# pvlan mapping 902 ethernet 3/2
BigIron(config-vlan-7)# pvlan mapping 903 ethernet 3/2
```

This example assumes that the port in the primary private VLAN is untagged. If the port in the primary VLAN is tagged, you must add the port as a tagged port to the isolated and community VLANs, as in the following example:

```
BigIron(config)# vlan 901
BigIron(config-vlan-901)# tagged ethernet 3/5 to 3/6
BigIron(config-vlan-901)# tagged ethernet 3/2
BigIron(config-vlan-901)# pvlan type community
BigIron(config-vlan-901)# exit
BigIron(config)# vlan 902
BigIron(config-vlan-902)# tagged ethernet 3/9 to 3/10
BigIron(config-vlan-902)# tagged ethernet 3/2
BigIron(config-vlan-902)# pvlan type isolated
BigIron(config-vlan-902)# exit
BigIron(config)# vlan 903
BigIron(config-vlan-903)# tagged ethernet 3/5 to 3/6
BigIron(config-vlan-903)# tagged ethernet 3/2
BigIron(config-vlan-903)# pvlan type community
BigIron(config-vlan-903)# exit
BigIron(config)# vlan 7
BigIron(config-vlan-7)# tagged ethernet 3/2
BigIron(config-vlan-7)# pvlan type primary
BigIron(config-vlan-7)# pvlan mapping 901 ethernet 3/2
BigIron(config-vlan-7)# pvlan mapping 902 ethernet 3/2
BigIron(config-vlan-7)# pvlan mapping 903 ethernet 3/2
```


NOTE: You also can specify the primary port and other ports on the same command line. In this example, the command **tagged ethernet 3/2 ethernet 3/5 to 3/6** is equivalent to the pair of **tagged** commands shown above for the same ports.

Dual-Mode VLAN Ports

Configuring a tagged port as a **dual-mode** port allows it to accept and transmit both tagged traffic and untagged traffic at the same time. A dual-mode port accepts and transmits frames belonging to VLANs configured for the port, as well as frames belonging to the default VLAN (that is, untagged traffic).

For example, in Figure 22.25, port 2/11 is a dual-mode port belonging to VLAN 20. Traffic for VLAN 20, as well as traffic for the default VLAN, flows from a hubs to this port. The dual-mode feature allows traffic for VLAN 20 and untagged traffic to go through the port at the same time.

Figure 22.25 Dual-mode VLAN port example



To enable the dual-mode feature on port 2/11 in Figure 22.25:

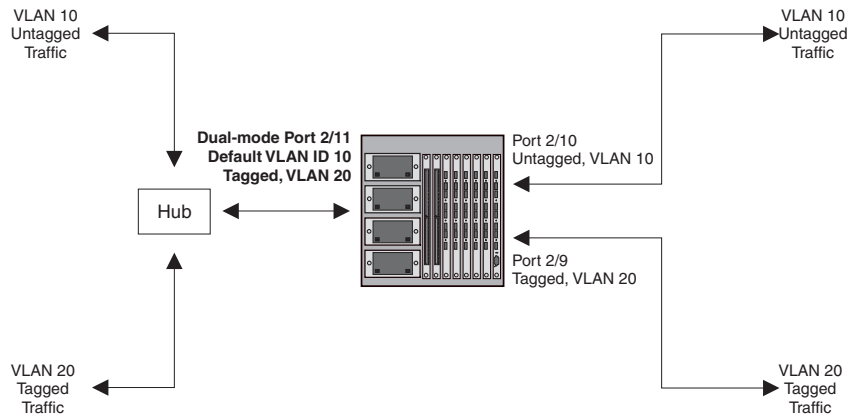
```
BigIron(config)# vlan 20
BigIron(config-vlan-20)# tagged e 2/11
BigIron(config-vlan-20)# tagged e 2/9
BigIron(config-vlan-20)# int e 2/11
BigIron(config-if-e100-2/11)# dual-mode
BigIron(config-if-e100-2/11)# exit
```

Syntax: [no] dual-mode

In releases prior to 07.6.01, a dual-mode port accepts and transmits frames belonging to VLANs configured for the port, as well as frames belonging to the DEFAULT-VLAN (VLAN 1). Traffic for the DEFAULT-VLAN is transmitted untagged, and traffic for other VLANs is tagged.

Starting with release 07.6.01, you can configure a dual-mode port to transmit traffic for a specified VLAN (other than the DEFAULT-VLAN) as untagged, while transmitting traffic for other VLANs as tagged. Figure 22.26 illustrates this enhancement.

Figure 22.26 Specifying a default VLAN ID for a dual-mode port



In Figure 22.26, tagged port 2/11 is a dual-mode port belonging to VLANs 10 and 20. The default VLAN assigned to this dual-mode port is 10. This means that the port transmits tagged traffic on VLAN 20 (and all other VLANs to which the port belongs) and transmits untagged traffic on VLAN 10.

The dual-mode feature allows tagged traffic for VLAN 20 and untagged traffic for VLAN 10 to go through port 2/11 at the same time. A dual-mode port transmits only untagged traffic on its default VLAN (that is, either VLAN 1, or a user-specified VLAN ID), and only tagged traffic on all other VLANs.

The following commands configure VLANs 10 and 20 in Figure 22.26. Tagged port 2/11 is added to VLANs 10 and 20, then designated a dual-mode port whose specified default VLAN is 10. In this configuration, port 2/11 transmits only untagged traffic on VLAN 10 and only tagged traffic on VLAN 20.

```
BigIron(config)# vlan 10 by port
BigIron(config-vlan-10)# untagged e 2/10
BigIron(config-vlan-10)# tagged e 2/11
BigIron(config-vlan-10)# exit

BigIron(config)# vlan 20 by port
BigIron(config-vlan-20)# tagged e 2/9
BigIron(config-vlan-20)# tagged e 2/11
BigIron(config-vlan-20)# exit

BigIron(config)# int e 2/11
BigIron(config-if-e100-2/11)# dual-mode 10
BigIron(config-if-e100-2/11)# exit
```

Syntax: [no] dual-mode [<vlan-id>]

Notes:

- If you do not specify a <vlan-id> in the **dual mode** command, the port's default VLAN is set to 1. The port transmits untagged traffic on the DEFAULT-VLAN.
- The dual-mode feature is disabled by default. Only tagged ports can be configured as dual-mode ports.
- In trunk group, either all of the ports must be dual-mode, or none of them can be.

The **show vlan** command displays a separate row for dual-mode ports on each VLAN. For example:

```
BigIron(config)# show vlan
Total PORT-VLAN entries: 3
Maximum PORT-VLAN entries: 16

legend: [S=Slot]

PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S1) 1 2 3 4 5 6 7 8
  Untagged Ports: (S2) 1 2 3 4 5 6 7 8 12 13 14 15 16 17 18 19
  Untagged Ports: (S2) 20 21 22 23 24
  Tagged Ports: None
  Uplink Ports: None
DualMode Ports: None
PORT-VLAN 10, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: (S2) 10
  Tagged Ports: None
  Uplink Ports: None
DualMode Ports: (S2) 11
PORT-VLAN 20, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: None
  Tagged Ports: (S2) 9
  Uplink Ports: None
DualMode Ports: (S2) 11
```

JetCore Module Hardware Flooding for Layer 2 Multicast and Broadcast Packets

You can configure JetCore modules to perform hardware flooding for Layer 2 multicast and broadcast packets. Layer 2 multicast packets have a multicast address in the destination MAC address field.

You enable hardware flooding for Layer 2 multicast and broadcast packets on a per-VLAN basis. For example:

```
BigIron(config)#
BigIron(config)# vlan 2
BigIron(config-vlan-2)# multicast-flooding
BigIron(config-vlan-2)# exit
```

Syntax: multicast-flooding

After entering the **multicast-flooding** command for a VLAN, you must reboot the Foundry device to activate the feature.

Notes:

- This feature is supported only on JetCore modules and the 10 Gigabit Ethernet module.
- This feature cannot be enabled on an empty VLAN; the VLAN must already have ports assigned to it prior to enabling this feature.
- This feature is not supported on Layer 3 protocol-based VLANs.
- This feature is not supported on private VLANs.
- You cannot enable this feature on the designated management VLAN for the device.
- If you enable this feature on a VLAN that includes a trunk group, hardware flooding for Layer 2 multicast and broadcast packets occurs only on the trunk group's primary port. Multicast and broadcast traffic for the other ports in the trunk group is handled by software.

JetCore Broadcast Suppression

NOTE: This feature is supported on JetCore devices running Service Provider release 09.1.00 and higher.

The broadcast suppression feature allows you to configure a Foundry device to rate limit broadcast traffic that is processed by the CPU.

Normally, Foundry devices process traffic that has a destination MAC address of FFFF.FFFF.FFFF in software; that is, by sending it to the CPU. If a VLAN is subject to a large amount of broadcast traffic, it could place excessive strain on the CPU. When excessive strain is placed on the CPU, it can cause the performance of the Foundry device to degrade, and the device could stop forwarding traffic altogether. By configuring the broadcast suppression feature to rate limit broadcast traffic sent to the CPU, you can prevent the CPU from becoming overloaded during broadcast storms.

You can specify a rate limit (in bits per second) for broadcast traffic sent to the CPU, and also specify a VLAN ID to which the rate should apply on a given port.

Broadcast suppression can be configured for incoming traffic on a physical port only. It cannot be used to rate limit outgoing traffic, nor can it be applied to a VE or a trunk port.

To specify a rate limit for broadcast traffic on interface 1/1, enter the following commands:

```
BigIron(config)# int e 1/1
BigIron(config-if-e1000-1/1)# bcast-suppress any 100000000
```

Syntax: [no] bcast-suppress <vlan-id> | any <avg-rate-in-bps>

The <vlan-id> parameter specifies the VLAN ID to which the rate applies. You can enter a specific VLAN ID or use the keyword **any**, which causes the rate limit to apply to all VLANs to which the port belongs.

The <avg-rate-in-bps> is the rate limit for incoming broadcast traffic on the port. You can specify a value can from 262144 (256Kbps) up to the maximum line rate of the port. For example, for a 100Mbps port, the maximum value is 100,000,000 (100Mbps).

Notes:

- You can configure the broadcast suppression feature on ports that also have one or more Layer 2 ACL rate limiting policies configured.
- The feature cannot be configured on a port that also has Layer 2 ACLs configured.
- You cannot configure the feature on ports that also have IP ACLs or IP ACL-based rate limiting policies configured.

Unicast Flooding on VLAN Ports (BigIron MG8 and NetIron 40G Software Release 02.0.00 and Later)

BigIron MG8 and NetIron 40G software release 02.0.00 and later allows Terathon devices to perform hardware flooding for Layer 2 unknown unicast packets on all ports on a VLAN. When this feature is enabled on a VLAN a “catch-all” CAM entry is added for the VLAN entry.

This CAM entry matches all unicast packets that have not been matched in other CAM entries. This CAM entry forces the packet to be flooded in hardware to the VLAN broadcast domain. In order for software to add CAM entries for MAC addresses that are eventually learned, a few packets need to be sent to the CPU from time to time. This is done by removing and adding the match-all CAM entry at fixed intervals.

To enable unicast flooding on a VLAN ports, enter commands such as the following:

```
BigIron MG8(config)# vlan 2
BigIron MG8(config-vlan-2)# unknown-unicast-flooding
BigIron MG8(config-vlan-2)# exit
BigIron MG8(config)# reload
```

Syntax: [no] unknown-unicast-flooding

You must reboot the Terathon device to activate the feature.

Configuration Considerations

Note the following configuration limitations for this feature:

- This feature is not supported on Layer 3 protocol-based VLANs.
- You cannot enable this feature on the designated management VLAN for the device.
- The **system-max vlan-multicast-flooding** command needs to be set to reserve CAM space for the unknown-unicast flooding CAM entries. Only when this is done can the configuration proceed.

Configuring VLAN Translation (BigIron MG8 and NetIron 40G Software Release 02.0.00 and later)

This feature is supported on all BigIron MG8 Interface modules except the following:

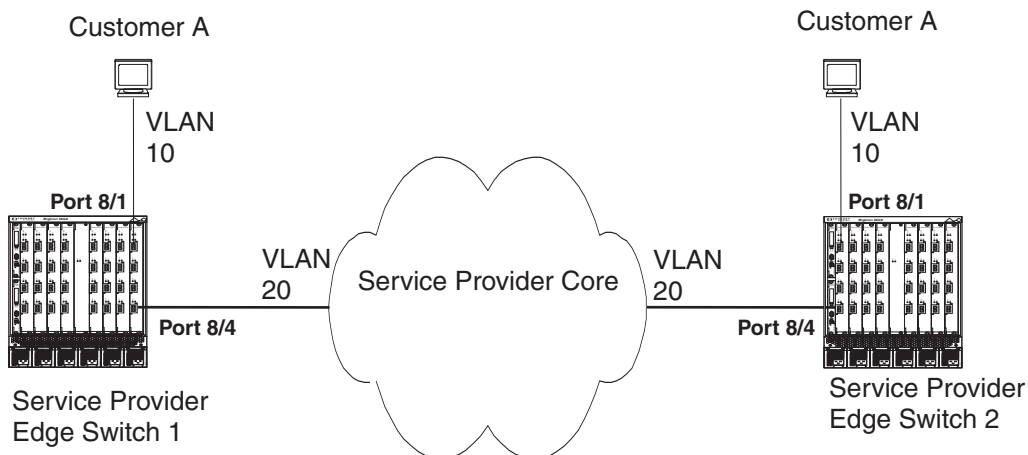
- 10G x 2 Interface modules of the type BIMG8-10 x 2 or BIMG8-10 x 2-v6
- 10G x 4 Interface modules of the type BIMG8-10 x 4 or BIMG8-10 x 4-v6

If using a 10G x 2 or 10G x 4 Interface module, it must be of the type BIMG8-10 x 2-v6-A or BIMG8-10 x 4-v6-A. For a more detailed explanation of the 10G x 2 or 10G x 4 Interface module versions, see the table in the “Support for Outbound ACLs and IPv6” section in the Product Overview chapter of the *Foundry BigIron Mg8 Switch Installation and Basic Configuration Guide*; and for the method to determine the version of the 10G x 2 or 10G x 4 Interface module that you have installed, see the “Determining the Type of Interface Module That You Have Installed” section in the Product Overview chapter of the *Foundry BigIron Mg8 Switch Installation and Basic Configuration Guide*.

VLAN Translation allows traffic from one VLAN to be transported across a different VLAN. Under this feature, packets from the original VLAN have their VLAN ID changed at the ingress port of the VLAN that is performing the translation. When they reach the egress point on the VLAN that performed the translation, the VLAN ID is translated back to its original ID.

This feature is useful for service providers who need to carry traffic from different customers across their network while preserving the VLAN ID and priority information of the customer’s network. For instance, in the following example Customer A has two geographically divided networks in the same IP subnet that are both in VLAN 10. The service provider uses VLAN 20 to route the traffic between these two geographically divided portions of VLAN 10. Each of the service provider edge switches perform VLAN translation to translate the VLAN ID between VLAN 10 and VLAN 20.

Figure 22.27 VLAN Translation Example



Configuration Considerations

1. A port must be a member of the translated VLAN before it can be used in its VLAN translation group.
2. A port-VLAN pair can only be used in one VLAN translation group.
3. Up to 4096 VLAN translation groups can be configured on a switch.
4. VLAN translation should not be combined on the same port with any Layer 4 features such as ACLs, policy-based routing, or ACL-based rate-limiting.
5. Only the primary port of a trunk group can be added to a VLAN translation group. Other ports are then automatically included in the VLAN translation group.
6. If VLAN translation is enabled on a port, hardware forwarding of unknown unicast packets should not be enabled on that port.
7. This feature is currently only supported on 40-port modules.
8. VLAN translation cannot be configured on virtual ports.

CLI Command for VLAN Translation

The following command required for VLAN Translation configures a VLAN Translation group and assigns interfaces to it.

This command creates a VLAN Translation Group. Packets that arrive on a port that is configured to be in a VLAN translation group are forwarded based on the destination MAC address. First the destination MAC address in the translated VLAN is used. If the port on which the destination MAC address is learned is a member of the translated VLAN and configured in the same VLAN translation group, then the packet is forwarded to that port. The VLAN ID is replaced with the translated VLAN ID. If the port is not part of the VLAN translation group then the destination MAC address in the ingress port's VLAN is used for packet forwarding. If the destination MAC address does not exist, then the packet is flooded to the ingress port's VLAN as well as the translated VLAN.

Syntax:

```
vlan-translate-group <number>  
(config-vlan-translate-group)#port <port_id> vlan-id <vlan_id>
```

number is the decimal number that you assign to a VLAN translation group.

port_id is the slot/port number that you want to configure in the VLAN translation group.

vlan_id is the VLAN number that the port specified in port_id is assigned to. The port must be separately configured in this VLAN.

EXAMPLE:

The following example creates vlan-translate-group number 1 and adds port 1/2 in VLAN 10 and port 1/5 in VLAN 20 to it.

```
MG8(config)# vlan-translate-group 1  
MG8(config-vlan-translate-group)# port 1/2 vlan-id 10  
MG8(config-vlan-translate-group)# port 1/5 vlan-id 20
```

Configuration Example

This section describes the configuration required to enable the configuration described in Figure 22.27 for Service provider edge switches 1 and 2.

Service Provider Edge Switch 1 Configuration

Each port used for the VLAN translation must first be configured in its VLAN as shown below.

```
MG8(config)# vlan 10  
MG8(config-vlan-10)# untagged ethernet 8/1  
MG8(config)# vlan 20  
MG8(config-vlan-20)# tagged ethernet 8/4
```

Each port used for VLAN translation must be added to a VLAN Translate group as shown below.

```
MG8(config)# vlan-translate-group 1
MG8(config-vlan-translate-group-1)# port 8/1 vlan-id 10
MG8(config-vlan-translate-group-1)# port 8/4 vlan-id 20
```

Service Provider Edge Switch 2 Configuration

Each port used for the VLAN translation must first be configured in it's VLAN as shown below.

```
MG8(config)# vlan 10
MG8(config-vlan-10)# untagged ethernet 8/1
MG8(config)# vlan 20
MG8(config-vlan-20)# tagged ethernet 8/4
```


Each port used for VLAN translation must be added to a VLAN Translate group as shown below.

```
MG8(config)# vlan-translate-group 1
MG8(config-vlan-translate-group-1)# port 8/1 vlan-id 10
MG8(config-vlan-translate-group-1)# port 8/4 vlan-id 20
```

Configuring Inner VLAN Translation with Super Aggregated VLANs (BigIron MG8 and NetIron 40G Software Release 02.0.00 and Later)

Inner VLAN translation is supported for packets with two VLAN tags. VLAN translation can be performed on the inner VLAN tag. In the following example, packets from customers A and B are tagged with VLAN 10 and etype 8100. Packets from customer A enter Service Provider Edge Switch 1 in VLAN 100, and packets from customer B enter Service Provider Edge Switch 1 in VLAN 200. The etype of both the ingress ports is set to 9100. The egress port on Service Provider Edge Switch 1 is contained within both VLANs 100 and 200 with etype set to 9100. Packets sent out on the egress port have two VLAN tags. On the ingress port of Service Provider Edge Switch 2 inner VLAN translation is set to translate traffic tagged with an outer VLAN tag of 100 and an inner tag of VLAN 10 to VLAN 101. Inner VLAN translation is also set to translate traffic tagged with an outer VLAN tag of 200 and an inner tag of VLAN 10 to VLAN 102. When the traffic from Service Provider Edge Switch 1 arrives at Service Provider Edge Switch 2, packets with outer VLAN tag 100 and inner VLAN tag 10 are translated to inner VLAN tag 101. Packets with outer VLAN tag 200 and inner VLAN tag 10 are translated to inner VLAN tag 102. The outer tag remains unchanged in both cases. The packet forwarding is done based on the outer VLAN tag.

Figure 23 VLAN Translation with Super Aggregated VLANs Example



Configuration Considerations

1. Inner-VLAN translation cannot be configured on virtual ports.
2. The **cam-partition block** command must be set for inner-VLAN translation to work.
3. The port on which the inner-VLAN translation is configured, must be a member of the outer VLAN.
4. VLAN translation and inner-VLAN translation cannot be enabled on a port at the same time.

5. If inner-VLAN translation is enabled on a port, hardware forwarding of unknown unicast packets should not be enabled on that port.
6. For a given interface, the (outer-VLAN, inner-VLAN) pair in the translation rule must be unique.
7. For trunk ports, inner-VLAN translation can be configured on the primary ports only. The configuration then applies to all ports of the trunk port.
8. There is no limit on the number of inner VLAN translation policies that can be applied to a port.
9. The trunk is rejected if any of the trunk's have VLAN or inner-VLAN translation configured.

CLI Command to Configure an Interface for VLAN Translation on a Super Aggregated VLAN

The following command is required to apply VLAN Translation for a Super Aggregated VLAN.

This command creates a VLAN translation rule on an interface used in a Super Aggregated VLAN.

Syntax: `inner-vlan-translate <outer-vlan-tag> <inner-vlan-tag> <translation-vlan-tag>`

outer-vlan-tag specifies outer vlan tag of the packet with two VLAN tags. This VLAN tag is maintained with the packets through the translation process.

inner-vlan-tag specifies inner vlan tag of the packet that needs to be translated.

translation-vlan-tag specifies vlan tag that the inner VLAN tag will be translated to.

EXAMPLE:

The following example applies a VLAN translation rule to interface 1/2 to translate traffic with an outer VLAN tag of 100 and an inner VLAN tag of 10 to an outer VLAN tag of 101.

```
BigIron MG8(config)# interface ethernet 1/2
BigIron MG8(config-if-e1000-1/2)# inner-vlan-translate 100 10 101
```

Configuration Example

This section describes the syntax required to enable the configuration described in Figure 23 for Service provider edge switches 1 and 2.

Service Provider Edge Switch 1 Configuration

Each port used for the VLAN translation must first be configured in its VLAN as shown below.

```
MG8(config)# vlan 100
MG8(config-vlan-100)# untagged ethernet 1/1
MG8(config-vlan-100)# tagged ethernet 8/2
MG8(config)# vlan 200
MG8(config-vlan-200)# untagged ethernet 1/4
MG8(config-vlan-200)# tagged ethernet 8/2
```

For Super Aggregated VLANs (SAV), VLAN translation is configured under an interface as an inbound feature. For SAVs, the outer VLAN, inner VLAN and translation VLAN must be configured. The configuration for interface 8/2 in the example in Figure 23 is shown below.

```
MG8(config)# interface ethernet 8/2
MG8(config-if-e1000-8/2)# inner-vlan-translate 100 101 10
MG8(config-if-e1000-8/2)# inner-vlan-translate 200 102 10
```

Service Provider Edge Switch 2 Configuration

Each port used for the VLAN translation must first be configured in its VLAN as shown below.

```
MG8(config)# vlan 100 by port
MG8(config-vlan-100)# untagged ethernet 8/1
MG8(config-vlan-100)# tagged ethernet 1/2
MG8(config)# vlan 200 by port
MG8(config-vlan-200)# untagged ethernet 8/4
```

```
MG8(config-vlan-200)# tagged ethernet 1/2
```

For Super Aggregated VLANs (SAV), VLAN translation is configured under an interface as an inbound feature. For SAVs, the outer VLAN, inner VLAN and translation VLAN must be configured. The configuration for interface 1/2 in the example in Figure 23 is shown below.

```
MG8(config)# interface ethernet 1/2
MG8(config-if-e1000-1/2)# inner-vlan-translate 100 10 101
MG8(config-if-e1000-1/2)# inner-vlan-translate 200 10 102
```

Configuring MAC VLANs (Stackable FastIron Backbone Layer 2 Switch only)

On the Stackable FastIron Backbone Layer 2 Switch, you can configure a list of source MAC addresses and switch input ports, and associate each address with a VLAN and, optionally, a QoS priority. When the Layer 2 Switch receives a packet that matches an entry in the list, the switch assigns that packet to the specified VLAN. In addition, if the destination port for the packet is a tagged port, the switch adds the 802.1q tag to the packet.


To configure the switch for MAC VLAN support, you create a text file containing a list of MAC VLAN entries. The MAC VLAN list is a text file that the switch reads from a TFTP server during startup. The MAC VLAN file contains entries in the following format:

Syntax: ext mac-vlan <source-mac-address> <vlan-id> ethernet <input-port> [priority normal | high]

The switch adds the MAC VLAN entries in the file to its MAC table. When the switch receives a packet, the switch checks the packet's contents against the MAC VLAN configuration list.

If a packet matches the source MAC address and input port of an entry in the list, the switch uses the VLAN ID of the MAC VLAN entry as the packet's VLAN ID. In addition, if the packet is destined an output port that is tagged, the switch tags the packet, adding the four-byte 802.1p tag field to the end of the packet.

Further, if the MAC VLAN list specifies a priority for the packet, the switch places the packet in the specified queue on the output port. For example, if the entry is marked with high priority, the switch places the packet in the high-priority queue. Figure 23.1 shows an example of how the MAC VLAN list is used.

Figure 23.1 VLAN list on Foundry Stackable Layer 2 Switch**MAC VLAN File:**

Source MAC, VLAN ID, input port, priority

```
ext mac-vlan 1145.5563.a101 10 ethernet 1
ext mac-vlan 1145.5563.a102 10 ethernet 1
ext mac-vlan 1145.5563.a103 20 ethernet 4
ext mac-vlan 1145.5563.c11f 20 ethernet 4 priority high
```

As shown in this example, the MAC VLAN list contains four entries. For simplicity, each of the MAC addresses in the list in this example belongs to a host attached to the switch. Each of the MAC addresses belongs to a host attached to the switch port indicated in the file. When the switch receives a packet from one of the MAC addresses, the packet matches an entry in the MAC VLAN list. When this occurs, the switch adds the packet to VLAN specified in the MAC VLAN entry. Moreover, if the output port for the packet is a tagged port, the switch adds the four-byte tag field to the end of the packet thus tagging the packet.

NOTE: If the switch receives a packet with a source MAC address that is contained in the MAC VLAN list, but that switch receives that packet on a different port than the one specified in the MAC VLAN list, the packet does not match the list and is switched normally. The MAC VLAN switching described in this section is not performed for the packet.

In this example, suppose host 1145.5563.a101 sends a packet destined for a host on the router attached to port 10. The switch compares the packet's source MAC address and input port against the entries in the MAC VLAN field and finds a match. The switch then assigns that packet to the VLAN indicated in the MAC VLAN file. In this case, the VLAN ID is 10. Since the packet's destination is reached through switch port 10, a tagged port, the switch also adds the tag to the packet before placing the packet in the output queue on port 10. The MAC VLAN list does not specify a priority for the packet, so the switch uses the priority associated with the port. In this case, the priority is normal so the switch places the packet in the normal priority queue.

If the switch receives a packet from host 1145.5563.C11f on port 4, the switch assigns the packet to VLAN 20 and places the packet in the high priority queue of its output port. If the output port is tagged, the switch also adds the tag to the packet.

Configuring a MAC VLAN List

To configure the MAC VLAN list:

1. Create a new text file. For convenience, you might want to create the text file on the TFTP server from which the switch will download the MAC VLAN list. Otherwise, you can copy the file onto the TFTP server after you add the commands to the file.

2. Enter the following command in the file to add a MAC VLAN entry to the file:

```
ext mac-vlan <mac-addr> <vlan-id> ethernet <portnum> [priority <num>]
```

The <mac-addr> parameter specifies the source MAC address you want the switch to check for.

The <vlan-id> parameter specifies the VLAN to which you want the switch to assign packets that match the other values in this command.

The <portnum> parameter specifies the source port. A packet matches this VLAN entry only if the source MAC address and source port match the values you specify.

The **priority** <num> parameter optionally changes the QoS priority for packets that match this entry. You can specify normal or high-priority. The default is normal.

3. Repeat step 2 for each source MAC address and source port.
4. Save the file and copy it to the TFTP server.
5. Enter the following command on the switch at the global CONFIG level of the CLI to identify the MAC VLAN file name and location:

```
ext get config-file <ip-addr> <external-file-name>
```

The <ip-addr> specifies the IP address of the TFTP server.

The <external-file-name> parameter specifies the name of the text file containing the **ext mac-vlan** commands.

6. Enter the following command on the switch to save the configuration change to flash memory:

```
write memory
```

Loading a MAC VLAN List

The switch automatically uses the information you enter with the **ext get config-file** command to load the MAC VLAN list the next time you reload the switch. However, you can load a MAC VLAN list at any time by entering the following command at the Privileged EXEC level of the CLI.

Syntax: ext refresh config-file <ip-addr> <external-file-name>

The <ip-addr> parameter specifies the IP address of the TFTP server on which you placed the MAC VLAN file.

The <external-file-name> parameter specifies the name of the MAC VLAN file.

This command adds the entries in the MAC VLAN list to the MAC table. Existing entries in the table are not cleared.

Specifying a Default VLAN for MAC Addresses That Are Not in the MAC VLAN List

You can specify a default port-based VLAN for MAC addresses that the switch receives from sources other than the MAC VLAN list you configure. For example, if a host plugs into the network and that host's MAC address is not in the MAC VLAN list, the switch assigns that VLAN to the default VLAN you specify.

You can specify the default VLAN on an individual interface basis. You can specify the same VLAN or different VLANs on each interface.

To specify the default VLAN for an interface, enter commands such as the following:

```
FastIron(config)# int e 1
FastIron(config-if-1)# mvlan-mode 2
FastIron(config-if-1)# write memory
```

Syntax: `mvlan-mode <vlan-id>`

Clearing MAC VLAN Entries from the MAC Table

You can clear entries that have been added to the MAC table from a MAC VLAN list. To clear entries, enter the following command at the Privileged EXEC level of the CLI.

Syntax: `ext clear mac-vlan [mac <mac-addr> <mask>] [vlan <vlan-id>]`

If you enter **ext clear mac-vlan** without any of the optional parameters, all the entries added from MAC VLAN lists are cleared.

The **mac** <mac-addr> <mask> parameter clears only the entries that match the specified address and mask.

The **vlan** <vlan-id> parameter clears only the entries that match the specified VLAN.

Configuring VLAN Byte Accounting

Beginning with release 02.2.01 of the NetIron 40G or BigIron MG8, you can configure the device to account for the number of bytes received by all the member ports of a VLAN. The bytes will include the preamble and the minimum inter-frame gap in Ethernet. The byte counts can then be viewed using the **show vlan** command. VLAN byte accounting is disabled by default.

Considerations When Configuring VLAN Byte Accounting

- VLAN byte accounting cannot be enabled for the default and control VLANs.
- The number of VLANs on which byte accounting can be enabled system-wide is restricted by the number of VLANs with byte accounting enabled on a given packet processor and the number of rate limiting policies enabled on the same packet processor ports.
- On a given packet processor, the total number of VLANs with byte accounting enabled and number of rate limiting policies based on ACLs and VLANs is dependent on the interface module. See Table 23.2 for details.
- If a port's VLAN has byte accounting enabled, you cannot enable rate limiting on that port. Similarly, if a port has rate limiting enabled, you cannot enable VLAN byte accounting on that port's VLAN.
- Clearing the rate limiting counters using **clear rate-limit counters** will also clear VLAN byte-accounting counters. It is recommended that when using rate limiting along with VLAN byte accounting, use individual port rate limiting counter clear command.

Configuring VLAN Byte Accounting

To enable VLAN accounting on a specified VLAN, use the following commands:

```
BigIron MG8(config)# vlan 10
BigIron MG8(config-vlan-10)# byte-accounting
```

Syntax: `[no] byte-accounting`

Displaying VLAN Byte Accounting Information

To display VLAN accounting information for all VLANs configured on a router, use the **show vlan** command as shown:

```
BigIron MG8# show vlan

Configured PORT-VLAN entries: 2
Maximum PORT-VLAN entries: 512
Default PORT-VLAN id: 1

PORT-VLAN 1, Name DEFAULT-VLAN, Priority Level0
L2 protocols   : NONE
Untagged Ports : ethe 1/1 to 1/40 ethe 2/1 to 2/4

PORT-VLAN 10, Name [None], Priority Level0
L2 protocols   : NONE
Tagged Ports   : ethe 1/2 to 1/5
Bytes received : 18527
```

To display VLAN accounting information for a specific VLAN, use the **show vlan <vlan>** command as shown:

```
BigIron MG8# show vlan 10

PORT-VLAN 10, Name [None], Priority Level0
L2 protocols   : NONE
Tagged Ports   : ethe 1/2 to 1/5
Bytes received : 5626
```

Syntax: show vlan [<vlan-id>]

The byte accounting statistics are displayed above in **bold**. Table 23.1 describes the Byte Accounting statistics displayed when using the **show vlan** command. For information about all other parameters, see the *Foundry Switch and Router Installation and Basic Configuration Guide*.

Table 23.1: VLAN Byte Accounting in Show VLAN

This Field...	Displays...
Bytes received	This field displays the number of bytes received by all member ports of all VLANs configured on the router if the command show vlan is used or if the <VLAN> variable is used with the show vlan command

Maximum Number of Rate Limiting Policies and VLANs with Byte Accounting

In this release, the maximum number of ACL-based, and VLAN-based rate limiting policies that can be configured on ports controlled by the same packet processor also depends on the number of VLANs with byte accounting enabled on the same packet processor.

On a given packet processor (PPCR), the total of:

- Number of VLANs with byte accounting enabled**
- +**
- Number of rate limiting policies based on ACLs and VLANs**

cannot exceed the maximum number per-PPCR as specified in Table 23.2.

Table 23.2: Maximum # of rate limiting policies and VLANs w/ byte accounting permitted per-PPCR

Module Type	PPCR Number	Ports supported by PPCR	Max # of rate limiting policies based on ACLs and VLANs + number of VLANs w/ byte accounting enabled
4 x 10G	PPCR 1	1	126
	PPCR 2	2	126
	PPCR 3	3	126
	PPCR 4	4	126
8 x 10G	PPCR 1	1 - 2	125
	PPCR 2	3 - 4	125
	PPCR 3	5 - 6	125
	PPCR 4	7 - 8	125
40 x 1G	PPCR 1	1 - 10	117
	PPCR 2	11 - 20	117
	PPCR 3	21 - 30	117
	PPCR 4	31 - 40	117
60 x 1G	PPCR 1	1 - 20	107
	PPCR 2	21 - 40	107
	PPCR 3	41 - 60	107

Clearing Counters

To clear the byte counter for a VLAN, enter a command such as the following:

```
BigIron MG8(config)#clear vlan byte-accounting 10
```

You can also enter the following command to clear the byte counter for all VLANs

```
BigIron MG8(config)#clear vlan byte-accounting all-vlans
```

Syntax: clear vlan byte-accounting <vlan-id> | all-vlans

Enter a VLAN ID if you want to clear the byte counters for a specific VLAN. Enter **all-vlans** to clear the byte counters for all VLANs.

Configuring VLANs Using the Web Management Interface

Use the procedures in the following sections to configure VLANs using the Web management interface.

Configuring a Port-Based VLAN

1. Log on to the device using a valid user name and password for read-write access.

2. If you have not already enabled OSPF, enable it by clicking on the Enable radio button next to OSPF on the System configuration dialog, then clicking Apply to apply the change.
3. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
4. Click on the plus sign next to VLAN in the tree view to expand the list of VLAN option links.
5. Click on the [Port](#) link.
 - If the device does not have any port-based VLANs, the Port VLAN configuration panel is displayed, as shown in the following example.
 - If at least one port-based VLAN is already configured and you are adding a new one, click on the [Add Port VLAN](#) link to display the Port VLAN configuration panel, as shown in the following example.
 - If you are modifying an existing port-based VLAN, click on the Modify button to the right of the row describing the VLAN to display the Port VLAN configuration panel, as shown in the following example.

Port VLAN

VLAN Id:	<input type="text" value="2"/>
Name:	<input type="text"/>
QoS:	<input type="text" value="0"/>
Spanning Tree	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Router Interface:	<input type="text" value="None"/>
Port members:	<input type="button" value="Select Port Members"/>

[\[Show\]](#)
[\[Protocol VLAN\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

6. Enter the VLAN ID and optionally the name.
7. If you want to assign the VLAN to a different Quality of Service (QoS) priority, select the priority from the QoS field's pulldown menu. For more information, see the "Configuring IronClad Quality of Service" chapter in the *Foundry Enterprise Configuration and Management Guide*.
8. Select Enable or Disable next to Spanning Tree to enable or disable the feature on this VLAN.
9. Select the virtual routing interface (router interface) if applicable.
10. Click the Select Port Members button to display the following panel.

Port Members

Row 1 <input type="checkbox"/>	<input checked="" type="checkbox"/> 1/1	<input checked="" type="checkbox"/> 1/2	<input checked="" type="checkbox"/> 1/3	<input checked="" type="checkbox"/> 1/4	<input checked="" type="checkbox"/> 1/5	<input checked="" type="checkbox"/> 1/6	<input type="checkbox"/> 1/7	<input type="checkbox"/> 1/8
Row 2 <input type="checkbox"/>	<input type="checkbox"/> 3/1	<input type="checkbox"/> 3/2	<input type="checkbox"/> 3/3	<input type="checkbox"/> 3/4	<input type="checkbox"/> 3/5	<input type="checkbox"/> 3/6	<input type="checkbox"/> 3/7	<input type="checkbox"/> 3/8
Row 3 <input type="checkbox"/>	<input type="checkbox"/> 3/9	<input type="checkbox"/> 3/10	<input type="checkbox"/> 3/11	<input type="checkbox"/> 3/12	<input type="checkbox"/> 3/13	<input type="checkbox"/> 3/14	<input type="checkbox"/> 3/15	<input type="checkbox"/> 3/16
Row 4 <input type="checkbox"/>	<input type="checkbox"/> 3/17	<input type="checkbox"/> 3/18	<input type="checkbox"/> 3/19	<input type="checkbox"/> 3/20	<input type="checkbox"/> 3/21	<input type="checkbox"/> 3/22	<input type="checkbox"/> 3/23	<input type="checkbox"/> 3/24
Row 5 <input type="checkbox"/>	<input type="checkbox"/> 4/1	<input type="checkbox"/> 4/2	<input type="checkbox"/> 4/3	<input type="checkbox"/> 4/4	<input type="checkbox"/> 4/5	<input type="checkbox"/> 4/6	<input type="checkbox"/> 4/7	<input type="checkbox"/> 4/8
Row 6 <input type="checkbox"/>	<input type="checkbox"/> 4/9	<input type="checkbox"/> 4/10	<input type="checkbox"/> 4/11	<input type="checkbox"/> 4/12	<input type="checkbox"/> 4/13	<input type="checkbox"/> 4/14	<input type="checkbox"/> 4/15	<input type="checkbox"/> 4/16
Row 7 <input type="checkbox"/>	<input type="checkbox"/> 4/17	<input type="checkbox"/> 4/18	<input type="checkbox"/> 4/19	<input type="checkbox"/> 4/20	<input type="checkbox"/> 4/21	<input type="checkbox"/> 4/22	<input type="checkbox"/> 4/23	<input type="checkbox"/> 4/24

- Select the ports you are placing in the VLAN. To select a row, click on the checkbox next to the row number, then click on the Select Row button.

NOTE: Ports highlighted in grey are members of a trunk group. The port right before the grey ports is the master port for that trunk group.

- When you finish selecting the ports, click on the Continue button to return to the Port VLAN configuration dialog.
- Click the Add button (to add a new VLAN) or the Modify button (if you are modifying an existing VLAN) to save the change to the device's running-config file.
- Select the [Save](#) link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring a Protocol-Based VLAN

This procedure describes how to configure a protocol-based VLAN. To configure an IP subnet VLAN, IPX network VLAN, or AppleTalk cable VLAN, see the sections following this one.

- Log on to the device using a valid user name and password for read-write access.
- Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
- Click on the plus sign next to VLAN in the tree view to expand the list of VLAN option links.
- Click on the [Protocol](#) link.
 - If the device does not have any protocol VLANs, the Protocol VLAN configuration panel is displayed, as shown in the following example.
 - If at least one protocol VLAN is already configured and you are adding a new one, click on the [Protocol](#) link to display the Protocol VLAN configuration panel.
 - If you are modifying an existing protocol VLAN, click on the Modify button to the right of the row describing the VLAN to display the configuration panel for the type of VLAN you are modifying. The following example shows the Protocol VLAN configuration dialog, used for configuring a protocol VLAN (not an IP subnet, IPX network, or AppleTalk cable VLAN).

VLAN Id:	1
VLAN Port members:	1/7, 1/8, 3/1, 3/2, 3/3, 3/4, 3/5, 3/6, 3/7, 3/8, 3/9, 3/10, 3/11, 3/12, 3/13, 3/14 3/15, 3/16, 3/17, 3/18, 3/19, 3/20, 3/21, 3/22, 3/23, 3/24, 4/1, 4/2, 4/3, 4/4, 4/5, 4/6 4/7, 4/8, 4/9, 4/10, 4/11, 4/12, 4/13, 4/14, 4/15, 4/16, 4/17, 4/18, 4/19, 4/20, 4/21, 4/22 4/23, 4/24
Protocol_VLAN_Name:	
Router Interface:	None
Protocol Type:	<input type="radio"/> IP <input type="radio"/> IPX <input type="radio"/> AppleTalk <input type="radio"/> Decnet <input type="radio"/> NetBIOS <input checked="" type="radio"/> Others
Selected Port Members:	<input type="checkbox"/> Dynamic Port Static Port: Change Static Members Exclude Port: Change Exclude Members

Clear Add Modify Delete Reset

[Show][Protocol][IP Subnet][IPX Network][AppleTalk Cable]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

- Enter the VLAN ID that will contain the protocol VLAN in the VLAN ID field.

6. Enter a name for the VLAN in the Protocol_VLAN_Name field.
7. Select the virtual routing interface from the Router_Interface pulldown list if you configured a virtual routing interface for routing into and out of the VLAN.
8. Select the protocol type.
9. Specify the port that are members for the VLAN:
 - Select Dynamic Port if you want the port membership to be dynamic. For information, see “Dynamic Ports” on page 22-12.
 - Click the Change Static Members button if you want to configure static ports. For information, see “Static Ports” on page 22-13.
 - Click the Change Exclude Members button if you want to explicitly exclude some ports. For information, see “Excluded Ports” on page 22-13.

NOTE: All the ports must be members of the port-based VLAN that contains this IP subnet VLAN. See “Layer 3 Protocol-Based VLANs” on page 22-3.

10. Click the Add button (if you are adding a new VLAN) or the Modify button (if you are modifying an existing VLAN) to save the change to the device’s running-config file.
11. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Configuring an IP Subnet VLAN

1. Log on to the device using a valid user name and password for read-write access.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to VLAN in the tree view to expand the list of VLAN option links.
4. Click on the Protocol link.
 - If the device does not have any protocol VLANs, the Protocol VLAN configuration panel is displayed, as shown in the following example.
 - If at least one protocol VLAN is already configured and you are adding a new one, click on the IP Subnet link to display the IP Subnet Protocol VLAN configuration panel.
 - If you are modifying an existing protocol VLAN, click on the Modify button to the right of the row describing the VLAN to display the configuration panel for the type of VLAN you are modifying. The following example shows the IP Subnet Protocol VLAN configuration dialog, used for configuring an IP

subnet protocol VLAN (not a protocol, IPX network, or AppleTalk cable VLAN).

VLAN Id:	<input type="text" value="2"/>
VLAN Port_members:	1/7,1/8, 3/1,3/2,3/3,3/4,3/5,3/6,3/7,3/8, 3/9,3/10,3/11,3/12,3/13,3/14 3/15,3/16, 3/17,3/18,3/19,3/20,3/21,3/22,3/23,3/24, 4/1,4/2,4/3,4/4,4/5,4/6 4/7,4/8, 4/9,4/10,4/11,4/12,4/13,4/14,4/15,4/16, 4/17,4/18,4/19,4/20,4/21,4/22 4/23,4/24
Protocol_VLAN_Name:	<input type="text"/>
Router_Interface:	None <input type="button" value="v"/>
IP_Address:	<input type="text" value="209.157.22.4"/>
Mask:	<input type="text" value="255.255.255.0"/>
Selected Port Members:	<input type="checkbox"/> Dynamic Port Static Port: <input type="button" value="Change Static Members"/> Exclude Port: <input type="button" value="Change Exclude Members"/>

[\[Show\]](#)
[\[Protocol\]](#)
[\[IP Subnet\]](#)
[\[IPX Network\]](#)
[\[AppleTalk Cable\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

5. Enter the VLAN ID that will contain the IP subnet VLAN in the VLAN ID field.
6. Enter a name for the VLAN in the Protocol_VLAN_Name field.
7. Select the virtual routing interface from the Router_Interface pulldown list if you configured a virtual routing interface for routing into and out of the VLAN.
8. Enter the IP address of the VLAN in the IP_Address field.
9. Enter the network mask in the Mask field.
10. Specify the port that are members for the VLAN:
 - Select Dynamic Port if you want the port membership to be dynamic. For information, see “Dynamic Ports” on page 22-12.
 - Click the Change Static Members button if you want to configure static ports. For information, see “Static Ports” on page 22-13.
 - Click the Change Exclude Members button if you want to explicitly exclude some ports. For information, see “Excluded Ports” on page 22-13.

NOTE: All the ports must be members of the port-based VLAN that contains this IP subnet VLAN. See “Layer 3 Protocol-Based VLANs” on page 22-3.

11. Click the Add button (if you are adding a new VLAN) or the Modify button (if you are modifying an existing VLAN) to save the change to the device’s running-config file.
12. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Configuring an IPX Network VLAN

1. Log on to the device using a valid user name and password for read-write access.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to VLAN in the tree view to expand the list of VLAN option links.

4. Click on the [Protocol](#) link.
 - If the device does not have any protocol VLANs, the Protocol VLAN configuration panel is displayed, as shown in the following example.
 - If at least one protocol VLAN is already configured and you are adding a new one, click on the [IPX Network](#) link to display the IP Subnet Protocol VLAN configuration panel.
 - If you are modifying an existing protocol VLAN, click on the Modify button to the right of the row describing the VLAN to display the configuration panel for the type of VLAN you are modifying. The following example shows the IPX Network Protocol VLAN configuration dialog, used for configuring an IPX network protocol VLAN (not a protocol, IP subnet, or AppleTalk cable VLAN).

VLAN Id:	<input type="text" value="1"/>
VLAN Port_members:	1/7, 1/8, 3/1, 3/2, 3/3, 3/4, 3/5, 3/6, 3/7, 3/8, 3/9, 3/10, 3/11, 3/12, 3/13, 3/14 3/15, 3/16, 3/17, 3/18, 3/19, 3/20, 3/21, 3/22, 3/23, 3/24, 4/1, 4/2, 4/3, 4/4, 4/5, 4/6 4/7, 4/8, 4/9, 4/10, 4/11, 4/12, 4/13, 4/14, 4/15, 4/16, 4/17, 4/18, 4/19, 4/20, 4/21, 4/22 4/23, 4/24
Protocol_VLAN_Name:	<input type="text"/>
Router_Interface:	None <input type="button" value="v"/>
Frame_Type:	Ethernet_802.2 <input type="button" value="v"/>
Network:	<input type="text" value="00000200"/>
Selected Port Members:	<input type="checkbox"/> Dynamic Port Static Port: <input type="button" value="Change Static Members"/> Exclude Port: <input type="button" value="Change Exclude Members"/>

[\[Show\]](#)
[\[Protocol\]](#)
[\[IP Subnet\]](#)
[\[IPX Network\]](#)
[\[AppleTalk Cable\]](#)

[\[Home\]](#)
[\[Site Map\]](#)
[\[Logout\]](#)
[\[Save\]](#)
[\[Frame Enable\]](#)
[\[Disable\]](#)
[\[TELNET\]](#)

5. Enter the VLAN ID that will contain the IPX network VLAN in the VLAN ID field.
6. Enter a name for the VLAN in the Protocol_VLAN_Name field.
7. Select the virtual routing interface from the Router_Interface pulldown list if you configured a virtual routing interface for routing into and out of the VLAN.
8. Select the encapsulation type from the Frame_Type field's pulldown list.
9. Enter the IPX network address of the VLAN in the Network field.
10. Specify the port that are members for the VLAN:
 - Select Dynamic Port if you want the port membership to be dynamic. For information, see “Dynamic Ports” on page 22-12.
 - Click the Change Static Members button if you want to configure static ports. For information, see “Static Ports” on page 22-13.
 - Click the Change Exclude Members button if you want to explicitly exclude some ports. For information, see “Excluded Ports” on page 22-13.

NOTE: All the ports must be members of the port-based VLAN that contains this IPX network VLAN. See “Layer 3 Protocol-Based VLANs” on page 22-3.

11. Click the Add button (if you are adding a new VLAN) or the Modify button (if you are modifying an existing VLAN) to save the change to the device's running-config file.
12. Select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Configuring an AppleTalk Cable VLAN

1. Log on to the device using a valid user name and password for read-write access.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to VLAN in the tree view to expand the list of VLAN option links.
4. Click on the [Protocol](#) link.
 - If the device does not have any protocol VLANs, the Protocol VLAN configuration panel is displayed, as shown in the following example.
 - If at least one protocol VLAN is already configured and you are adding a new one, click on the [AppleTalk Cable](#) link to display the AppleTalk Cable VLAN configuration panel.
 - If you are modifying an existing protocol VLAN, click on the Modify button to the right of the row describing the VLAN to display the configuration panel for the type of VLAN you are modifying. The following example shows the AppleTalk Cable VLAN configuration dialog, used for configuring an AppleTalk cable VLAN (not a protocol, IP subnet, or IPX network VLAN).

VLAN Id:	<input type="text" value="1"/>
VLAN Port_members:	1/7,1/8, 3/1,3/2,3/3,3/4, 3/5,3/6,3/7,3/8, 3/9,3/10,3/11,3/12,3/13,3/14 3/15,3/16, 3/17,3/18,3/19,3/20,3/21,3/22,3/23,3/24, 4/1,4/2,4/3,4/4,4/5,4/6 4/7,4/8, 4/9,4/10,4/11,4/12,4/13,4/14,4/15,4/16, 4/17,4/18,4/19,4/20,4/21,4/22 4/23,4/24
Protocol_VLAN_Name:	<input type="text"/>
Router_Interface:	None <input type="button" value="v"/>
AppleTalk Cable:	1 <input type="button" value="v"/>
Selected Port Members:	<input type="checkbox"/> Dynamic Port Static Port: <input type="button" value="Change Static Members"/> Exclude Port: <input type="button" value="Change Exclude Members"/>

[\[Show\]](#)[\[Protocol\]](#)[\[IP Subnet\]](#)[\[IPX Network\]](#)[\[AppleTalk Cable\]](#)

[\[Home\]](#)[\[Site Map\]](#)[\[Logout\]](#)[\[Save\]](#)[\[Frame Enable\]](#)[\[Disable\]](#)[\[TELNET\]](#)

5. Enter the VLAN ID that will contain the AppleTalk cable VLAN in the VLAN ID field.
6. Enter a name for the VLAN in the Protocol_VLAN_Name field.
7. Select the virtual routing interface from the Router_Interface pulldown list if you configured a virtual routing interface for routing into and out of the VLAN.
8. Select the AppleTalk cable ID from the AppleTalk Cable field's pulldown list.
9. Specify the port that are members for the VLAN:
 - Select Dynamic Port if you want the port membership to be dynamic. For information, see "Dynamic Ports" on page 22-12.
 - Click the Change Static Members button if you want to configure static ports. For information, see "Static

Ports” on page 22-13.

- Click the Change Exclude Members button if you want to explicitly exclude some ports. For information, see “Excluded Ports” on page 22-13.

NOTE: All the ports must be members of the port-based VLAN that contains this AppleTalk cable VLAN. See “Layer 3 Protocol-Based VLANs” on page 22-3.

10. Click the Add button (if you are adding a new VLAN) or the Modify button (if you are modifying an existing VLAN) to save the change to the device’s running-config file.
11. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device’s flash memory.

Displaying VLAN Information

After you configure the VLANs, you can verify the configuration using the following methods.

NOTE: If a VLAN name begins with “GVRP_VLAN_“, the VLAN was created by the GARP VLAN Registration Protocol (GVRP). If a VLAN name begins with “STATIC_VLAN_“, the VLAN was created by GVRP and then was converted into a statically configured VLAN.

Displaying System-Wide VLAN Information

Use one of the following methods to display VLAN information for all the VLANs configured on the device.

USING THE CLI

Enter the following command at any CLI level. This example shows the display for the IP subnet and IPX network VLANs configured in the examples in “Configuring an IP Subnet VLAN with Dynamic Ports” on page 22-39 and “Configuring an IPX Network VLAN with Dynamic Ports” on page 22-40.

```
BigIron(config)# show vlans

Total PORT-VLAN entries: 2
Maximum PORT-VLAN entries: 8
legend: [S=Slot]

PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S2) 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
  Untagged Ports: (S2) 17 18 19 20 21 22 23 24
  Untagged Ports: (S4) 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
  Untagged Ports: (S4) 17 18 19 20 21 22 23 24
  Tagged Ports: None

PORT-VLAN 10, Name IP_VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S1) 1 2 3 4 5 6
  Tagged Ports: None

IP-subnet VLAN 1.1.1.0 255.255.255.0, Dynamic port enabled
  Name: Mktg-LAN
  Static ports: None
  Exclude ports: None
  Dynamic ports: (S1) 1 2 3 4 5 6
PORT-VLAN 20, Name IPX_VLAN, Priority level0, Spanning tree Off
  Untagged Ports: (S2) 1 2 3 4 5 6
  Tagged Ports: None

IPX-network VLAN 0000ABCD, frame type ethernet_ii, Dynamic port enabled
  Name: Eng-LAN
  Static ports: None
  Exclude ports: None
  Dynamic ports: (S2) 1 2 3 4 5 6
```

Syntax: show vlans [*<vlan-id>* | ethernet *<portnum>* | pos *<portnum>*]

The *<vlan-id>* parameter specifies a VLAN for which you want to display the configuration information.

The **ethernet** *<portnum>* | **pos** *<portnum>* parameter specifies a port. If you use this parameter, the command lists all the VLAN memberships for the port.

USING THE WEB MANAGEMENT INTERFACE

To display VLAN configuration information:

1. Log on to the device using a valid user name and password for read-only or read-write access. The System configuration dialog is displayed.
2. Click on the plus sign next to Configure in the tree view to expand the list of configuration options.
3. Click on the plus sign next to VLAN in the tree view to expand the list of VLAN option links.
4. Click on the [Port](#) link to display the Port-based VLAN table or the [Protocol](#) link to display the Protocol-based VLAN table.

Displaying VLAN Information for Specific Ports

Use one of the following methods to display VLAN information for specific ports.

USING THE CLI

To display VLAN information for all the VLANs of which port 7/1 is a member, enter the following command:

```
BigIron(config)# show vlans e 7/1

Total PORT-VLAN entries: 3
Maximum PORT-VLAN entries: 8

legend: [S=Slot]

PORT-VLAN 100, Name [None], Priority level0, Spanning tree Off
  Untagged Ports: (S7) 1 2 3 4
  Tagged Ports: None

IP-subnet VLAN 207.95.11.0 255.255.255.0, Dynamic port disabled
  Static ports: (S7) 1 2
  Exclude ports: None
  Dynamic ports: None
```

Syntax: show vlans [<vlan-id> | ethernet <portnum> | pos <portnum>]

The <vlan-id> parameter specifies a VLAN for which you want to display the configuration information.

The **ethernet** <portnum> | **pos** <portnum> parameter specifies a port. If you use this parameter, the command lists all the VLAN memberships for the port.

USING THE WEB MANAGEMENT INTERFACE

You cannot display port-specific VLAN information using the Web management interface.

Using show run to Display VLAN Information

.You can use the **show running-config** command to display information for VLANs and VLAN groups. The **show running-config** command displays VLANs in the order they were configured on the device. Device running Enterprise software release 08.0.00 and later displays VLANs in numerical order (for example, VLAN 100, VLAN 200, and so on in the **show running-config** output).

Chapter 24

Configuring IP Multicast Traffic Reduction

Foundry Layer 2 Switches and Layer 3 Switches forward all IP multicast traffic by default based on the Layer 2 information in the packets. Optionally, you can enable these Foundry devices to make forwarding decisions in hardware, based on multicast group by enabling the IP Multicast Traffic Reduction feature.

When this feature is enabled, these Foundry devices examine the MAC address in an IP multicast packet and forward the packet only on the ports from which the device has received Group Membership reports for that group, instead of forwarding all multicast traffic to all ports. The device sends traffic for other groups out all ports.

When you enable IP Multicast Traffic Reduction, you also can configure the following features:

- IGMP mode – When you enable IP Multicast Traffic Reduction, the device passively listens for IGMP Group Membership reports by default. If the multicast domain does not have a router to send IGMP queries to elicit these Group Membership reports, you can enable the device to actively send the IGMP queries.
- Query interval – The query interval specifies how often the device sends Group Membership queries. This query interval applies only to the active IGMP mode. The default is 60 seconds. You can change the interval to a value from 10 – 600 seconds.
- Age interval – The age interval specifies how long an IGMP group can remain in the IGMP group table without the device receiving a Group Membership report for the group. If the age interval expires before the device receives another Group Membership report for the group, the device removes the entry from the table. The default is 140 seconds. You can change the interval to a value from 10 – 1220 seconds.
- Forwarding policy – The device forwards all IP multicast traffic by default but you can enable the device to forward IP multicast traffic only for groups for which the device has received a Group Membership report, and drop traffic for all other groups.

The following sections describe how to configure IP multicast traffic reduction and PIM SM Traffic Snooping parameters on a Foundry device.

NOTE: Beginning with software release 07.7.00, IP multicast traffic reduction and PIM SM Traffic Snooping are available on Layer 3 Switches.

Enabling IP Multicast Traffic Reduction

By default, Foundry devices forward all IP multicast traffic out all ports except the port on which the traffic was received. To reduce multicast traffic through the device, you can enable IP Multicast Traffic Reduction. This feature configures the device to forward multicast traffic only on the ports attached to multicast group members, instead of forwarding all multicast traffic to all ports. The device determines the ports that are attached to multicast group members based on entries in the IGMP table. Each entry in the table consists of MAC addresses and the Foundry device ports from which the device has received Group Membership reports for that group.

By default, the device broadcasts traffic addressed to an IP multicast group that doesn't have any entries in the IGMP table. When you enable IP Multicast Traffic Reduction, the device determines the ports that are attached to multicast group members based on entries in the IGMP table. The IGMP table entries are created when the VLAN receives a group membership report for a group. Each entry in the table consists of an IP multicast group address and the Foundry device ports from which the device has received Group Membership reports.

When the device receives traffic for an IP multicast group, the device looks in the IGMP table for an entry corresponding to that group. If the device finds an entry, the device forwards the group traffic out the ports listed in the corresponding entries, as long as the ports are members of the same VLAN. If the table does not contain an entry corresponding to the group or if the port is a member of the default VLAN, the device broadcasts the traffic.

NOTE: When one or more Foundry devices are running Layer 2 IP Multicast Traffic reduction, configure one of the devices for active IGMP and leave the other devices configured for passive IGMP. However, if the IP multicast domain contains a multicast-capable router, configure all the Foundry devices for passive IGMP and allow the router to actively send the IGMP queries.

To enable IP Multicast Traffic Reduction, use either of the following methods.

USING THE CLI

To enable IP Multicast Traffic Reduction, enter the following command:

```
BigIron(config)# ip multicast
```

Syntax: [no] ip multicast

NOTE: If the "route-only" feature is enabled on the Layer 3 Switch, then IP Multicast Traffic Reduction will not be supported.

NOTE: This feature is not supported on the default VLAN of Layer 3 Switches.

To verify that IP Multicast Traffic Reduction is enabled, enter the following command at any level of the CLI:

```
BigIron(config)# show ip multicast
IP multicast is enabled - Active
```

Syntax: show ip multicast

NOTE: In software releases earlier than 07.1.09, this command does not display a message if you have enabled IP Multicast Traffic Reduction but you have not yet reloaded the software.

USING THE WEB MANAGEMENT INTERFACE

To enable IP Multicast Traffic Reduction on a device:

NOTE: This feature is not available on Layer 3 Switches.

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select Enable next to IP Multicast.
3. Click the Apply button to save the change to the device's running-config file.
4. Select the Save link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Changing the IGMP Mode

When you enable IP Multicast Traffic Reduction on the device, IGMP also is enabled. The device uses IGMP to maintain a table of the Group Membership reports received by the device. You can use active or passive IGMP mode. The default mode is passive.

- Active – When active IGMP mode is enabled, a Foundry device actively sends out IGMP queries to identify IP multicast groups on the network and makes entries in the IGMP table based on the Group Membership reports received from the network.

NOTE: Routers in the network generally handle this operation. Use the active IGMP mode only when the device is in a stand-alone Layer 2 Switched network with no external IP multicast router attachments. In this case, enable the active IGMP mode on only one of the devices and leave the other devices configured for passive IGMP mode.

- Passive – When passive IGMP mode is enabled, the device listens for IGMP Group Membership reports but does not send IGMP queries. The passive mode is sometimes called “IGMP snooping”. Use this mode when another device in the network is actively sending queries.

To set change the IGMP mode, use either of the following methods.

NOTE: Layer 2 Switches only: In software releases earlier than 07.1.10, you must reload the software after making this configuration change and saving it to the startup-config file. If you are using software release 07.1.10 or later, you do not need to reload the software.

USING THE CLI

To enable active IGMP, enter the following command:

```
BigIron(config)# ip multicast active
BigIron(config)# write memory
BigIron(config)# end
BigIron# reload
```

Syntax: [no] ip multicast active | passive

To enable passive IGMP, enter the following command:

```
BigIron(config)# ip multicast passive
BigIron(config)# write memory
BigIron(config)# end
BigIron# reload
```

USING THE WEB MANAGEMENT INTERFACE

To change the IGMP mode:

NOTE: This feature is not available on Layer 3 Switches.

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select Active or Passive next to IGMP.
3. Click the Apply button to save the change to the device's running-config file.
4. Select the [Save](#) link at the bottom of the dialog, then select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Disabling IGMP on Individual Ports

NOTE: IP Multicast Traffic Reduction cannot be disabled on individual ports of a Layer 3 Switch. You cannot use the **ip-multicast-disable** command that is available on Layer 2 Switches. IP multicast must be enabled and disabled globally on Layer 3 Switches.

By default, when you enable IP multicast on a Foundry device, all ports on the device are configured for IGMP. If you are using active IGMP, all ports can send IGMP queries and receive IGMP reports. If you are using passive IGMP, all ports can receive IGMP queries.

You can disable IGMP on individual ports of a Layer 2 Switch if you want to block all IP multicast traffic on those ports. When you disable IGMP on an individual port, the device does not forward any multicast traffic out the port, but other ports can still send and receive multicast traffic.

To disable IGMP on a port, use the following CLI method.

NOTE: Layer 2 Switches only: In software releases earlier than 07.1.10, you must reload the software after making this configuration change and saving it to the startup-config file. If you are using software release 07.1.10 or later, you do not need to reload the software.

USING THE CLI

```
BigIron(config)# int e 1/5
BigIron(config-if-1/5)# ip-multicast-disable
```

Syntax: [no] ip-multicast-disable

The command in this example disables IGMP on port 1/5 but does not affect the state of IGMP on other ports.

USING THE WEB MANAGEMENT INTERFACE

You cannot disable IGMP on a port using the Web management interface.

Modifying the Query Interval

If IP Multicast Traffic Reduction is set to active mode, you can modify the query interval, which specifies how often a Foundry device enabled for active IP Multicast Traffic Reduction sends Group Membership queries.

NOTE: The query interval applies only to the active mode of IP Multicast Traffic reduction.

To modify the query interval, use the following CLI method.

NOTE: In software releases earlier than 07.1.10, you must reload the software after making this configuration change and saving it to the startup-config file. If you are using software release 07.1.10 or later, you do not need to reload the software.

USING THE CLI

To modify the query interval, enter a command such as the following:

```
BigIron(config)# ip multicast query-interval 120
```

Syntax: [no] ip multicast query-interval <interval>

The <interval> parameter specifies the interval between queries. You can specify a value from 10 – 600 seconds. The default is 60 seconds.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure this feature using the Web management interface.

Modifying the Age Interval

When the device receives a Group Membership report, the device makes an entry in the IGMP group table for the group in the report. The age interval specifies how long the entry can remain in the table without the device receiving another Group Membership report.

To modify the age interval, use the following CLI method.

NOTE: In software releases earlier than 07.1.10, you must reload the software after making this configuration change and saving it to the startup-config file. If you are using software release 07.1.10 or later, you do not need to reload the software.

USING THE CLI

To modify the age interval, enter a command such as the following:

```
BigIron(config)# ip multicast age-interval 280
```

Syntax: [no] ip multicast age-interval <interval>

The <interval> parameter specifies the interval between queries. You can specify a value from 10 – 1220 seconds. The default is 140 seconds.

USING THE WEB MANAGEMENT INTERFACE

You cannot configure this feature using the Web management interface.

Filtering Multicast Groups

By default, Foundry devices forward multicast traffic for all valid multicast groups. You can configure a Foundry device to filter out all multicast traffic for groups other than the ones for which the device has received Group Membership reports.

When device starts up, it forwards all multicast groups even though multicast traffic filters are configured. This process continues until the device receives a group membership report. Once the group membership report is received, the device drops all multicast packets for groups other than the ones for which the device has received the group membership report.

To enable IP multicast filtering, use the following CLI method.

NOTE: In software releases earlier than 07.1.10, you must reload the software after making this configuration change and saving it to the startup-config file. If you are using software release 07.1.10 or later, you do not need to reload the software.

USING THE CLI

To enable IP multicast filtering, enter the following command:

```
BigIron(config)# ip multicast filter
```

Syntax: [no] ip multicast filter

USING THE WEB MANAGEMENT INTERFACE

You cannot configure this feature using the Web management interface.

PIM SM Traffic Snooping

By default, when a Foundry device receives an IP multicast packet, the device does not examine the multicast information in the packet. Instead, the device simply forwards the packet out all ports except the port that received the packet. In some networks, this method can cause unnecessary traffic overhead in the network. For example, if the Foundry device is attached to only one group source and two group receivers, but has devices attached to every port, the device forwards group traffic out all ports in the same broadcast domain except the port attached to the source, even though there are only two receivers for the group.

PIM SM traffic snooping eliminates the superfluous traffic by configuring the device to forward IP multicast group traffic only on the ports that are attached to receivers for the group.

PIM SM traffic snooping requires IP multicast traffic reduction to be enabled on the device. IP multicast traffic reduction configures the device to listen for IGMP messages. PIM SM traffic snooping provides a finer level of multicast traffic control by configuring the device to listen specifically for PIM SM join and prune messages sent from one PIM SM router to another through the device.

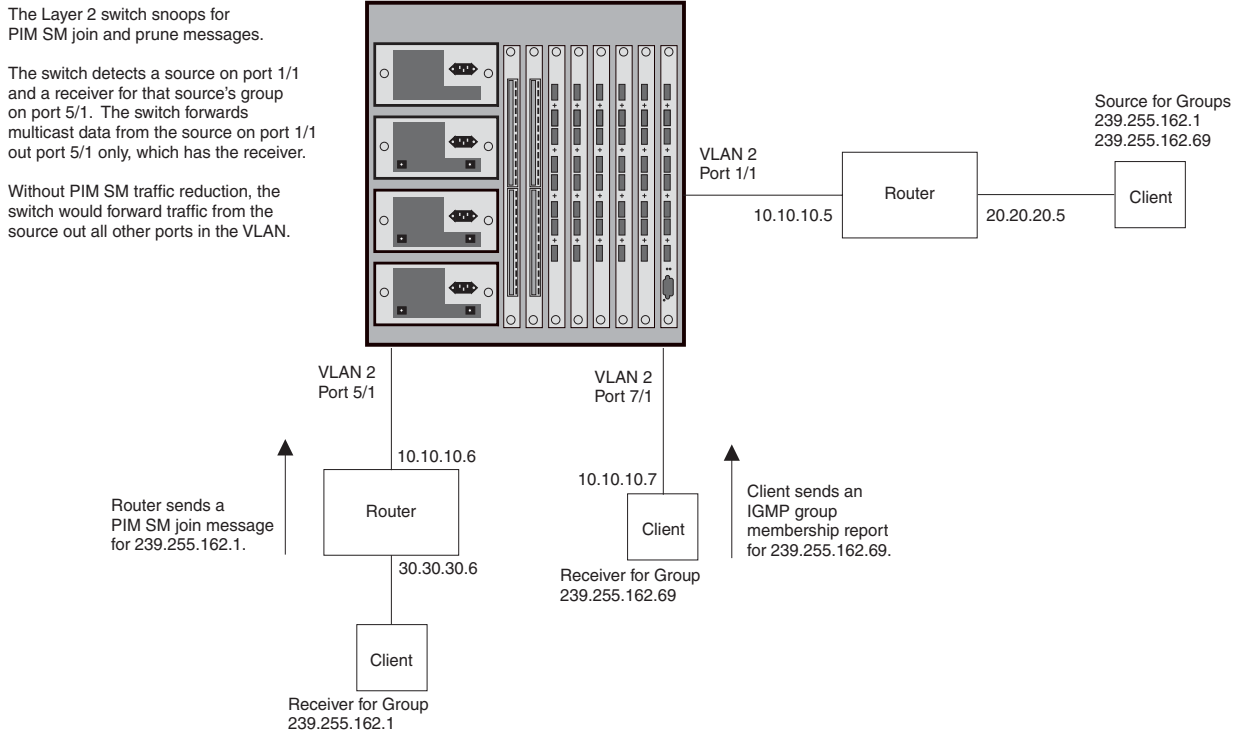
NOTE: This feature applies only to PIM SM version 2 (PIM V2).

NOTE: Software release 03.4.00 introduces PIM SM traffic snooping on the Layer 2 switch (FES0xxxx.bin software image).

Application Examples

Figure 24.1 shows an example application of the PIM SM traffic snooping feature. In this example, a device is connected through an IP router to a PIM SM group source that is sending traffic for two PIM SM groups. The device also is connected to a receiver for each of the groups.

Figure 24.1 PIM SM traffic reduction in enterprise network



When PIM SM traffic snooping is enabled, the device starts listening for PIM SM join and prune messages and IGMP group membership reports. Until the device receives a PIM SM join message or an IGMP group membership report, the device forwards IP multicast traffic out all ports. Once the device receives a join message or group membership report for a group, the device forwards subsequent traffic for that group only on the ports from which the join messages or IGMP reports were received.

In this example, the router connected to the receiver for group 239.255.162.1 sends a join message toward the group's source. Since PIM SM traffic snooping is enabled on the device, the device examines the join message to learn the group ID, then makes a forwarding entry for the group ID and the port connected to the receiver's router. The next time the device receives traffic for 239.255.162.1 from the group's source, the device forwards the traffic only on port 5/1, since that is the only port connected to a receiver for the group.

Notice that the receiver for group 239.255.162.69 is directly connected to the device. As result, the device does not see a join message on behalf of the client. However, since IP multicast traffic reduction also is enabled, the device uses the IGMP group membership report from the client to select the port for forwarding traffic to group 239.255.162.69 receivers.

The IP multicast traffic reduction feature and the PIM SM traffic snooping feature together build a list of groups and forwarding ports for the VLAN. The list includes PIM SM groups learned through join messages as well as MAC addresses learned through IGMP group membership reports. In this case, even though the device never sees a join message for the receiver for group 239.255.162.69, the device nonetheless learns about the receiver and forwards group traffic to the receiver.

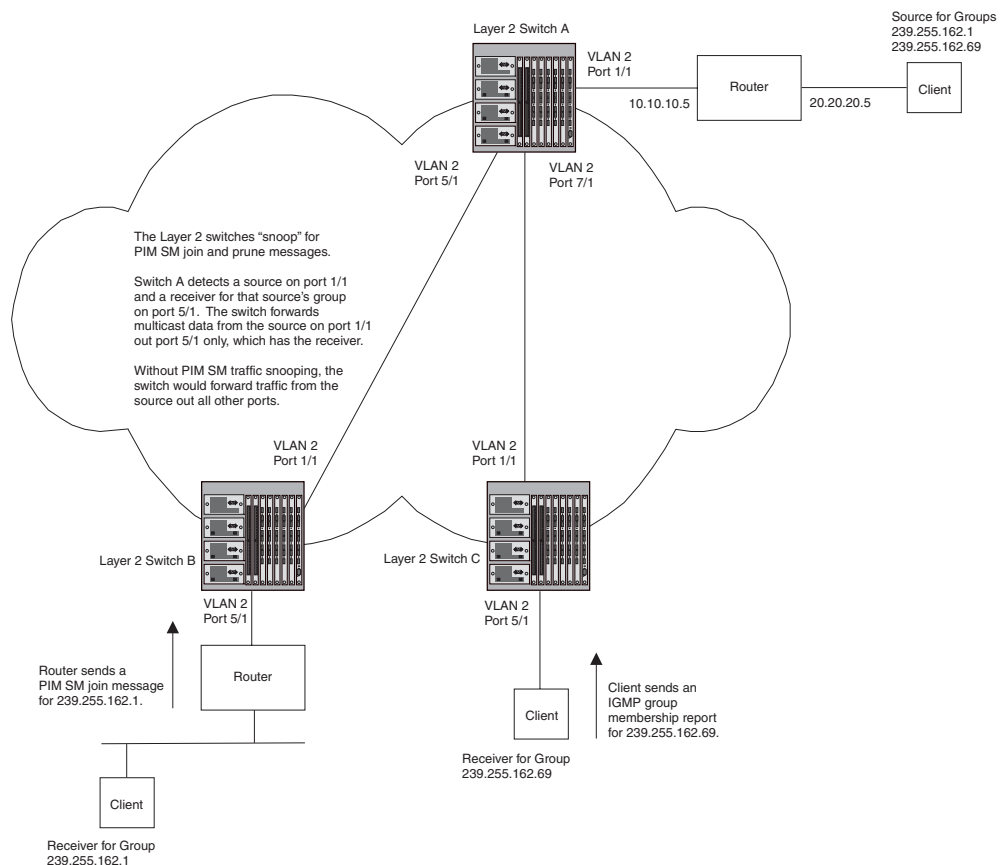
The device stops forwarding IP multicast traffic on a port for a group if the port receives a prune message for the group.

Notice that the ports connected to the source and the receivers are all in the same port-based VLAN on the device. This is required for the PIM SM snooping feature. The feature also requires the source and the downstream router to be on different IP subnets, as shown in Figure 24.1.

Figure 24.2 shows another example application for PIM SM traffic snooping. This example shows devices on the edge of a Global Ethernet cloud (a Layer 2 Packet over SONET cloud). Assume that each device is attached to numerous other devices such as other Layer 2 Switches and Layer 3 Switches (routers).

NOTE: This example assumes that the devices are actually BigIron Chassis devices running Layer 2 Switch software. Global Ethernet requires POS interfaces, which are supported on the BigIron and NetIron Chassis devices but not the FastIron Chassis devices.

Figure 24.2 PIM SM traffic reduction in Global Ethernet environment



The devices on the edge of the Global Ethernet cloud are configured for IP multicast traffic reduction and PIM SM traffic snooping. Although this application uses multiple devices, the feature has the same requirements and works the same way as it does on a single device.

Configuration Requirements

- IP multicast traffic reduction must be enabled on the device that will be running PIM SM snooping. The PIM SM traffic snooping feature requires IP multicast traffic reduction.

NOTE: Use the passive mode of IP multicast traffic reduction instead of the active mode. The passive mode assumes that a router is sending group membership queries as well as join and prune messages on behalf of receivers. The active mode configures the device to send group membership queries.

- All the device ports connected to the source and receivers or routers must be in the same port-based VLAN.
- The PIM SM snooping feature assumes that the group source and the device are in different subnets and communicate through a router. The source must be in a different IP subnet than the receivers. A PIM SM router sends PIM join and prune messages on behalf of a multicast group receiver only when the router and the source are in different subnets. When the receiver and source are in the same subnet, they do not need the router in order to find one another. They find one another directly within the subnet.

The device forwards all IP multicast traffic by default. Once you enable IP multicast traffic reduction and PIM SM traffic snooping, the device initially blocks all PIM SM traffic instead of forwarding it. The device forwards PIM SM traffic to a receiver only when the device receives a join message from the receiver. Consequently, if the source and the downstream router are in the same subnet, and PIM SM traffic snooping is enabled, the device blocks the PIM SM traffic and never starts forwarding the traffic. This is because the device never receives a join message from the downstream router for the group. The downstream router and group find each other without a join message because they are in the same subnet.

NOTE: If the “route-only” feature is enabled on a Layer 3 Switch, PINM SM traffic snooping will not be supported.

Enabling PIM SM Traffic Snooping

To enable PIM SM traffic snooping, you must enable IP multicast traffic reduction, then enable snooping. Use the following CLI method.

USING THE CLI

To enable PIM SM traffic snooping, enter the following commands at the global CONFIG level of the CLI:

```
BigIron(config)# ip multicast
BigIron(config)# ip pimsm-snooping
```

The first command enables IP multicast traffic reduction. This feature is similar to PIM SM traffic snooping but listens only for IGMP information, not PIM SM information. You must enable both IP multicast traffic reduction and PIM SM traffic snooping to enable the device to listen for PIM SM join and prune messages.

Syntax: [no] ip multicast [active | passive]

This command enables IP multicast traffic reduction. The **active | passive** parameter specifies the mode. The PIM SM traffic snooping feature assumes that the network has routers that are running PIM SM.

Syntax: [no] ip pimsm-snooping

This command enables PIM SM traffic snooping.

To disable the feature, enter the following command:

```
BigIron(config)# no ip pimsm-snooping
```

If you also want to disable IP multicast traffic reduction, enter the following command:

```
BigIron(config)# no ip multicast
```

Displaying IP Multicast Information

The following sections show how to display and clear IP multicast reduction information.

Displaying Multicast Information on Layer 2 Switches

To display IP multicast information on Layer 2 Switches, including the state of the traffic reduction and traffic snooping features, use the following CLI methods.

USING THE CLI

To display IP multicast information, enter the following command at any level of the CLI:

```
BigIron(config)# show ip multicast
IP multicast is enabled - Passive
VLAN ID 22
Active 5.5.5.1 Router Ports 3/4 3/10 5/3
Total number of Multicast Group: 1
  1 Multicast Group: 239.255.162.1, Port: 3/4 3/10 5/3
  IGMP Group Port:
  PIMv2 Group Port: 3/4 3/10 5/3
```

Syntax: show ip multicast

This display shows the following information.

This Field...	Displays...
The IP multicast traffic snooping state	The first line of the display indicates whether IP multicast traffic snooping is enabled or disabled. The PIM SM traffic snooping feature requires the IP multicast traffic reduction feature.
VLAN ID	The port-based VLAN to which the information listed below the VLAN ID applies. Each port-based VLAN is a separate Layer 2 broadcast domain. Note: PIM SM traffic snooping requires the source and the receivers to be in the same port-based VLAN on the device. If the source and receivers are in different port-based VLANs, the device blocks the multicast traffic.
Active	The IP address of the device that actively sends IGMP queries.
Router Ports	The ports that are connected to routers that support IP multicast.
Total Number of Multicast Group	The number of groups for which the VLAN's ports have received IGMP group membership reports, join messages, or prune messages.
Multicast Group	An IP multicast group.
Port	The ports attached to receivers for the IP multicast group.
IGMP Group Port	The port(s) in this VLAN on which the device has received IGMP group membership reports.
PIMv2 Group Port	The port(s) in this VLAN on which the device has received join messages for IP multicast.

You also can display PIM SM information on Layer 2 Switches by entering the following command, at any level of the CLI:

```
BigIron(config)# show ip pim
PIMSM snooping is enabled
VLAN ID 22
PIMSM Neighbor list:
    5.5.5.2 : 3/4   expire 95 s
    5.5.5.3 : 3/10  expire 180 s
    5.5.5.1 : 5/3   expire 160 s
Multicast Group: 239.255.162.1, fid 000026bc camindex 2058
Forwarding Port: 3/4 3/10 5/3
PIMv2 Group Port: 3/4 3/10 5/3
(Source, Port) list:
    55.55.55.2, port: 3/10 5/3
    42.42.42.42, port: 3/4 3/10
    5.5.5.1, port: 3/10
    162.162.162.162, port: 3/4 5/3
```

Syntax: show ip pim

This display shows the following information.

This Field...	Displays...
The PIM SM traffic snooping state	The first line of the display indicates whether the feature is enabled or disabled.
VLAN ID	The port-based VLAN to which the neighbors and groups listed below the VLAN ID apply. Each port-based VLAN is a separate Layer 2 broadcast domain. Note: PIM SM traffic snooping requires the source and the receivers to be in the same port-based VLAN on the device. If the source and receivers are in different port-based VLANs, the device blocks the multicast traffic.
PIM SM Neighbor list	The PIM SM routers that are attached to the device's ports in the VLAN. The value following "expire" indicates how many seconds the device will wait for a hello message from the neighbor before determining that the neighbor is no longer present and removing the neighbor from the list.
Multicast Group	The IP multicast group ID. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Forwarding Port	The port(s) attached to the group's receivers. A port is listed here when it receives a join message for the group, an IGMP membership report for the group, or both.
PIMv2 Group Port	The port(s) on which the device has received PIM SM join messages for the group.
Source, Port list	The IP address of each PIM SM source and the device ports connected to the receivers of the source.

Beginning with software release 07.8.00, the show ip multicast output has been changed to display the following information:

```
BigIron# show ip multicast
IP multicast is enabled - Passive
VLAN ID 100
  Querier: 1.100.100.7, (port: 3/1)
  Router Ports: 3/1 3/2 3/3
  Total number of Multicast Group in vlan: 3
1   Group: 224.0.1.22, fid 08ac, NO cam
   Forwarding Port: 3/3
2   Group: 239.255.162.2, fid 08aa, cam 8
   Forwarding Port: 3/1 3/2
3   Group: 239.255.163.2, fid 08a9, cam 10
   Forwarding Port: 3/1 3/2

VLAN ID 4008
  Querier: 1.1.5.1, (port: 3/48)
  Router Ports: 3/48
  Total number of Multicast Group in vlan: 0
```

Syntax: show ip multicast

This display shows the following information.

This Field...	Displays...
The IP multicast traffic snooping state	The first line of the display indicates whether IP multicast traffic snooping is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.
VLAN ID	The port-based VLAN to which the information listed applies.
Querier	The IP address of the device that actively sends IGMP queries.
(port)	The port on which the queries are being sent out.
Router Ports	The ports that are connected to routers that support IP multicast.
Total Number of Multicast Group in VLAN	The total number of groups for which the VLAN's ports have received IGMP group membership reports, join messages, or prune messages.
Multicast Group	Address of the IP multicast group. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Forwarding Port	The forwarding ports for the IP multicast group.

Displaying Multicast Information for a Specific Group

Beginning with software release 07.8.00, you can display multicast information for a specific group, by entering a command such as the following at any level of the CLI:

```
BigIron# show ip multicast 239.255.162.2
VLAN ID 100
Active 1.100.100.7 Router Ports 3/1 3/2 3/3
  Group: 239.255.162.2, fid 08aa, cam 8
    Forwarding Port: 3/1 3/2
group 239.255.162.2 in 1 vlans
```

Syntax: show ip multicast <group-address>

This display shows the following information.

This Field...	Displays...
VLAN ID	The port-based VLAN to which the information listed applies.
Active	The IP address of the device that actively sends IGMP queries.
Router Ports	The ports that are connected to routers that support IP multicast.
Group	Address of the IP multicast group. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Forwarding Port	The forwarding ports for the IP multicast group.

Displaying Usage of Hardware Resource by Multicast Groups

Beginning with software release 07.8.00, you can display how much hardware resource (CAM and FID) is currently being used by multicast groups by entering a command such as the following at any level of the CLI:

```
BigIron# show ip multicast hardware
Hw resource is shared by groups with the same lower 23 bits
VLAN ID 100
Total number of HW resource in vlan: 3
1   G=XXX.0/128.1.22, ref_cnt 1, fid 08ac, NO cam, dma=
    Forwarding Port: 3/3
2   G=XXX.127/255.162.2, ref_cnt 1, fid 08aa, cam 8, dma=8,
    Forwarding Port: 3/1 3/2
3   G=XXX.127/255.163.2, ref_cnt 1, fid 08a9, cam 10, dma=8,
    Forwarding Port: 3/1 3/2
VLAN ID 4008
Total number of HW resource in vlan: 0
```

If you want to display the amount of hardware resource that is currently being used by a specific group, enter a command such as the following at any level of the CLI:

```
BigIron# show ip multicast hardware 239.255.163.2
VLAN ID 100
G=XXX.127/255.163.2, ref_cnt 1, fid 08a9, cam 10, dma=8,
  Forwarding Port: 3/1 3/2
group 239.255.163.2 in 1 vlans
```

Syntax: show ip multicast hardware [<group-address> | vlan <vlan-id>]

Enter the address of a group for <group-address> if you want to display the hardware resource usage of a particular group.

Likewise, enter the ID of a VLAN for <vlan-id> if you want display the hardware resource usage of groups in a VLAN.

The display shows the following information:

This Field...	Displays...
VLAN ID	The port-based VLAN to which the information listed below applies.
Total number of HW resource in VLAN	The number of resources in the VLAN.
G=	<p>Address of the IP multicast group that is using the entry. In the display above, "XXX.0/128.1.22" means that either group XXX.0.1.22 or XXX 128.1.22 or both is using this entry.</p> <p>The field ref_cnt shows the number of groups that are sharing this entry. Multiple groups could share one entry because only low 23 bits are significant.</p> <p>Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.</p>
Forwarding Port	The forwarding ports for the IP multicast group.

Displaying Software Resource Usage

Beginning with software release 07.8.00, you can display the amount of software resources used by each IGMP and PIM process that is enabled on a Layer 2 Switch by entering the following command at any level of the CLI:

```
BigIron# show ip multicast resource
          alloc in-use  avail  allo-fail  up-limit  size
pim neighbor      32     3    29       0     512    19
pim source-hash   256    3   253       0   10000   484
pim source       1024    7  1017       0  400000    6
pim source port   1024    7  1017       0  200000   13
igmp vlan struct    16    2    14       0    255   479
igmp mdb          256    3   253       0   10000   385
igmp hw resource   256    3   253       0   10000  5786
igmp port-age     2048    8  2040       0  100000    8
igmp leave        512    0   512       0 no-limit    8
In use: hw-res: 3, cam: 2, fid: 6
cpu forwarded packets: 411
```

Syntax: show ip multicast resource

Displaying Multicast Traffic Statistics

The **show ip multicast statistics** command has been available on Layer 2 Switches in previous releases. However, in software release 07.8.00, the report has been updated to show the following message if the Layer 2 Switch receives an IGMP V3 report :

```
*** Warning! IGMPv3 reports: 10
```

The warning shows the count of IGMP V3 that were received by a Layer 2 Switch. Refer to the *Foundry Switch and Router Command Line Interface Reference* for information on this command.

Displaying Multicast Information by VLAN

Beginning with software release 07.8.00, you can display multicast information for a specific VLAN by entering command such as the following at any level of the CLI:

```
BigIron# show ip multicast vlan 100
Only display vlan 100
VLAN ID 100
  Querier: 1.100.100.7, (port: 3/1)
  Router Ports: 3/1 3/2
  Total number of Multicast Group in vlan: 3
1   Group: 224.0.1.22, fid 08ac, NO cam
   Forwarding Port: 3/3
2   Group: 239.255.162.2, fid 08aa, cam 8
   Forwarding Port: 3/1 3/2
3   Group: 239.255.163.2, fid 08a9, cam 10
   Forwarding Port: 3/1 3/2
```

Syntax: show ip multicast vlan <vlan-id>

Enter the ID of the VLAN for <vlan-id>.

This Field...	Displays...
VLAN ID	The port-based VLAN to which the information listed below the VLAN ID applies. Each port-based VLAN is a separate Layer 2 broadcast domain.
Querier (port)	The IP address of the device that actively sends IGMP queries. The port on which the queries are being sent out.
Router Ports	The ports that are connected to a switch that support IP multicast.
Total Number of Multicast Group in VLAN	The total number of groups for which the VLAN's ports have received IGMP group membership reports, join messages, or prune messages.
Multicast Group	Address of the IP multicast group. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Forwarding Port	The forwarding ports for the IP multicast group.

Displaying PIM SM Snooping Information

Beginning with software release 07.8.00, you can display PIM SM snooping information for all groups by entering the following command at any level of the CLI on a Layer 2 Switch:

```
BigIron# show ip pimsm-snooping vlan 100
VLAN ID 100, total 3 entries
PIMSM Neighbor list:
    1.100.100.12      : 3/3 expire 120 s
    1.100.100.10      : 3/2 expire 170 s
    1.100.100.7       : 3/1 expire 160 s
1   Group: 224.0.1.22, fid 08ac, NO cam
    Forwarding Port: 3/3
    PIMv2 Group Port: 3/3
    (Source, Port) list: 1 entries
2   Group: 239.255.162.2, fid 08aa, cam 8
    Forwarding Port: 3/1 3/2
    PIMv2 Group Port: 3/1 3/2
    (Source, Port) list: 3 entries
3   Group: 239.255.163.2, fid 08a9, cam 10
    Forwarding Port: 3/1 3/2
    PIMv2 Group Port: 3/1 3/2
    (Source, Port) list: 3 entries
VLAN ID 4008, total 0 entries
PIMSM Neighbor list:
```

Syntax: show ip pimsm-snooping vlan <vlan-id>

Enter the ID of the VLAN for the **vlan** <vlan-id> parameter.

If you want to display PIM SM snooping information for one source or one group, enter a command as in the following example. The command also displays the (source, port) list of the group.

```
BigIron# show ip pimsm-snooping 239.255.163.2
Show pimsm snooping group 239.255.163.2 in all vlan
VLAN ID 100
Group: 239.255.163.2, fid 08a9, cam 10
Forwarding Port: 3/1 3/2
PIMv2 Group Port: 3/1 3/2
(Source, Port) list: 3 entries
  1 192.168.176.44, age=0, port: 3/2
  2 158.158.158.158, age=0, port: 3/1
  3 1.1.7.1, age=0, port: 3/2
```

Syntax: show ip pimsm-snooping <group-address> | <source-address>

If the address you entered is within the range of source addresses, then the router treats it as the source address. Likewise, if the address falls in the range of group addresses, then the router assumes that you are requesting a report for that group.

This display shows the following information.

This Field...	Displays...
VLAN ID	The port-based VLAN to which the information listed below apply and the number of members in the VLAN.
PIM SM Neighbor list	The PIM SM routers that are attached to the Layer 2 Switch's ports in the VLAN. The value following "expires" indicates how many seconds the Layer 2 Switch will wait for a hello message from the neighbor before determining that the neighbor is no longer present and removing the neighbor from the list.
Multicast Group	The IP address of the multicast group. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Forwarding Port	The port(s) attached to the group's receivers. A port is listed here when it receives a join message for the group, an IGMP membership report for the group, or both.
PIMv2 Group Port	The port(s) on which the Layer 2 Switch has received PIM SM join messages for the group.
Source, Port list	The IP address of each PIM SM source and the Layer 2 Switch ports connected to the receivers of the source.

Displaying PIM SM Snooping Information for a Specific Source in a Group

Beginning with software release 07.8.00, you can display PIM SM snooping information for a specific (source, group) pair by entering commands such as the following at any level of the CLI:

```
BigIron# show ip pimsm-snooping 239.255.163.2 192.168.176.44
Show pimsm snooping source 192.168.176.44, group 239.255.163.2 in all vlan
VLAN ID 100, G=239.255.163.2, S=192.168.176.44, age=0, port: 3/2
```

Syntax: show ip pimsm-snooping <group-address> <source-ip-address>

The router determines which address is the group address and which one is the source address based on the ranges that the address fall into. If the address is within the range of source addresses, then the router treats it as the source address. Likewise, if the address falls in the range of group addresses, then the router assumes it is a group address.

The output shows the following information.

This Field...	Displays...
VLAN ID	VLAN membership of the source
Group	Address of the group
Source	IP address of the source
Age	Age of the source.
Port	Port on which the source is sending traffic

Displaying Multicast Information on Layer 3 Switches

To display IP multicast traffic reduction information on Layer 3 Switches, enter the following command at any level of the CLI:

```
BigIron(config)# show ip multicast igmp-snooping
IP multicast is enabled - Passive

VLAN ID 100
Active 10.10.10.10
Router Ports 4/45
Number of Multicast Groups: 2
1   Group: 239.0.0.1 Num SG 1 Ports: 4/48 4/47 4/45
    IGMP report ports : 4/47 4/48
    1 Source: (10.10.10.2, 4/45) FID 0x08a4
2   Group: 239.0.0.10 Num SG 1 Ports: 4/48
    IGMP report ports : 4/48
    1 Source: (10.10.10.2, 4/45) FID 0x08a5
```

Syntax: show ip multicast igmp-snooping

This display shows the following information.

This Field...	Displays...
The IP multicast traffic snooping state	The first line of the display indicates whether IP multicast traffic snooping is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.
VLAN ID	The port-based VLAN to which the information listed below the VLAN ID applies. Each port-based VLAN is a separate Layer 2 broadcast domain.
Active	The IP address of the device that actively sends IGMP queries.
Router Ports	The ports that are connected to routers that support IP multicast.
Number of Multicast Group	The total number of groups for which the VLAN's ports have received IGMP group membership reports, join messages, or prune messages.
Multicast Group	An IP multicast group. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Port	The ports attached to receivers for the IP multicast group.
IGMP Report Port	The port(s) in this VLAN on which the Layer 3 Switch has received IGMP group membership reports for IP multicast groups.
Source, Port list	The IP address of each IGMP source and the Layer 3 Switch ports connected to the receivers of the source.

You also can display PIM SM information on Layer 3 Switches by entering the following command, at any level of the CLI:

```
BigIron(config)# show ip multicast pimsm-snooping
PIMSM snooping is enabled

VLAN ID 100
  PIMSM neighbour list:
    31.31.31.4 : 12/2 expires 142 s
    31.31.31.13 : 10/7 expires 136 s
    31.31.31.2 : 3/1 expires 172 s
Number of Multicast Groups: 2
1  Group: 239.255.162.4 Num SG 4
   Forwarding ports : 3/1 12/2
   PIMv2 *G join ports : 3/1 12/2
   1  Source: (165.165.165.165, 10/7) FID 0x0bb3
     SG join ports: 12/2 10/7
   2  Source: (161.161.161.161, 10/7) FID 0x0bb2
     SG join ports: 12/2 3/1
   3  Source: (158.158.158.158, 10/7) FID 0x0bb1
     SG join ports: 12/2 3/1
   4  Source: (170.170.170.170, 10/7) FID 0x0baf
     SG join ports: 3/1 10/7
     (S, G) age 0 s
2  Group: 239.255.163.2 Num SG 1
   Forwarding ports : 10/7 12/2
   PIMv2 *G join ports : 10/7 12/2
   1  Source: (165.165.165.165, 3/1) FID 0x0bb5
     SG join ports: 12/2 10/7
```

Syntax: show ip multicast pimsm-snooping

This display shows the following information.

This Field...	Displays...
The PIM SM traffic snooping state	The first line of the display indicates whether the feature is enabled or disabled; and if it is enabled, if it is passive or active. The PIM SM traffic snooping feature requires the IP multicast traffic reduction feature.
VLAN ID	The port-based VLAN to which the neighbors and groups listed below the VLAN ID apply. Each port-based VLAN is a separate Layer 2 broadcast domain. Note: PIM SM traffic snooping requires the source and the receivers to be in the same port-based VLAN on the Layer 3 Switch. If the source and receivers are in different port-based VLANs, the Layer 3 Switch blocks the multicast traffic.
PIM SM Neighbor list	The PIM SM routers that are attached to the Layer 3 Switch's ports in the VLAN. The value following "expires" indicates how many seconds the Layer 3 Switch will wait for a hello message from the neighbor before determining that the neighbor is no longer present and removing the neighbor from the list.

This Field...	Displays...
Number of Multicast Group	The total number of groups for which the VLAN's ports have received PIM join or prune messages and IGMP group membership reports.
Multicast Group	The IP address of the multicast group. The "Num SG" entry indicates how many Source to Group flows are created for that Multicast Group as there can be more than one source for a given group. Note: The fid and camindex values are used by Foundry Technical Support for troubleshooting.
Forwarding Port	The port(s) attached to the group's receivers. A port is listed here when it receives a join message for the group, an IGMP membership report for the group, or both.
PIMv2 Group Port	The port(s) on which the Layer 3 Switch has received PIM SM join messages for the group.
Source, Port list	The IP address of each PIM SM source and the Layer 3 Switch ports connected to the receivers of the source.
SG join ports:	Ports from which a join message was received. The Layer 3 Switch forwards the traffic only on this port.
(S, G) age	The actual aging value. If this entry shows the value 0 seconds, software age value is still 0 and the flow is programmed in the CAM. If the entry shows a value other than 0 seconds, then the CAM entry has aged out and the software aging has begun. Once this age value reaches the Group Age value the entry will be deleted from the table. Group age value can be can be from 10 – 1220 seconds. The default is 140 seconds.

Displaying IP Multicast Statistics

To display IP multicast statistics on a device, enter the following commands at any level of the CLI:

```
BigIron# show ip multicast statistics
IP multicast is enabled - Passive
```

```
VLAN ID 1
Reports Received:          34
Leaves Received:          21
General Queries Received: 60
Group Specific Queries Received: 2
Others Received:          0
General Queries Sent:     0
Group Specific Queries Sent: 0
```

```
VLAN ID 2
Reports Received:          0
Leaves Received:          0
General Queries Received: 60
Group Specific Queries Received: 2
Others Received:          0
General Queries Sent:     0
Group Specific Queries Sent: 0
```

The command in this example shows statistics for two port-based VLANs.

Syntax: show ip multicast statistics

Clearing IP Multicast Statistics

To clear IP multicast statistics on a device, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# clear ip multicast statistics
```

This command resets statistics counters for all the statistics displayed by the **show ip multicast statistics** command to zero.

Syntax: clear ip multicast statistics

Clearing IGMP Group Flows

To clear all the IGMP flows learned by the device, enter the following command at the Privileged EXEC level of the CLI:

```
BigIron# clear ip multicast all
```

The following example shows IGMP flows information listed by the **show ip multicast** command, followed by removal of the information by the **clear ip multicast all** command.

```
BigIron# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
Multicast Group: 239.255.162.5, Port: 4/4 4/13
Multicast Group: 239.255.162.4, Port: 4/10 4/13
```

```
BigIron# clear ip multicast all
```

```
BigIron# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
```

To clear the learned IGMP flows for a specific IP multicast group, enter a command such as the following:

```
BigIron# clear ip multicast group 239.255.162.5
```

The following example shows how to clear the IGMP flows for a specific group and retain reports for other groups.

```
BigIron# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
Multicast Group: 239.255.162.5, Port: 4/4 4/13
Multicast Group: 239.255.162.4, Port: 4/10 4/13
```

```
BigIron# clear ip multicast group 239.255.162.5
```

```
BigIron# show ip multicast
IP multicast is enabled - Active
VLAN ID 1
Active 192.168.2.30 Router Ports 4/13
Multicast Group: 239.255.162.4, Port: 4/10 4/13
```

Syntax: clear ip multicast all | group <group-id>

The **all** parameter clears the learned flows for all groups.

The **group** <group-id> parameter clears the flows for the specified group but does not clear the flows for other groups.

Chapter 25

Configuring GARP VLAN Registration Protocol (GVRP)

GARP VLAN Registration Protocol (GVRP) is a Generic Attribute Registration Protocol (GARP) application that provides VLAN registration service by means of dynamic configuration (registration) and distribution of VLAN membership information.

NOTE: This feature is supported in the B2R, B2S, and B2P images only.

A Foundry device enabled for GVRP can do the following:

- Learn about VLANs from other Foundry devices and configure those VLANs on the ports that learn about the VLANs. The device listens for GVRP Protocol Data Units (PDUs) from other devices, and implements the VLAN configuration information in the PDUs.
- Advertise VLANs configured on the device to other Foundry devices. The device sends GVRP PDUs advertising its VLANs to other devices. GVRP advertises statically configured VLANs and VLANs learned from other devices through GVRP.

GVRP enables a Foundry device to dynamically create 802.1Q-compliant VLANs on links with other devices that are running GVRP. GVRP reduces the chances for errors in VLAN configuration by automatically providing VLAN ID consistency across the network. You can use GVRP to propagate VLANs to other GVRP-aware devices automatically, without the need to manually configure the VLANs on each device. In addition, if the VLAN configuration on a device changes, GVRP automatically changes the VLAN configurations of the affected devices.


The Foundry implementation of GARP and GVRP is based on the following standards:

- ANSI/IEEE standard 802.1D, 1998 edition
- IEEE standard 802.1Q, 1998 edition; approved December 8, 1998
- IEEE draft P802.1w/D10, March 26, 2001
- IEEE draft P802.1u/D9, November 23, 2000
- IEEE draft P802.1t/D10, November 20, 2000

Application Examples

Figure 25.1 shows an example of a network that uses GVRP. This section describes various ways you can use GVRP in a network such as this one. “CLI Examples” on page 25-17 lists the CLI commands to implement the applications of GVRP described in this section.

Figure 25.1 Example of GVRP



In this example, a core device is attached to three edge devices. Each of the edge devices is attached to other edge devices or host stations (represented by the clouds).

The effects of GVRP in this network depend on which devices the feature is enabled on, and whether both learning and advertising are enabled. In this type of network (a core device and edge devices), you can have the following four combinations:

- Dynamic core and fixed edge
- Dynamic core and dynamic edge
- Fixed core and dynamic edge
- Fixed core and fixed edge

Dynamic Core and Fixed Edge

In this configuration, all ports on the core device are enabled to learn and advertise VLAN information. The edge devices are configured to advertise their VLAN configurations on the ports connected to the core device. GVRP learning is disabled on the edge devices.

Core Device	Edge Device A	Edge Device B	Edge Device C
<p>GVRP is enabled on all ports.</p> <p>Both learning and advertising are enabled.</p> <p>Note: Since learning is disabled on all the edge devices, advertising on the core device has no effect in this configuration.</p>	<p>GVRP is enabled on port 4/24. Learning is disabled.</p> <p>VLAN 20</p> <p>Port 2/1 (untagged)</p> <p>Port 4/24 (tagged)</p> <p>VLAN 40</p> <p>Port 4/1 (untagged)</p> <p>Port 4/24 (tagged)</p>	<p>GVRP is enabled on port 4/1. Learning is disabled.</p> <p>VLAN 20</p> <p>Port 2/24 (untagged)</p> <p>Port 4/1 (tagged)</p> <p>VLAN 30</p> <p>Port 4/24 (untagged)</p> <p>Port 4/1 (tagged)</p>	<p>GVRP is enabled on port 4/1. Learning is disabled.</p> <p>VLAN 30</p> <p>Port 2/24 (untagged)</p> <p>Port 4/1 (tagged)</p> <p>VLAN 40</p> <p>Port 4/24 (untagged)</p> <p>Port 4/1 (tagged)</p>

In this configuration, the edge devices are statically (manually) configured with VLAN information. The core device dynamically configures itself to be a member of each of the edge device's VLANs. The operation of GVRP on the core device results in the following VLAN configuration on the device:

- VLAN 20
 - 1/24 (tagged)
 - 6/24 (tagged)
- VLAN 30
 - 6/24 (tagged)
 - 8/17 (tagged)
- VLAN 40
 - 1/24 (tagged)
 - 8/17 (tagged)

VLAN 20 traffic can now travel through the core between edge devices A and B. Likewise, VLAN 30 traffic can travel between B and C and VLAN 40 traffic can travel between A and C. If an edge device is moved to a different core port or the VLAN configuration of an edge device is changed, the core device automatically reconfigures itself to accommodate the change.

Notice that each of the ports in the dynamically created VLANs is tagged. All GVRP VLAN ports configured by GVRP are tagged, to ensure that the port can be configured for additional VLANs.

NOTE: This example assumes that the core device has no static VLANs configured. However, you can have static VLANs on a device that is running GVRP. GVRP can dynamically add other ports to the statically configured VLANs but cannot delete statically configured ports from the VLANs.

Dynamic Core and Dynamic Edge

GVRP is enabled on the core device and on the edge devices. This type of configuration is useful if the devices in the edge clouds are running GVRP and advertise their VLANs to the edge devices. The edge devices learn the VLANs and also advertise them to the core. In this configuration, you do not need to statically configure the VLANs on the edge or core devices, although you can have statically configured VLANs on the devices. The devices learn the VLANs from the devices in the edge clouds.

Fixed Core and Dynamic Edge

GVRP learning is enabled on the edge devices. The VLANs on the core device are statically configured, and the core device is enabled to advertise its VLANs but not to learn VLANs. The edge devices learn the VLANs from the core.

Fixed Core and Fixed Edge

The VLANs are statically configured on the core and edge devices. On each edge device, VLAN advertising is enabled but learning is disabled. GVRP is not enabled on the core device. This configuration enables the devices in the edge clouds to learn the VLANs configured on the edge devices.

VLAN Names

The **show vlans** command lists VLANs created by GVRP as “GVRP_VLAN_<vlan-id>”. VLAN names for statically configured VLANs are not affected. To distinguish between statically-configured VLANs that you add to the device and VLANs that you convert from GVRP-configured VLANs into statically-configured VLANs, the **show vlans** command displays a converted VLAN’s name as “STATIC_VLAN_<vlan-id>”.

Configuration Considerations

- If you disable GVRP, all GVRP configuration information is lost if you save the configuration change (**write memory** command) and then reload the software. However, if you reload the software without first saving the configuration change, the GVRP configuration is restored following a software reload.
- The maximum number of VLANS supported on a device enabled for GVRP is the same as the maximum number on a device that is not enabled for GVRP.
 - To display the maximum number of VLANs allowed on your device, enter the **show default values** command. See the “vlan” row in the System Parameters section. Make sure you allow for the default VLAN (1), the GVRP base VLAN (4093), and the Single STP VLAN (4094). These VLANs are maintained as “Registration Forbidden” in the GVRP database. Registration Forbidden VLANs cannot be advertised or learned by GVRP.
 - To increase the maximum number of VLANs supported on the device, enter the **system-max vlan <num>** command at the global CONFIG level of the CLI, then save the configuration and reload the software. The maximum number you can specify is listed in the Maximum column of the **show default values** display.
- The default VLAN (VLAN 1) is not advertised by the Foundry implementation of GVRP. The default VLAN contains all ports that are not members of statically configured VLANs or VLANs enabled for GVRP.

NOTE: The default VLAN has ID 1 by default. You can change the VLAN ID of the default VLAN, but only before GVRP is enabled. You cannot change the ID of the default VLAN after GVRP is enabled.

- Single STP must be enabled on the device. Foundry’s implementation of GVRP requires Single STP. If you do not have any statically configured VLANs on the device, you can enable Single STP as follows:

```
BigIron(config)# vlan 1
BigIron(config-vlan-1)# exit
BigIron(config)# span
BigIron(config)# span single
```

These commands enable configuration of the default VLAN (VLAN 1), which contains all the device’s ports, and enable STP and Single STP.

- All VLANs that are learned dynamically through GVRP are added to the single spanning tree.
- All ports that are enabled for GVRP become tagged members of the GVRP base VLAN (4093). If you need to use this VLAN ID for another VLAN, you can change the GVRP VLAN ID. See “Changing the GVRP Base VLAN ID” on page 25-5. The software adds the GVRP base VLAN to the single spanning tree.

- All VLAN ports added by GVRP are tagged.
- GVRP is supported only for tagged ports or for untagged ports that are members of the default VLAN. GVRP is not supported for ports that are untagged and are members of a VLAN other than the default VLAN.
- To configure GVRP on a trunk group, enable the protocol on the primary port in the trunk group. The GVRP configuration of the primary port is automatically applied to the other ports in the trunk group.
- You can use GVRP on a device even if the device has statically configured VLANs. GVRP does not remove any ports from the statically configured VLANs, although GVRP can add ports to the VLANs. GVRP advertises the statically configured VLANs. Ports added by GVRP do not appear in the running-config and will not appear in the startup-config file when save the configuration. You can manually add a port to make the port a permanent member of the VLAN. After you manually add the port, the port will appear in the running-config and be saved to the startup-config file when you save the configuration.
- VLANs created by GVRP do not support virtual routing interfaces or protocol-based VLANs. virtual routing interfaces and protocol-based VLANs are still supported on statically configured VLANs even if GVRP adds ports to those VLANs.
- You cannot manually configure any parameters on a VLAN that is created by GVRP. For example, you cannot change STP parameters for the VLAN.
- The GVRP timers (Join, Leave, and Leaveall) must be set to the same values on all the devices that are exchanging information using GVRP.
- If the network has a large number of VLANs, the GVRP traffic can use a lot of CPU resources. If you notice high CPU utilization after enabling GVRP, set the GVRP timers to longer values. In particular, set the Leaveall timer to a longer value. See “Changing the GVRP Timers” on page 25-7.
- The feature is supported only on Ethernet ports.

NOTE: If you plan to change the GVRP base VLAN ID (4093) or the maximum configurable value for the Leaveall timer (300000 ms by default), you must do so before you enable GVRP.

Configuring GVRP

To configure a device for GVRP, globally enable support for the feature, then enable the feature on specific ports. Optionally, you can disable VLAN learning or advertising on specific interfaces.

You also can change the protocol timers and change the GVRP base VLAN ID.

Changing the GVRP Base VLAN ID

By default, GVRP uses VLAN 4093 as a base VLAN for the protocol. All ports that are enabled for GVRP become tagged members of this VLAN. If you need to use VLAN ID 4093 for a statically configured VLAN, you can change the GVRP base VLAN ID.

NOTE: If you want to change the GVRP base VLAN ID, you must do so before enabling GVRP.

To change the GVRP base VLAN ID, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# gvrp-base-vlan-id 1001
```

This command changes the GVRP VLAN ID from 4093 to 1001.

Syntax: [no] gvrp-base-vlan-id <vlan-id>

The <vlan-id> parameter specifies the new VLAN ID. You can specify a VLAN ID from 2 – 4092 or 4095.

Increasing the Maximum Configurable Value of the Leaveall Timer

By default, the highest value you can specify for the Leaveall timer is 300000 ms. You can increase the maximum configurable value of the Leaveall timer to 1000000 ms.

NOTE: You must enter this command before enabling GVRP. Once GVRP is enabled, you cannot change the maximum Leaveall timer value.

NOTE: This command does not change the default value of the Leaveall timer itself. The command only changes the maximum value to which you can set the Leaveall timer.

To increase the maximum value you can specify for the Leaveall timer, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# gvrp-max-leaveall-timer 1000000
```

Syntax: [no] gvrp-max-leaveall-timer <ms>

The <ms> parameter specifies the maximum number of ms to which you can set the Leaveall timer. You can specify from 300000 – 1000000 (one million) ms. The value must be a multiple of 100 ms. The default is 300000 ms.

Enabling GVRP

To enable GVRP, enter commands such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable all
```

The first command globally enables support for the feature and changes the CLI to the GVRP configuration level. The second command enables GVRP on all ports on the device.

The following command enables GVRP on ports 1/24, 6/24, and 8/17:

```
BigIron(config-gvrp)# enable ethernet 1/24 ethernet 6/24 ethernet 8/17
```

Syntax: [no] gvrp-enable

Syntax: [no] enable all | ethernet <portnum> [ethernet <portnum> | to <portnum>]

The **all** parameter enables GVRP on all ports.

The **ethernet <portnum>** [**ethernet <portnum>** | **to <portnum>**] parameter enables GVRP on the specified list or range of Ethernet ports.

- To specify a list, enter each port as **ethernet <portnum>** followed by a space. For example, to enable GVRP on three Ethernet ports, enter the following command: **enable ethernet 1/24 ethernet 6/24 ethernet 8/17**
- To specify a range, enter the first port in the range as **ethernet <portnum>** followed by **to** followed by the last port in the range. For example, to add ports 1/1 – 1/8, enter the following command: **enable ethernet 1/1 to 1/8**

You can combine lists and ranges in the same command. For example: **enable ethernet 1/1 to 1/8 ethernet 1/24 ethernet 6/24 ethernet 8/17**

Disabling VLAN Advertising

To disable VLAN advertising on a port enabled for GVRP, enter a command such as the following at the GVRP configuration level:

```
BigIron(config-gvrp)# block-applicant ethernet 1/24 ethernet 6/24 ethernet 8/17
```

This command disables advertising of VLAN information on ports 1/24, 6/24, and 8/17.

Syntax: [no] block-applicant all | ethernet <portnum> [ethernet <portnum> | to <portnum>]

NOTE: Leaveall messages are still sent on the GVRP ports.

Disabling VLAN Learning

To disable VLAN learning on a port enabled for GVRP, enter a command such as the following at the GVRP configuration level:

```
BigIron(config-gvrp)# block-learning ethernet 6/24
```

This command disables learning of VLAN information on port 6/24.

NOTE: The port still advertises VLAN information unless you also disable VLAN advertising.

Syntax: [no] block-learning all | ethernet <portnum> [ethernet <portnum> | to <portnum>]

Changing the GVRP Timers

GVRP uses the following timers:

- Join – The maximum number of milliseconds (ms) a device's GVRP interfaces wait before sending VLAN advertisements on the interfaces. The actual interval between Join messages is randomly calculated to a value between 0 and the maximum number of milliseconds specified for Join messages. You can set the Join timer to a value from 200 – one third the value of the Leave timer. The default is 200 ms.
- Leave – The number of ms a GVRP interface waits after receiving a Leave message on the port to remove the port from the VLAN indicated in the Leave message. If the port receives a Join message before the Leave timer expires, GVRP keeps the port in the VLAN. Otherwise, the port is removed from the VLAN. When a port receives a Leave message, the port's GVRP state is changed to Leaving. Once the Leave timer expires, the port's GVRP state changes to Empty. You can set the Leave timer to a value from three times the Join timer – one fifth the value of the Leaveall timer. The default is 600 ms.

NOTE: When all ports in a dynamically created VLAN (one learned through GVRP) leave the VLAN, the VLAN is immediately deleted from the device's VLAN database. However, this empty VLAN is still maintained in the GVRP database for an amount of time equal to the following:

$(\text{number-of-GVRP-enabled-up-ports}) * (2 * \text{join-timer})$

While the empty VLAN is in the GVRP database, the VLAN does not appear in the **show vlans** display but does still appear in the **show gvrp vlan all** display.

- Leaveall – The minimum interval at which GVRP sends Leaveall messages on all GVRP interfaces. Leaveall messages ensure that the GVRP VLAN membership information is current by aging out stale VLAN information and adding information for new VLAN memberships, if the information is missing. A Leaveall message instructs the port to change the GVRP state for all its VLANs to Leaving, and remove them unless a Join message is received before the Leave timer expires. By default, you can set the Leaveall timer to a value from five times the Leave timer – maximum value allowed by software (configurable from 300000 – 1000000 ms). The default is 10000.

NOTE: The actual interval is a random value between the Leaveall interval and $1.5 * \text{the Leaveall time}$ or the maximum Leaveall time, whichever is lower.

NOTE: You can increase the maximum configurable value of the Leaveall timer from 300000 ms up to 1000000 ms using the **gvrp-max-leaveall-timer** command. (See "Increasing the Maximum Configurable Value of the Leaveall Timer" on page 25-6.)

Timer Configuration Requirements

- All timer values must be in multiples of 100 ms.
- The Leave timer must be $\geq 3 \times$ the Join timer.
- The Leaveall timer must be $\geq 5 \times$ the Leave timer.
- The GVRP timers must be set to the same values on all the devices that are exchanging information using GVRP.

Changing the Join, Leave, and Leaveall Timers

The same CLI command controls changes to the Join, Leave, and Leaveall timers. To change values to the timers, enter a command such as the following:

```
BigIron(config-gvrp)# join-timer 1000 leave-timer 3000 leaveall-timer 15000
```

This command changes the Join timer to 1000 ms, the Leave timer to 3000 ms, and the Leaveall timer to 15000.

Syntax: [no] join-timer <ms> leave-timer <ms> leaveall-timer <ms>

NOTE: When you enter this command, all the running GVRP timers are canceled and restarted using the new times specified by the command.

Resetting the Timers to Their Defaults

To reset the Join, Leave, and Leaveall timers to their default values, enter the following command:

```
BigIron(config-gvrp)# default-timers
```

Syntax: default-timers

This command resets the timers to the following values:

- Join – 200 ms
- Leave – 600 ms
- Leaveall – 10000 ms

Converting a VLAN Created by GVRP into a Statically-Configured VLAN

You cannot configure VLAN parameters on VLANs created by GVRP. Moreover, VLANs and VLAN ports added by GVRP do not appear in the running-config and cannot be saved in the startup-config file.

To be able to configure and save VLANs or ports added by GVRP, you must convert the VLAN ports to statically-configured ports.

To convert a VLAN added by GVRP into a statically-configured VLAN, add the ports using commands such as the following:

```
BigIron(config)# vlan 22
BigIron(config-vlan-22)# tagged ethernet 1/1 to 1/8
```

These commands convert GVRP-created VLAN 22 containing ports 1/1 through 1/8 into statically-configured VLAN 22.

Syntax: [no] vlan <vlan-id>

Syntax: [no] tagged ethernet <portnum> [to <portnum> | ethernet <portnum>]

Use the same commands to statically add ports that GVRP added to a VLAN.

NOTE: You cannot add the VLAN ports as untagged ports.

NOTE: After you convert the VLAN, the VLAN name changes from “GVRP_VLAN_<vlan-id>” to “STATIC_VLAN_<vlan-id>”.

Displaying GVRP Information

You can display the following GVRP information:

- GVRP configuration information
- GVRP VLAN information
- GVRP statistics
- CPU utilization statistics
- GVRP diagnostic information

Displaying GVRP Configuration Information

To display GVRP configuration information, enter a command such as the following:

```
BigIron(config)# show gvrp
GVRP is enabled on the system
```

```
GVRP BASE VLAN ID      : 4093
GVRP MAX Leaveall Timer : 300000 ms
```

```
GVRP Join Timer        : 200 ms
GVRP Leave Timer       : 600 ms
GVRP Leave-all Timer  : 10000 ms
```

```
=====  
Configuration that is being used:
```

```
block-learning ethe 1/3
block-applicant ethe 2/7 ethe 2/11
enable ethe 1/1 to 1/7 ethe 2/1 ethe 2/7 ethe 2/11
```

```
=====  
Spanning Tree: SINGLE SPANNING TREE
Dropped Packets Count: 0
```

```
=====  
Number of VLANs in the GVRP Database: 15
Maximum Number of VLANs that can be present: 4095
```

```
=====  
Syntax: show gvrp [ethernet <port-num>]
```

This display shows the following information.

Table 25.1: CLI Display of Summary GVRP Information

This Field...	Displays...
Protocol state	The state of GVRP. The display shows one of the following: <ul style="list-style-type: none"> • GVRP is disabled on the system • GVRP is enabled on the system
GVRP BASE VLAN ID	The ID of the base VLAN used by GVRP.
GVRP MAX Leaveall Timer	The maximum number of ms to which you can set the Leaveall timer. <p>Note: To change the maximum value, see “Increasing the Maximum Configurable Value of the Leaveall Timer” on page 25-6.</p>
GVRP Join Timer	The value of the Join timer. <p>Note: For descriptions of the Join, Leave, and Leaveall timers or to change the timers, see “Changing the GVRP Timers” on page 25-7.</p>
GVRP Leave Timer	The value of the Leave timer.
GVRP Leave-all Timer	The value of the Leaveall timer.
Configuration that is being used	The configuration commands used to enable GVRP on individual ports. If GVRP learning or advertising is disabled on a port, this information also is displayed.
Spanning Tree	The type of STP enabled on the device. <p>Note: The current release supports GVRP only with Single STP.</p>
Dropped Packets Count	The number of GVRP packets that the device has dropped. A GVRP packet can be dropped for either of the following reasons: <ul style="list-style-type: none"> • GVRP packets are received on a port on which GVRP is not enabled. <p>Note: If GVRP support is not globally enabled, the device does not drop the GVRP packets but instead forwards them at Layer 2.</p> • GVRP packets are received with an invalid GARP Protocol ID. The protocol ID must always be 0x0001.
Number of VLANs in the GVRP Database	The number of VLANs in the GVRP database. <p>Note: This number includes the default VLAN (1), the GVRP base VLAN (4093), and the single STP VLAN (4094). These VLANs are not advertised by GVRP but are maintained as “Registration Forbidden”.</p>
Maximum Number of VLANs that can be present	The maximum number of VLANs that can be configured on the device. This number includes statically configured VLANs, VLANs learned through GVRP, and VLANs 1, 4093, and 4094. <p>To change the maximum number of VLANs the device can have, use the system-max vlan <num> command. See “Displaying and Modifying System Parameter Default Settings” on page 9-48.</p>

To display detailed GVRP information for an individual port, enter a command such as the following:

```
BigIron(config)# show gvrp ethernet 2/1
Port 2/1 -
  GVRP Enabled      : YES
  GVRP Learning     : ALLOWED
  GVRP Applicant    : ALLOWED
  Port State        : UP
  Forwarding        : YES

VLAN Membership:      [VLAN-ID]          [MODE ]
                     1                  FORBIDDEN
                     2                  FIXED
                     1001               NORMAL
                     1003               NORMAL
                     1004               NORMAL
                     1007               NORMAL
                     1009               NORMAL
                     1501               NORMAL
                     2507               NORMAL
                     4001               NORMAL
                     4093               FORBIDDEN
                     4094               FORBIDDEN
```

This display shows the following information.

Table 25.2: CLI Display of Detailed GVRP Information for a Port

This Field...	Displays...
Port number	The port for which information is being displayed.
GVRP Enabled	Whether GVRP is enabled on the port.
GVRP Learning	Whether the port can learn VLAN information from GVRP.
GVRP Applicant	Whether the port can advertise VLAN information into GVRP.
Port State	The port's link state, which can be UP or DOWN.
Forwarding	Whether the port is in the GVRP Forwarding state: <ul style="list-style-type: none"> • NO – The port is in the Blocking state. • YES – The port is in the Forwarding state.

Table 25.2: CLI Display of Detailed GVRP Information for a Port (Continued)

This Field...	Displays...
VLAN Membership	<p>The VLANs of which the port is a member. For each VLAN, the following information is shown:</p> <ul style="list-style-type: none"> • VLAN ID – The VLAN's ID. • Mode – The type of VLAN, which can be one of the following: <ul style="list-style-type: none"> • FIXED – The port will always be a member of this VLAN and the VLAN will always be advertised on this port by GVRP. A port becomes FIXED when you configure the port as a tagged member of a statically configured VLAN. • FORBIDDEN – The VLAN is one of the special VLANs that is not advertised or learned by GVRP. In the current release, the following VLANs are forbidden: the default VLAN (1), the GVRP base VLAN (4093), or the Single STP VLAN (4094). • NORMAL – The port became a member of this VLAN after learning about the VLAN through GVRP. The port's membership in the VLAN depends on GVRP. If the VLAN is removed from the ports that send GVRP advertisements to this device, then the port will stop being a member of the VLAN.

Displaying GVRP VLAN Information

To display information about all the VLANs on the device, enter the following command:

```
BigIron(config)# show gvrp vlan brief
```

```
Number of VLANs in the GVRP Database: 7
Maximum Number of VLANs that can be present: 4095
```

[VLAN-ID]	[MODE]	[VLAN-INDEX]
1	STATIC-DEFAULT	0
7	STATIC	2
11	STATIC	4
1001	DYNAMIC	7
1003	DYNAMIC	8
4093	STATIC-GVRP-BASE-VLAN	6
4094	STATIC-SINGLE-SPAN-VLAN	5

=====

Syntax: show gvrp vlan all | brief | <vlan-id>

This display shows the following information.

Table 25.3: CLI Display of Summary VLAN Information for GVRP

This Field...	Displays...
Number of VLANs in the GVRP Database	The number of VLANs in the GVRP database. Note: This number includes the default VLAN (1), the GVRP base VLAN (4093), and the single STP VLAN (4094). These VLANs are not advertised by GVRP but are included in the total count.
Maximum Number of VLANs that can be present	The maximum number of VLANs that can be configured on the device. This number includes statically configured VLANs, VLANs learned through GVRP, and VLANs 1, 4093, and 4094. To change the maximum number of VLANs the device can have, use the system-max vlan <num> command. See “Displaying and Modifying System Parameter Default Settings” on page 9-48.
VLAN-ID	The VLAN ID.
MODE	The type of VLAN, which can be one of the following: <ul style="list-style-type: none"> • STATIC – The VLAN is statically configured and cannot be removed by GVRP. This includes VLANs you have configured as well as the default VLAN (1), base GVRP VLAN (4093), and Single STP VLAN (4094). • DYNAMIC – The VLAN was learned through GVRP.
VLAN-INDEX	A number used as an index into the internal database.

To display detailed information for a specific VLAN, enter a command such as the following:

```
BigIron(config)# show gvrp vlan 1001

VLAN-ID: 1001, VLAN-INDEX: 7, STATIC: NO, DEFAULT: NO, BASE-VLAN: NO
Timer to Delete Entry Running: NO
Legend: [S=Slot]

Forbidden Members: None

Fixed Members: None

Normal(Dynamic) Members: (S2) 1
```

This display shows the following information.

Table 25.4: CLI Display of Summary VLAN Information for GVRP

This Field...	Displays...
VLAN-ID	The VLAN ID.
VLAN-INDEX	A number used as an index into the internal database.
STATIC	Whether the VLAN is a statically configured VLAN.

Table 25.4: CLI Display of Summary VLAN Information for GVRP (Continued)

This Field...	Displays...
DEFAULT	Whether this is the default VLAN.
BASE-VLAN	Whether this is the base VLAN for GVRP.
Timer to Delete Entry Running	Whether all ports have left the VLAN and the timer to delete the VLAN itself is running. The timer is described in the note for the Leave timer in "Changing the GVRP Timers" on page 25-7.
Legend	The meanings of the letter codes used in other parts of the display.
Forbidden Members	The ports that cannot become members of a VLAN advertised or learned by GVRP.
Fixed Members	The ports that are statically configured members of the VLAN. GVRP cannot remove these ports.
Normal(Dynamic) Members	The ports that were added by GVRP. These ports also can be removed by GVRP.
MODE	The type of VLAN, which can be one of the following: <ul style="list-style-type: none"> • STATIC – The VLAN is statically configured and cannot be removed by GVRP. This includes VLANs you have configured as well as the default VLAN (1), base GVRP VLAN (4093), and Single STP VLAN (4094). • DYNAMIC – The VLAN was learned through GVRP.

To display detailed information for all VLANs, enter the **show gvrp vlan all** command.

Displaying GVRP Statistics

To display GVRP statistics for a port, enter a command such as the following:

```
BigIron(config)# show gvrp statistics ethernet 2/1
PORT 2/1 Statistics:
  Leave All Received           : 147
  Join Empty Received         : 4193
  Join In Received            : 599
  Leave Empty Received        : 0
  Leave In Received           : 0
  Empty Received              : 588
  Leave All Transmitted       : 157
  Join Empty Transmitted      : 1794
  Join In Transmitted         : 598
  Leave Empty Transmitted     : 0
  Leave In Transmitted        : 0
  Empty Transmitted           : 1248
  Invalid Messages/Attributes Skipped : 0
  Failed Registrations        : 0
```

Syntax: show gvrp statistics all | ethernet <port-num>

This display shows the following information for the port.

Table 25.5: CLI Display of GVRP Statistics

This Field...	Displays...
Leave All Received	The number of Leaveall messages received.
Join Empty Received	The number of Join Empty messages received.
Join In Received	The number of Join In messages received.
Leave Empty Received	The number of Leave Empty messages received.
Leave In Received	The number of Leave In messages received.
Empty Received	The number of Empty messages received.
Leave All Transmitted	The number of Leaveall messages sent.
Join Empty Transmitted	The number of Join Empty messages sent.
Join In Transmitted	The number of Join In messages sent.
Leave Empty Transmitted	The number of Leave Empty messages sent.
Leave In Transmitted	The number of Leave In messages sent.
Empty Transmitted	The number of Empty messages sent.
Invalid Messages/Attributes Skipped	<p>The number of invalid messages or attributes received or skipped. This can occur in the following cases:</p> <ul style="list-style-type: none"> • The incoming GVRP PDU has an incorrect length. • "End of PDU" was reached before the complete attribute could be parsed. • The Attribute Type of the attribute that was being parsed was not the GVRP VID Attribute Type (0x01). • The attribute that was being parsed had an invalid attribute length. • The attribute that was being parsed had an invalid GARP event. • The attribute that was being parsed had an invalid VLAN ID. The valid range is 1 – 4095.
Failed Registrations	<p>The number of failed registrations that have occurred. A failed registration can occur for the following reasons:</p> <ul style="list-style-type: none"> • Join requests were received on a port that was blocked from learning dynamic VLANs (GVRP Blocking state). • An entry for a new GVRP VLAN could not be created in the GVRP database.

To display GVRP statistics for all ports, enter the **show gvrp statistics all** command.

Displaying CPU Utilization Statistics

You can display CPU utilization statistics for GVRP.

To display CPU utilization statistics for GVRP for the previous one-second, one-minute, five-minute, and fifteen-minute intervals, enter the following command at any level of the CLI:

```
BigIron# show process cpu
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP             0.01       0.03       0.09       0.22        9
BGP             0.00       0.00       0.00       0.00        0
GVRP          0.00      0.03      0.04      0.07       4
ICMP            0.00       0.00       0.00       0.00        0
IP              0.00       0.00       0.00       0.00        0
OSPF            0.00       0.00       0.00       0.00        0
RIP             0.00       0.00       0.00       0.00        0
STP             0.00       0.00       0.00       0.00        0
VRRP            0.00       0.00       0.00       0.00        0
```

If the software has been running less than 15 minutes (the maximum interval for utilization statistics), the command indicates how long the software has been running. Here is an example:

```
BigIron# show process cpu
The system has only been up for 6 seconds.
Process Name    5Sec(%)    1Min(%)    5Min(%)    15Min(%)    Runtime(ms)
ARP             0.01       0.00       0.00       0.00        0
BGP             0.00       0.00       0.00       0.00        0
GVRP            0.00       0.00       0.00       0.00        0
ICMP            0.01       0.00       0.00       0.00        1
IP              0.00       0.00       0.00       0.00        0
OSPF            0.00       0.00       0.00       0.00        0
RIP             0.00       0.00       0.00       0.00        0
STP             0.00       0.00       0.00       0.00        0
VRRP            0.00       0.00       0.00       0.00        0
```

To display utilization statistics for a specific number of seconds, enter a command such as the following:

```
BigIron# show process cpu 2
Statistics for last 1 sec and 80 ms
Process Name    Sec(%)    Time(ms)
ARP             0.00      0
BGP             0.00      0
GVRP            0.01      1
ICMP            0.00      0
IP              0.00      0
OSPF            0.00      0
RIP             0.00      0
STP             0.01      1
VRRP            0.00      0
```

When you specify how many seconds' worth of statistics you want to display, the software selects the sample that most closely matches the number of seconds you specified. In this example, statistics are requested for the previous two seconds. The closest sample available is actually for the previous 1 second plus 80 milliseconds.

Syntax: show process cpu [<num>]

The <num> parameter specifies the number of seconds and can be from 1 – 900. If you use this parameter, the command lists the usage statistics only for the specified number of seconds. If you do not use this parameter, the command lists the usage statistics for the previous one-second, one-minute, five-minute, and fifteen-minute intervals.

Displaying GVRP Diagnostic Information

To display diagnostic information, enter the following command:

```
BigIron# debug gvrp packets
      GVRP:  Packets debugging is on
GVRP: 0x2095ced4: 01 80 c2 00 00 21 00 e0 52 ab 87 40 00 3a 42 42
GVRP: 0x2095cee4: 03 00 01 01 02 00 04 05 00 02 04 05 00 07 04 05
GVRP: 0x2095cef4: 00 09 04 05 00 0b 04 02 03 e9 04 01 03 eb 04 01
GVRP: 0x2095cf04: 03 ec 04 01 03 ef 04 01 03 f1 04 01 05 dd 04 01
GVRP: 0x2095cf14: 09 cb 04 01 0f a1 00 00
GVRP: Port 2/1 RCV
GVRP: 0x2095ced4: 01 80 c2 00 00 21 00 e0 52 ab 87 40 00 28 42 42
GVRP: 0x2095cee4: 03 00 01 01 04 02 03 e9 04 01 03 eb 04 01 03 ec
GVRP: 0x2095cef4: 04 01 03 ef 04 01 03 f1 04 01 05 dd 04 01 09 cb
GVRP: 0x2095cf04: 04 01 0f a1 00 00
GVRP: Port 2/1 TX
GVRP: 0x207651b8: 01 80 c2 00 00 21 00 04 80 2c 0e 20 00 3a 42 42
GVRP: 0x207651c8: 03 00 01 01 02 00 04 05 03 e9 04 05 03 eb 04 05
GVRP: 0x207651d8: 03 ec 04 05 03 ef 04 05 03 f1 04 05 05 dd 04 05
GVRP: 0x207651e8: 09 cb 04 05 0f a1 04 02 00 02 04 01 00 07 04 01
GVRP: 0x207651f8: 00 09 04 01 00 0b 00 00
GVRP: Port 2/1 TX
GVRP: 0x207651b8: 01 80 c2 00 00 21 00 04 80 2c 0e 20 00 18 42 42
GVRP: 0x207651c8: 03 00 01 01 04 02 00 02 04 01 00 07 04 01 00 09
GVRP: 0x207651d8: 04 01 00 0b 00 00
```

Syntax: debug gvrp packets

Clearing GVRP Statistics

To clear the GVRP statistics counters, enter a command such as the following:

```
BigIron# clear gvrp statistics all
```

This command clears the counters for all ports. To clear the counters for a specific port only, enter a command such as the following:

```
BigIron# clear gvrp statistics ethernet 2/1
```

Syntax: clear gvrp statistics all | ethernet <portnum>

CLI Examples

The following sections show the CLI commands for implementing the applications of GVRP described in “Application Examples” on page 25-1.

NOTE: Although some of the devices in these configuration examples do not have statically configured VLANs, this is not a requirement. You always can have statically configured VLANs on a device that is running GVRP.

Dynamic Core and Fixed Edge

In this configuration, the edge devices advertise their statically configured VLANs to the core device. The core device does not have any statically configured VLANs but learns the VLANs from the edge devices.

Enter the following commands on the core device:

```
BigIron> enable
BigIron# configure terminal
```

```
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable all
```

These commands globally enable GVRP support and enable the protocol on all ports.

Enter the following commands on edge device A:

```
BigIron> enable
BigIron# configure terminal
BigIron(config)# vlan 20
BigIron(config-vlan-20)# untag ethernet 2/1
BigIron(config-vlan-20)# tag ethernet 4/24
BigIron(config-vlan-20)# vlan 40
BigIron(config-vlan-40)# untag ethernet 2/1
BigIron(config-vlan-40)# tag ethernet 4/24
BigIron(config-vlan-40)# exit
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable ethernet 4/24
BigIron(config-gvrp)# block-learning ethernet 4/24
```

These commands statically configure two port-based VLANs, enable GVRP on port 4/24, and block GVRP learning on the port. The device will advertise the VLANs but will not learn VLANs from other devices.

Enter the following commands on edge device B:

```
BigIron> enable
BigIron# configure terminal
BigIron(config)# vlan 20
BigIron(config-vlan-20)# untag ethernet 2/24
BigIron(config-vlan-20)# tag ethernet 4/1
BigIron(config-vlan-20)# vlan 30
BigIron(config-vlan-30)# untag ethernet 4/24
BigIron(config-vlan-30)# tag ethernet 4/1
BigIron(config-vlan-30)# exit
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable ethernet 4/1
BigIron(config-gvrp)# block-learning ethernet 4/1
```

Enter the following commands on edge device C:

```
BigIron> enable
BigIron# configure terminal
BigIron(config)# vlan 30
BigIron(config-vlan-30)# untag ethernet 2/24
BigIron(config-vlan-30)# tag ethernet 4/1
BigIron(config-vlan-20)# vlan 40
BigIron(config-vlan-40)# untag ethernet 4/24
BigIron(config-vlan-40)# tag ethernet 4/1
BigIron(config-vlan-40)# exit
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable ethernet 4/1
BigIron(config-gvrp)# block-learning ethernet 4/1
```

Dynamic Core and Dynamic Edge

In this configuration, the core and edge devices have no statically configured VLANs and are enabled to learn and advertise VLANs. The edge and core devices learn the VLANs configured on the devices in the edge clouds. To enable GVRP on all the ports, enter the following command on each edge device **and** on the core device.

```
BigIron> enable
BigIron# configure terminal
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable all
```

Fixed Core and Dynamic Edge

In this configuration, GVRP learning is enabled on the edge devices. The VLANs on the core device are statically configured, and the core device is enabled to advertise its VLANs but not to learn VLANs. The edge devices learn the VLANs from the core.

Enter the following commands on the core device:

```
BigIron> enable
BigIron# configure terminal
BigIron(config)# vlan 20
BigIron(config-vlan-20)# tag ethernet 1/24
BigIron(config-vlan-20)# tag ethernet 6/24
BigIron(config-vlan-20)# vlan 30
BigIron(config-vlan-30)# tag ethernet 6/24
BigIron(config-vlan-30)# tag ethernet 8/17
BigIron(config-vlan-30)# vlan 40
BigIron(config-vlan-40)# tag ethernet 1/5
BigIron(config-vlan-40)# tag ethernet 8/17
BigIron(config-vlan-40)# vlan 50
BigIron(config-vlan-50)# untag ethernet 6/1
BigIron(config-vlan-50)# tag ethernet 1/11
BigIron(config-vlan-50)# exit
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable ethernet 1/24 ethernet 6/24 ethernet 8/17
BigIron(config-gvrp)# block-learning ethernet 1/24 ethernet 6/24 ethernet 8/17
```

These VLAN commands configure VLANs 20, 30, 40, and 50. The GVRP commands enable the protocol on the ports that are connected to the edge devices, and disable VLAN learning on those ports. All the VLANs are advertised by GVRP.

Enter the following commands on edge devices A, B, and C:

```
BigIron> enable
BigIron# configure terminal
BigIron(config)# gvrp-enable
BigIron(config-gvrp)# enable all
BigIron(config-gvrp)# block-applicant all
```

Fixed Core and Fixed Edge

The VLANs are statically configured on the core and edge devices. On each edge device, VLAN advertising is enabled but learning is disabled. GVRP is not configured on the core device. This configuration enables the devices in the edge clouds to learn the VLANs configured on the edge devices.

This configuration does not use any GVRP configuration on the core device.

The configuration on the edge device is the same as in "Dynamic Core and Fixed Edge" on page 25-17.

Chapter 26

Configuring Base Layer 3

This chapter describes how to configure static IP and RIP in the Base Layer 3 software image. The Layer 2 with Base Layer 3 software image contains all the system-level features in the Layer 2 images, along with the following:

- Static IP routes
- RIPv1 and RIPv2
- Routing between directly connected subnets
- RIP advertisements of the directly connected subnets

NOTE: These images provide static RIP support. The device does not learn RIP routes from other Layer 3 devices. However, the device does advertise directly connected routes. Foundry Networks recommends that you deploy these devices only at the edge of your network, since incoming traffic can learn directly-connected routes advertised by the Foundry device, but outgoing traffic to other devices must use statically configured or default routes.

NOTE: The Base Layer 3 images do not support IP multicasting, OSPF, BGP4, IPX, or AppleTalk.

The procedures in this chapter describe how to perform the following tasks:

- Add a static IP route.
- Add a static entry to the ARP table.
- Configure RIP.

Adding a Static IP Route

To add a static IP route, enter a command such as the following at the global CONFIG level of the CLI:

```
FastIron(config)# ip route 209.157.2.0 255.255.255.0 192.168.2.1
```

This command adds a static IP route to the 209.157.2.x/24 subnet.

Syntax: [no] ip route <dest-ip-addr> <dest-mask> <next-hop-ip-addr> [<metric>]

or

Syntax: [no] ip route <dest-ip-addr>/<mask-bits> <next-hop-ip-addr> [<metric>]

The <dest-ip-addr> is the route's destination. The <dest-mask> is the network mask for the route's destination IP address. Alternatively, you can specify the network mask information by entering a forward slash followed by the number of bits in the network mask. For example, you can enter 192.0.0.0 255.255.255.0 as 192.0.0.0/24. To

configure a default route, enter 0.0.0.0 for <dest-ip-addr> and 0.0.0.0 for <dest-mask> (or 0 for the <mask-bits> if you specify the address in CIDR format). Specify the IP address of the default gateway using the <next-hop-ip-addr> parameter.

The <next-hop-ip-addr> is the IP address of the next-hop router (gateway) for the route.

The <metric> parameter specifies the cost of the route and can be a number from 1 – 16. The default is 1. The metric is used by RIP. If you do not enable RIP, the metric is not used.

NOTE: You cannot specify **null0** or another interface as the next hop in the Base Layer 3 image.

Adding a Static ARP Entry

Static entries are useful in cases where you want to pre-configure an entry for a device that is not connected to the FastIron 4802, or you want to prevent a particular entry from aging out. The software removes a dynamic entry from the ARP cache if the ARP aging interval expires before the entry is refreshed. Static entries do not age out, regardless of whether the FastIron 4802 receives an ARP request from the device that has the entry's address. The software places a static ARP entry into the ARP cache as soon as you create the entry.

To add a static ARP entry, enter a command such as the following at the global CONFIG level of the CLI:

```
FastIron(config)# arp 1 209.157.22.3 aaaa.bbbb.cccc ethernet 3
```

This command adds a static ARP entry that maps IP address 209.157.22.3 to MAC address aaaa.bbbb.cccc. The entry is for a MAC address connected to FastIron 4802 port 3.

Syntax: [no] arp <num> <ip-addr> <mac-addr> ethernet <portnum>

The <num> parameter specifies the entry number. You can specify a number from 1 up to the maximum number of static entries allowed on the device. You can allocate more memory to increase this amount. To do so, enter the **system-max ip-static-arp <num>** command at the global CONFIG level of the CLI.

The <ip-addr> command specifies the IP address of the device that has the MAC address of the entry.

The <mac-addr> parameter specifies the MAC address of the entry.

The **ethernet <portnum>** command specifies the port number attached to the device that has the MAC address of the entry.

NOTE: The **clear arp** command clears learned ARP entries but does not remove any static ARP entries.

Configuring RIP

RIP is disabled by default. If you want the FastIron 4802 to use RIP you must enable the protocol globally, then enable RIP on individual ports. When you enable RIP on a port, you also must specify the version (version 1 only, version 2 only, or version 1 compatible with version 2).

Optionally, you also can set or change the following parameters:

- Route redistribution – You can enable the software to redistribute static routes from the IP route table into RIP. Redistribution is disabled by default.
- Learning of default routes – The default is disabled.
- Loop prevention (split horizon or poison reverse) – The default is poison reverse.

Enabling RIP

RIP is disabled by default. To enable it, use the following CLI method. You must enable the protocol both globally and on the ports on which you want to use RIP.

To enable RIP globally, enter the following command:

```
FastIron(config)# router rip
```

Syntax: [no] router rip

To enable RIP on a port and specify the RIP version, enter commands such as the following:

```
FastIron(config-rip-router)# interface ethernet 1
FastIron(config-if-1)# ip rip v1-only
```

This command changes the CLI to the configuration level for port 1 and enables RIP version 1 on the interface. You must specify the version.

Syntax: interface ethernet <portnum>

Syntax: [no] ip rip v1-only | v1-compatible-v2 | v2-only

Enabling Redistribution of IP Static Routes into RIP

By default, the software does not redistribute the IP static routes in the route table into RIP. To configure redistribution, perform the following tasks:

- Configure redistribution filters (optional). You can configure filters to permit or deny redistribution for a route based on the route's metric. You also can configure a filter to change the metric. You can configure up to 64 redistribution filters. The software uses the filters in ascending numerical order and immediately takes the action specified by the filter. Thus, if filter 1 denies redistribution of a given route, the software does not redistribute the route, regardless of whether a filter with a higher ID permits redistribution of that route.

NOTE: The default redistribution action is permit, even after you configure and apply a permit or deny filter. To deny redistribution of specific routes, you must configure a deny filter.

NOTE: The option to set the metric is not applicable to static routes.

- Enable redistribution.

NOTE: If you plan to configure redistribution filters, do not enable redistribution until you have configured the filters.

When you enable redistribution, all IP static routes are redistributed by default. If you want to deny certain routes from being redistributed into RIP, configure deny filters for those routes before you enable redistribution. You can configure up to 64 RIP redistribution filters. They are applied in ascending numerical order.

NOTE: The default redistribution action is still permit, even after you configure and apply redistribution filters to the port. If you want to tightly control redistribution, apply a filter to deny all routes as the last filter (filter ID 64), then apply filters with lower filter IDs to allow specific routes.

To configure a redistribution filter, enter a command such as the following:

```
FastIron(config-rip-router)# deny redistribute 1 static address 207.92.0.0
255.255.0.0
```

This command denies redistribution of all 207.92.x.x IP static routes.

Syntax: [no] permit | deny redistribute <filter-num> static address <ip-addr> <ip-mask>
[match-metric <value> | set-metric <value>]

The <filter-num> specifies the redistribution filter ID. Specify a number from 1 – 64. The software uses the filters in ascending numerical order. Thus, if filter 1 denies a route from being redistributed, the software does not redistribute that route even if a filter with a higher ID permits redistribution of the route.

The **address** <ip-addr> <ip-mask> parameters apply redistribution to the specified network and subnet address. Use 0 to specify “any”. For example, “207.92.0.0 255.255.0.0” means “any 207.92.x.x subnet”. However, to specify any subnet (all subnets match the filter), enter “address 255.255.255.255 255.255.255.255”.

The **match-metric** <value> parameter applies redistribution to those routes with a specific metric value; possible values are from 1 – 15.

The **set-metric** <value> parameter sets the RIP metric value that will be applied to the routes imported into RIP.

NOTE: The **set-metric** parameter does not apply to static routes.

The following command denies redistribution of a 207.92.x.x IP static route only if the route's metric is 5.

```
FastIron(config-rip-router)# deny redistribute 2 static address 207.92.0.0
255.255.0.0 match-metric 5
```

The following commands deny redistribution of all routes except routes for 10.10.10.x and 20.20.20.x:

```
FastIron(config-rip-router)# deny redistribute 64 static address 255.255.255.255
255.255.255.255
FastIron(config-rip-router)# permit redistribute 1 static address 10.10.10.0
255.255.255.0
FastIron(config-rip-router)# permit redistribute 2 static address 20.20.20.0
255.255.255.0
```

Enabling Redistribution

After you configure redistribution parameters, you need to enable redistribution.

To enable RIP redistribution, enter the following command:

```
FastIron(config-rip-router)# redistribution
```

Syntax: [no] redistribution

Enabling Learning of Default Routes

By default, the software does not learn RIP default routes.

To enable learning of default RIP routes, enter commands such as the following:

```
FastIron(config)# interface ethernet 1
FastIron(config-if-1)# ip rip learn-default
```

Syntax: interface ethernet <portnum>

Syntax: [no] ip rip learn-default

Changing the Route Loop Prevention Method

RIP can use the following methods to prevent routing loops:

- Split horizon – The FastIron 4802 does not advertise a route on the same interface as the one on which the FastIron 4802 learned the route.
- Poison reverse – The FastIron 4802 assigns a cost of 16 (“infinite” or “unreachable”) to a route before advertising it on the same interface as the one on which the FastIron 4802 learned the route. This is the default.

NOTE: These methods are in addition to RIP's maximum valid route cost of 15.

To enable split horizon, enter commands such as the following:

```
FastIron(config)# interface ethernet 1
FastIron(config-if-1)# no ip rip poison-reverse
```


Syntax: [no] ip rip poison-reverse

Additional Features

For information about the other IP configuration commands in the Layer 2 with Base Layer 3 image, see the "Configuring IP" chapter of the *Foundry Enterprise Configuration and Management Guide*.

Chapter 27

Enabling the Foundry Discovery Protocol (FDP) and Reading Cisco Discovery Protocol (CDP) Packets

Using FDP

The Foundry Discovery Protocol (FDP) enables Foundry devices to advertise themselves to other Foundry devices on the network. When you enable FDP on a Foundry device, the device periodically advertises information including the following:

- Hostname (device ID)
- Product platform and capability
- Software version
- VLAN and Layer 3 protocol address information for the port sending the update. IP, IPX, and AppleTalk Layer 3 information is supported.

A Foundry device running FDP sends FDP updates on Layer 2 to MAC address 01-E0-52-CC-CC-CC. Other Foundry devices listening on that address receive the updates and can display the information in the updates. Foundry devices can send and receive FDP updates on Ethernet, POS, and ATM interfaces.

NOTE: On POS and ATM, FDP is supported for Layer 2 only.

FDP is disabled by default.

NOTE: If FDP is not enabled on a Foundry device that receives an FDP update or the device is running a software release that does not support FDP, the update passes through the device at Layer 2.

Configuring FDP

The following sections describe how to enable FDP and how to change the FDP update and hold timers.

Enabling FDP Globally

To enable a Foundry device to globally send FDP packets, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# fdp run
```

Syntax: [no] fdp run

The feature is disabled by default.

Enabling FDP at the Interface Level

Starting in software release 07.6.02, you can enable FDP at the interface level by entering commands such as the following:

```
BigIron(config)# int e 2/1
BigIron(config-if-2/1)# fdp enable
```

For an ATM interface, you can enable or disable FDP at the sub-interface level by entering commands such as the following:

```
BigIron(config)# int atm 2/1.1
BigIron(config-subif-2/1.1)# fdp enable
```

Syntax: [no] fdp enable

By default, the feature is enabled on an interface once FDP is enabled on the device.

Changing the FDP Update Timer

By default, a Foundry device enabled for FDP sends an FDP update every 60 seconds. You can change the update timer to a value from 5 – 900 seconds.

To change the FDP update timer, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# fdp timer 120
```

Syntax: [no] fdp timer <secs>

The <secs> parameter specifies the number of seconds between updates and can be from 5 – 900 seconds. The default is 60 seconds.

Changing the FDP Hold Time

By default, a Foundry device that receives an FDP update holds the information until one of the following events occurs:

- The device receives a new update.
- 180 seconds have passed since receipt of the last update. This is the hold time.

Once either of these events occurs, the device discards the update.

To change the FDP hold time, enter a command such as the following at the global CONFIG level of the CLI:

```
BigIron(config)# fdp holdtime 360
```

Syntax: [no] fdp holdtime <secs>

The <secs> parameter specifies the number of seconds a Foundry device that receives an FDP update can hold the update before discarding it. You can specify from 10 – 255 seconds. The default is 180 seconds.

Displaying FDP Information

You can display the following FDP information:

- FDP entries for Foundry neighbors
- Individual FDP entries
- FDP information for an interface on the device you are managing
- FDP packet statistics

NOTE: If the Foundry device has intercepted CDP updates, then the CDP information is also displayed.

Displaying Neighbor Information

To display a summary list of all the Foundry neighbors that have sent FDP updates to this Foundry device, enter the following command:

```
BigIronA# show fdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a CDP device
```

```
Device ID      Local Int      Holdtm Capability Platform      Port ID
-----
BigIronB      Eth 2/9       178   Router      BigIron Rou  Eth 2/9
```

Syntax: show fdp neighbor [ethernet | pos | atm <portnum>] [detail]

The **ethernet | pos | atm <portnum>** parameter lists the information only for updates received on the specified interface.

The **detail** parameter lists detailed information for each device.

The **show fdp neighbor** command, without optional parameters, displays the following information.

Table 27.1: Summary FDP and CDP Neighbor Information

This Line...	Displays...
Device ID	The hostname of the neighbor.
Local Int	The interface on which this Foundry device received an FDP or CDP update for the neighbor.
Holdtm	The maximum number of seconds this device can keep the information received in the update before discarding it.
Capability	The role the neighbor is capable of playing in the network.
Platform	The product platform of the neighbor.
Port ID	The interface through which the neighbor sent the update.

To display detailed information, enter the following command:

```
BigIronA# show fdp neighbor detail
Device ID: BigIronB configured as default VLAN1, tag-type8100
Entry address(es):
Platform: BigIron Router, Capabilities: Router
Interface: Eth 2/9
Port ID (outgoing port): Eth 2/9 is TAGGED in following VLAN(s):
 9 10 11
Holdtime : 176 seconds
Version :
Foundry Networks, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1
```

The **show fdp neighbor detail** command displays the following information.

Table 27.2: Detailed FDP and CDP Neighbor Information

This Line...	Displays...
Device ID	The hostname of the neighbor. In addition, this line lists the VLAN memberships and other VLAN information for the neighbor port that sent the update to this device.
Entry address(es)	The Layer 3 protocol addresses configured on the neighbor port that sent the update to this device. If the neighbor is a Layer 2 Switch, this field lists the management IP address.
Platform	The product platform of the neighbor.
Capabilities	The role the neighbor is capable of playing in the network.
Interface	The interface on which this Foundry device received an FDP or CDP update for the neighbor.
Port ID	The interface through which the neighbor sent the update.
Holdtime	The maximum number of seconds this device can keep the information received in the update before discarding it.
Version	The software version running on the neighbor.

Displaying FDP Entries

To display the detailed neighbor information for a specific device, enter a command such as the following:

```
BigIronA# show fdp entry BigIronB
Device ID: BigIronB configured as default VLAN1, tag-type8100
Entry address(es):
Platform: BigIron Router, Capabilities: Router
Interface: Eth 2/9
Port ID (outgoing port): Eth 2/9 is TAGGED in following VLAN(s):
 9 10 11
Holdtime : 176 seconds
Version :
Foundry Networks, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1
```

Syntax: show fdp entry * | <device-id>

The * | <device-id> parameter specifies the device ID. If you enter *, the detailed updates for all neighbor devices are displayed. If you enter a specific device ID, the update for that device is displayed. For information about the display, see Table 27.2 on page 27-4.

Displaying FDP Information for an Interface

To display FDP information for an interface, enter a command such as the following:

```
BigIronA# show fdp interface ethernet 2/3
FastEthernet2/3 is up, line protocol is up
  Encapsulation ethernet
  Sending FDP packets every 5 seconds
  Holdtime is 180 seconds
```

This example shows information for Ethernet port 2/3. The port sends FDP updates every 5 seconds. Neighbors that receive the updates can hold them for up to 180 seconds before discarding them.

Syntax: show fdp interface [ethernet | pos | atm <portnum>]

The **ethernet | pos | atm <portnum>** parameter lists the information only for the specified interface.

Displaying FDP and CDP Statistics

To display FDP and CDP packet statistics, enter the following command:

```
BigIronA# show fdp traffic
CDP/FDP counters:
  Total packets output: 6, Input: 5
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
  Internal errors: 0
```

Syntax: show fdp traffic

Clearing FDP and CDP Information

You can clear the following FDP and CDP information:

- Information received in FDP and CDP updates
- FDP and CDP statistics

The same commands clear information for both FDP and CDP.

Clearing FDP and CDP Neighbor Information

To clear the information received in FDP and CDP updates from neighboring devices, enter the following command:

```
BigIron# clear fdp table
```

Syntax: clear fdp table

NOTE: This command clears all the updates for FDP and CDP.

Clearing FDP and CDP Statistics

To clear FDP and CDP statistics, enter the following command:

```
BigIron# clear fdp counters
```

Syntax: clear fdp counters

Reading CDP Packets

Cisco Discovery Protocol (CDP) packets are used by Cisco devices to advertise themselves to other Cisco devices. By default, Foundry devices forward these packets without examining their contents. You can configure

a Foundry device to intercept and display the contents of CDP packets. This feature is useful for learning device and interface information for Cisco devices in the network.

Foundry software release 07.5.xx supports intercepting and interpreting CDP version 1 packets. Software release 07.6.01 extends CDP support to version 2 packets. In 07.6.01 and later, when you enable CDP support, support for both CDP versions is enabled.

NOTE: The Foundry device can interpret only the information fields that are common to both CDP version 1 and CDP version 2.

NOTE: When you enable interception of CDP packets, the Foundry device drops the packets. As a result, Cisco devices will no longer receive the packets.

Enabling Interception of CDP Packets Globally

To enable the Foundry device to intercept and display CDP packets, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# cdp run
```

Syntax: [no] cdp run

The feature is disabled by default.

Enabling Interception of CDP Packets on an Interface

Starting with software release 07.6.02, you can disable and enable CDP at the interface level.

You can enter commands such as the following:

```
BigIron(config)# int e 2/1
BigIron(config-if-2/1)# cdp enable
```

For an ATM interface, you can enable or disable CDP at the sub-interface level by entering commands such as the following:

```
BigIron(config)# int atm 2/1.1
BigIron(config-subif-2/1.1)# cdp enable
```

Syntax: [no] cdp enable

By default, the feature is enabled on an interface once CDP is enabled on the device.

Displaying CDP Information

You can display the following CDP information:

- Cisco neighbors
- CDP entries for all Cisco neighbors or a specific neighbor
- CDP packet statistics

Displaying Neighbors

To display the Cisco neighbors the Foundry device has learned from CDP packets, enter the following command:

```
BigIron# show fdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a Cisco device

   Device ID      Local Int    Holdtm Capability Platform    Port ID
-----
(*)Router        Eth 1/1     124    R          cisco RSP4
FastEthernet5/0/0
```

Syntax: show fdp neighbors [detail | ethernet <portnum>]

To display detailed information for the neighbors, enter the following command:

```
BigIron# show fdp neighbors detail
Device ID: Router
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 150 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

To display information about a neighbor attached to a specific port, enter a command such as the following:

```
BigIron# show fdp neighbors ethernet 1/1
Device ID: Router
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 127 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Displaying CDP Entries

To display CDP entries for all neighbors, enter the following command:

```
BigIron# show fdp entry *
Device ID: Router
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 124 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Syntax: show fdp entry * | <device-id>

To display CDP entries for a specific device, specify the device ID. Here is an example.

```
BigIron# show fdp entry Router1
Device ID: Router1
Entry address(es):
  IP address: 207.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 156 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1,  RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

Displaying CDP Statistics

To display CDP packet statistics, enter the following command:

```
BigIron# show fdp traffic
CDP counters:
  Total packets output: 0, Input: 3
  Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
  No memory: 0, Invalid packet: 0, Fragmented: 0
```

Syntax: show fdp traffic

Clearing CDP Information

You can clear the following CDP information:

- Cisco Neighbor information
- CDP statistics

To clear the Cisco neighbor information, enter the following command:

```
BigIron# clear fdp table
```

Syntax: clear fdp table

To clear CDP statistics, enter the following command:

```
BigIron# clear fdp counters
```

Syntax: clear fdp counters

Appendix A

Using Syslog

This appendix describes how to display Syslog messages and how to configure the Syslog facility, and lists the Syslog messages that a Foundry device can display during standard operation.

NOTE: This appendix does not list Syslog messages that can be displayed when a debug option is enabled. For information about Syslog messages that are displayed by a debug option, see the *Foundry Diagnostic Guide*.

NOTE: For information about ServerIron Syslog messages, see the “show logging” section in the “Show Commands” chapter of the *Foundry ServerIron Command Line Interface Reference*.

Overview

A Foundry device’s software can write syslog messages to provide information at the following severity levels:

- Emergencies
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

The device writes the messages to a local buffer. In software release 07.6.02 and earlier, the local buffer can hold up to 100 entries. Beginning with software release 07.6.03, the buffer can hold up to 1000 entries.

You also can specify the IP address or host name of up to six Syslog servers. When you specify a Syslog server, the Foundry device writes the messages both to the system log and to the Syslog server.

Using a Syslog server ensures that the messages remain available even after a system reload. The Foundry device’s local Syslog buffer is cleared during a system reload or reboot, but the Syslog messages sent to the Syslog server remain on the server.

The Syslog service on a Syslog server receives logging messages from applications on the local host or from devices such as a Layer 2 Switch or Layer 3 Switch. Syslog adds a time stamp to each received message and

directs messages to a log file. Most Unix workstations come with Syslog configured. Some third party vendor products also provide Syslog running on NT.

Syslog uses UDP port 514 and each Syslog message thus is sent with destination port 514. Each Syslog message is one line with Syslog message format. The message is embedded in the text portion of the Syslog format. There are several subfields in the format. Keywords are used to identify each subfield, and commas are delimiters. The subfield order is insensitive except that the text subfield should be the last field in the message. All the subfields are optional.

Displaying Syslog Messages

To display the Syslog messages in the device's local buffer, enter the following command at any level of the CLI:

```
BigIron> show logging
```

```
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
```

```
Static Log Buffer:
```

```
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
```

```
Dynamic Log Buffer (50 entries):
```

```
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
```

```
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
```

```
Dec 15 18:45:15:I:Warm start
```

For information about the Syslog configuration information, time stamps, and dynamic and static buffers, see "Displaying the Syslog Configuration" on page A-3.

Enabling Real-Time Display of Syslog Messages

By default, to view Syslog messages generated by a Foundry device, you need to display the Syslog buffer or the log on a Syslog server used by the Foundry device.

You can enable real-time display of Syslog messages on the management console. When you enable this feature, the software displays a Syslog message on the management console when the message is generated.

When you enable the feature, the software displays Syslog messages on the serial console when they occur. However, to enable display of real-time Syslog messages in Telnet or SSH sessions, you also must enable display within the individual sessions.

USING THE CLI

To enable real-time display of Syslog messages, enter the following command at the global CONFIG level of the CLI:

```
BigIron(config)# logging console
```

Syntax: [no] logging console

This command enables the real-time display of Syslog messages on the serial console. You can enter this command from the serial console or a Telnet or SSH session.

To also enable the real-time display for a Telnet or SSH session, enter the following command from the Privileged EXEC level of the session:

```
telnet@BigIron# terminal monitor
Syslog trace was turned ON
```

Syntax: terminal monitor

Notice that the CLI displays a message to indicate the status change for the feature. To disable the feature in the management session, enter the **terminal monitor** command again. The command toggles the feature on and off.

```
telnet@BigIron# terminal monitor
Syslog trace was turned OFF
```

Here is an example of how the Syslog messages are displayed:

```
telnet@BigIron# terminal monitor
Syslog trace was turned ON
SYSLOG: <9>BigIron, Power supply 2, power supply on left connector, failed

SYSLOG: <14>BigIron, Interface ethernet 1/6, state down

SYSLOG: <14>BigIron, Interface ethernet 1/2, state up
```

Configuring the Syslog Service

The procedures in this section describe how to perform the following Syslog configuration tasks:

- Specify a Syslog server. You can configure the Foundry device to use up to six Syslog servers. (Use of a Syslog server is optional. The system can hold up to 100 Syslog messages in an internal buffer.)
- Change the level of messages the system logs.
- Change the number of messages the local Syslog buffer can hold.
- Display the Syslog configuration.
- Clear the local Syslog buffer.

Logging is enabled by default, with the following settings:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- By default, up to 50 messages are retained in the local Syslog buffer. This can be changed.
- No Syslog server is specified.

Displaying the Syslog Configuration

To display the Syslog parameters currently in effect on a Foundry device, enter the following command from any level of the CLI:

```
BigIron> show logging
```

```

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
```

```
Static Log Buffer:
```

```
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
```

```
Dynamic Log Buffer (50 entries):
```

```
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
```

```
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
```

```
Dec 15 18:45:15:I:Warm start
```

Syntax: show logging

The Syslog display shows the following configuration information, in the rows above the log entries themselves.

Table A.1: CLI Display of Syslog Buffer Configuration

This Field...	Displays...
Syslog logging	The state (enabled or disabled) of the Syslog buffer.
messages dropped	The number of Syslog messages dropped due to user-configured filters. By default, the software logs messages for all Syslog levels. You can disable individual Syslog levels, in which case the software filters out messages at those levels. See “Disabling Logging of a Message Level” on page A-10. Each time the software filters out a Syslog message, this counter is incremented.
flushes	The number of times the Syslog buffer has been cleared by the clear logging command or equivalent Web management interface option. See “Clearing the Syslog Messages from the Local Buffer” on page A-12.
overruns	The number of times the dynamic log buffer has filled up and been cleared to hold new entries. For example, if the buffer is set for 100 entries, the 101st entry causes an overrun. After that, the 201st entry causes a second overrun.
level	The message levels that are enabled. Each letter represents a message type and is identified by the key (level code) below the value. If you disable logging of a message level, the code for that level is not listed.
messages logged	The total number of messages that have been logged since the software was loaded.
level code	The message levels represented by the one-letter codes.

Static and Dynamic Buffers

The software provides two separate buffers:

- Static – logs power supply failures, fan failures, and temperature warning or shutdown messages
- Dynamic – logs all other message types

In the static log, new messages replace older ones, so only the most recent message is displayed. For example, only the most recent temperature warning message will be present in the log. If multiple temperature warning messages are sent to the log, the latest one replaces the previous one. The static buffer is not configurable.

The message types that appear in the static buffer do not appear in the dynamic buffer. The dynamic buffer contains up to the maximum number of messages configured for the buffer (50 by default), then begins removing the oldest messages (at the bottom of the log) to make room for new ones.

The static and dynamic buffers are both displayed when you display the log.

```
BigIron(config)# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
```

Static Log Buffer:

```
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
```

Dynamic Log Buffer (50 entries):

```
Dec 15 18:46:17:I:Interface ethernet 1/4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Notice that the static buffer contains two separate messages for fan failures. Each message of each type has its own buffer. Thus, if you replace fan 1 but for some reason that fan also fails, the software replaces the first message about the failure of fan 1 with the newer message. The software does not overwrite the message for fan 2, unless the software sends a newer message for fan 2.

When you clear log entries, you can selectively clear the static or dynamic buffer, or you can clear both. For example, to clear only the dynamic buffer, enter the following command at the Privileged EXEC level:

```
BigIron# clear logging dynamic-buffer
```

Syntax: clear logging [dynamic-buffer | static-buffer]

You can specify **dynamic-buffer** to clear the dynamic buffer or **static-buffer** to clear the static buffer. If you do not specify a buffer, both buffers are cleared.

Time Stamps

The contents of the time stamp differ depending on whether you have set the time and date on the onboard system clock.

- If you have set the time and date on the onboard system clock, the date and time are shown in the following format:

mm dd hh:mm:ss

where:

- *mm* – abbreviation for the name of the month
- *dd* – day
- *hh* – hours
- *mm* – minutes
- *ss* – seconds

For example, "Oct 15 17:38:03" means October 15 at 5:38 PM and 3 seconds.

- If you have not set the time and date on the onboard system clock, the time stamp shows the amount of time that has passed since the device was booted, in the following format:

<num>d<num>h<num>m<num>s

where:

- <num>d – day
- <num>h – hours
- <num>m – minutes
- <num>s – seconds

For example, “188d1h01m00s” means the device had been running for 188 days, 11 hours, one minute, and zero seconds when the Syslog entry with this time stamp was generated.

Example of Syslog Messages on a Device Whose Onboard Clock Is Set

The example shows the format of messages on a device whose onboard system clock has been set. Each time stamp shows the month, the day, and the time of the system clock when the message was generated. For example, the system time when the most recent message (the one at the top) was generated was October 15 at 5:38 PM and 3 seconds.

```
BigIron(config)# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 38 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed

Dynamic Log Buffer (50 entries):
Oct 15 17:38:03:warning:list 101 denied tcp 209.157.22.191(0)(Ethernet 4/18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)

Oct 15 07:03:30:warning:list 101 denied tcp 209.157.22.26(0)(Ethernet 4/18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)

Oct 15 06:58:30:warning:list 101 denied tcp 209.157.22.198(0)(Ethernet 4/18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
```

Example of Syslog Messages on a Device Whose Onboard Clock Is Not Set

The example shows the format of messages on a device whose onboard system clock is not set. Each time stamp shows the amount of time the device had been running when the message was generated. For example, the most

recent message, at the top of the list of messages, was generated when the device had been running for 21 days, seven hours, two minutes, and 40 seconds.

```
BigIron(config)# show log

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 38 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:

Dynamic Log Buffer (50 entries):
21d07h02m40s:warning:list 101 denied tcp 209.157.22.191(0)(Ethernet 4/18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)

19d07h03m30s:warning:list 101 denied tcp 209.157.22.26(0)(Ethernet 4/18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)

17d06h58m30s:warning:list 101 denied tcp 209.157.22.198(0)(Ethernet 4/18
0010.5alf.77ed) -> 198.99.4.69(http), 1 event(s)
```

Displaying and Configuring Syslog Buffer Parameters Using the Web Management Interface

To configure Syslog parameters using the Web management interface, use the following procedure:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Select Management from the System configuration sheet to display the Management panel.
3. Select the System Log link to display the following panel.

System Log	
Logging:	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Buffer Size:	<input type="text" value="50"/>
Facility:	<input type="text" value="user"/>
Accept Severity:	<input checked="" type="checkbox"/> alert <input checked="" type="checkbox"/> critical <input checked="" type="checkbox"/> debugging <input checked="" type="checkbox"/> emergency <input checked="" type="checkbox"/> error <input checked="" type="checkbox"/> informational <input checked="" type="checkbox"/> notification <input checked="" type="checkbox"/> warning

[\[Show Log Server\]](#)

[\[Home\]](#) [\[Site Map\]](#) [\[Logout\]](#) [\[Save\]](#) [\[Frame Enable\]](#) [\[Disable\]](#) [\[TELNET\]](#)

4. Select Disable or Enable next to Logging to disable or enable the Syslog service on the device. The service is enabled by default.
5. Optionally change the number of entries the local Syslog buffer can hold. The buffer size can be from 1 – 100. The default is 50.

NOTE: A change in the buffer size takes effect only after you restart the system. The buffer size does not affect how many entries the device can log on a Syslog server. The number of entries the device can log on the server depends on the server's configuration.

6. Select the messages facility. The default is User. For a list of values, display the pulldown menu.
7. Select the message levels you want the device to log. All the levels are logged by default.
8. Click Apply to save the changes to the device's running-config file.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
10. To view a list of the Syslog servers that have been defined, click the Show Log Server link under the Apply and Reset buttons to display the Log Server panel.

Figure A.1 List of Log Servers

Log Server

IP Address	UDP port	
1.1.1.1	12	Delete
3.3.3.4	66	Delete
IP Address	UDP port	

[Add Log Server]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

The list shows the IP Addresses and UDP Ports of the Syslog Servers.

11. To delete an entry, click on the Delete button for that entry.
12. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
13. To add a Syslog server, click on the Add Log Server link under the dialog to display the System Log Server panel.

Figure A.2 System Log Server Panel

System Log Server

Server IP Address:	<input type="text" value="3.3.3.4"/>
Server Udp Port:	<input type="text" value="66"/>

[Show Log Server]

[Home][Site Map][Logout][Save][Frame Enable][Disable][TELNET]

14. Enter the IP address of the new Syslog server, if you want the device to log messages on the Syslog server as well as in the local buffer.
15. Enter the UDP port on the server that will be used for logging messages.

16. Click on the Add button to add the server to the list. You can add up to six Syslog servers.
17. When you have finished, select the [Save](#) link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.

Disabling or Re-Enabling Syslog

Syslog is enabled by default. To disable or re-enable it, use one of the following methods.

USING THE CLI

To disable it, enter the following command at the global CONFIG level:

```
BigIron(config)# no logging on
```

Syntax: [no] logging on [<udp-port>]

The <udp-port> parameter specifies the application port used for the Syslog facility. The default is 514.

To re-enable logging, enter the following command:

```
BigIron(config)# logging on
```

This command enables local Syslog logging with the following defaults:

- Messages of all severity levels (Emergencies – Debugging) are logged.
- Up to 50 messages are retained in the local Syslog buffer.
- No Syslog server is specified.

Specifying a Syslog Server

To specify a Syslog server, use one of the following methods.

USING THE CLI

For software releases earlier than 07.7.00, enter a command such as the following:

```
BigIron(config)# logging 10.0.0.99
```

For software releases 07.7.00 and later, enter a command such as the following:

```
BigIron(config)# logging host 10.0.0.99
```

For backward compatibility, the software reads the old command syntax from the startup configuration, and converts it to the new command syntax in the running configuration.

Syntax: logging <ip-addr> | <server-name> (software releases earlier than 07.7.00)

Syntax: logging host <ip-addr> | <server-name> (software release 07.7.00 and later)

USING THE WEB MANAGEMENT INTERFACE

See the section “Displaying and Configuring Syslog Buffer Parameters Using the Web Management Interface” on page A-7.

NOTE: You can specify a server name only if you have already configured the DNS Resolver feature. See the “Configuring IP” chapter in the *Foundry Enterprise Configuration and Management Guide*.

Specifying an Additional Syslog Server

USING THE CLI

To specify an additional Syslog server, enter the **logging host** <ip-addr> command again, as in the following example. You can specify up to six Syslog servers.

For software releases earlier than 07.7.00, enter a command such as the following:

```
BigIron(config)# logging 10.0.0.99
```

For software releases 07.7.00 and later, enter a command such as the following:

```
BigIron(config)# logging host 10.0.0.99
```

For backward compatibility, the software reads the old command syntax from the startup configuration, and converts it to the new command syntax in the running configuration.

Syntax: logging <ip-addr> | <server-name> (software releases earlier than 07.7.00)

Syntax: logging host <ip-addr> | <server-name> (software release 07.7.00 and later)

[USING THE WEB MANAGEMENT INTERFACE](#)

See the section “Displaying and Configuring Syslog Buffer Parameters Using the Web Management Interface” on page A-7.

Disabling Logging of a Message Level

To change the message level, disable logging of specific message levels. You must disable the message levels on an individual basis.

[USING THE CLI](#)

For example, to disable logging of debugging and informational messages, enter the following commands:

```
BigIron(config)# no logging buffered debugging
BigIron(config)# no logging buffered informational
```

Syntax: [no] logging buffered <level> | <num-entries>

The <level> parameter can have one of the following values:

- alerts
- critical
- debugging
- emergencies
- errors
- informational
- notifications
- warnings

The commands in the example above change the log level to notification messages or higher. The software will not log informational or debugging messages. The changed message level also applies to the Syslog servers.

[USING THE WEB MANAGEMENT INTERFACE](#)

See the section “Displaying and Configuring Syslog Buffer Parameters Using the Web Management Interface” on page A-7.

Changing the Number of Entries the Local Buffer Can Hold

You also can use the **logging buffered** command to change the number of entries the local Syslog buffer can store. For example:

```
BigIron(config)# logging buffered 100
```

The default number of messages is 50. The value can be from 1 – 1000 on Layer 2 Switches and Layer 3 Switches. The change takes effect immediately and does not require you to reload the software.

[USING THE WEB MANAGEMENT INTERFACE](#)

See the section “Displaying and Configuring Syslog Buffer Parameters Using the Web Management Interface” on page A-7.

NOTE: If you decrease the size of the buffer, the software clears the buffer before placing the change into effect. If you increase the size of the buffer, the software does not clear existing entries.

Changing the Log Facility

The Syslog daemon on the Syslog server uses a facility to determine where to log the messages from the Foundry device. The default facility for messages the Foundry device sends to the Syslog server is "user". You can change the facility using the following command.

NOTE: You can specify only one facility. If you configure the Foundry device to use two Syslog servers, the device uses the same facility on both servers.

```
BigIron(config)# logging facility local0
```

Syntax: logging facility <facility-name>

The <facility-name> can be one of the following:

- kern – kernel messages
- user – random user-level messages
- mail – mail system
- daemon – system daemons
- auth – security/authorization messages
- syslog – messages generated internally by Syslog
- lpr – line printer subsystem
- news – netnews subsystem
- uucp – uucp subsystem
- sys9 – cron/at subsystem
- sys10 – reserved for system use
- sys11 – reserved for system use
- sys12 – reserved for system use
- sys13 – reserved for system use
- sys14 – reserved for system use
- cron – cron/at subsystem
- local0 – reserved for local use
- local1 – reserved for local use
- local2 – reserved for local use
- local3 – reserved for local use
- local4 – reserved for local use
- local5 – reserved for local use
- local6 – reserved for local use
- local7 – reserved for local use

USING THE WEB MANAGEMENT INTERFACE

See the section "Displaying and Configuring Syslog Buffer Parameters Using the Web Management Interface" on page A-7.

Displaying the Interface Name in Syslog Messages

By default, an interface's slot number (if applicable) and port number are displayed when you display Syslog messages. If you want to display the name of the interface instead of its number, enter the following command:

```
BigIron(config)# ip show-portname
```

This command is applied globally to all interfaces on Layer 2 Switches and Layer 3 Switches.

Syntax: [no] ip show-portname

When you display the messages in the Syslog, you see the interface name under the Dynamic Log Buffer section. The actual interface number is appended to the interface name. For example, if the interface name is "lab" and its port number is "2", you see "lab2" displayed as in the example below:

```
BigIron# show logging

Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed

Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet Lab2, state up
Dec 15 18:45:15:I:Warm start
```

Clearing the Syslog Messages from the Local Buffer

To clear the Syslog messages stored in the Foundry device's local buffer, use one of the following methods:

USING THE CLI

```
BigIron# clear logging
```

Syntax: clear logging

USING THE WEB MANAGEMENT INTERFACE

To clear Syslog messages using the Web management interface, use the following procedure:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Command in the tree view to display the command options.
3. Select the Clear link to display the Clear panel.
4. Click on the checkbox next to System Logging to place a checkmark in the box.
5. Click Apply to clear the log.

Displaying TCP/UDP Port Numbers in Syslog Messages

The command **ip show-acl-service-number** allows you to change the display of TCP/UDP application information from the TCP/UDP well-known port name to the TCP/UDP port number. For example, entering the following command causes the Foundry device to display **http** (the well-known port name) instead of **80** (the port number) in the output of **show** commands, and other commands that contain application port information. By default, Foundry devices display TCP/UDP application information in named notation.

In this release, you can display TCP/UDP port number instead of their names in syslog messages by entering the following command:

```
BigIron(config)# ip show-service-number-in-log
```

Syntax: [no] ip show-service-number-in-log

Syslog Messages

Table A.2 lists all of the Syslog messages. Note that some of the messages apply only to Layer 3 Switches. The messages are listed by message level, in the following order:

- Emergencies (none)
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

Table A.2: Foundry Syslog Messages

Message Level	Message	Explanation
Alert	MGMT and ATM (slot <num>) modules SW incompatible, MGMT <version>, ATM <version>	<p>The flash code versions on the management module and an ATM module do not match. As a result, the device has shut down the ATM module.</p> <p>The slot <num> is the slot number that contains the module.</p> <p>The MGMT <version> is the flash code version running on the management module.</p> <p>The ATM <version> is the flash code version installed on the ATM module's flash memory.</p> <p>Note: You need to interactively boot the ATM module from a TFTP server, then copy the correct flash code onto its flash memory.</p>
Alert	MGMT and POS (slot <num>) modules SW incompatible, MGMT <version>, POS <version>	<p>The flash code versions on the management module and a POS module do not match. As a result, the device has shut down the POS module.</p> <p>The slot <num> is the slot number that contains the module.</p> <p>The MGMT <version> is the flash code version running on the management module.</p> <p>The POS <version> is the flash code version installed on the POS module's flash memory.</p> <p>Note: You need to interactively boot the POS module from a TFTP server, then copy the correct flash code onto its flash memory.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Alert	Power supply <num>, <location>, failed	<p>A power supply has failed.</p> <p>The <num> is the power supply number.</p> <p>The <location> describes where the failed power supply is in the chassis. The location can be one of the following:</p> <ul style="list-style-type: none"> • In 4-slot s: <ul style="list-style-type: none"> • left side power supply • right side power supply • In 8-slot s: <ul style="list-style-type: none"> • bottom power supply • middle bottom power supply • middle top power supply • top power supply • In 15-slot s: <ul style="list-style-type: none"> • left side power supply • second from left power supply • second from right power supply • right side power supply • In Stackable devices: <ul style="list-style-type: none"> • power supply on right connector • power supply on left connector

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Alert	Fan <num>, <location>, failed	<p>A fan has failed.</p> <p>The <num> is the power supply number.</p> <p>The <location> describes where the failed power supply is in the chassis. The location can be one of the following:</p> <ul style="list-style-type: none"> • In Stackable devices: <ul style="list-style-type: none"> • fan on right connector • fan on left connector • In 4-slot s: <ul style="list-style-type: none"> • left side panel, back fan • left side panel, front fan • rear/back panel, left fan • rear/back panel, right fan • In 8-slot and 15-slot s: <ul style="list-style-type: none"> • rear/back panel, top fan • rear/back panel, bottom fan • top panel, fan
Alert	Management module at slot <slot-num> state changed from <module-state> to <module-state>.	<p>Indicates a state change in a management module.</p> <p>The <slot-num> indicates the chassis slot containing the module.</p> <p>The <module-state> can be one of the following:</p> <ul style="list-style-type: none"> • active • standby • crashed • coming-up • unknown
Alert	Temperature <degrees> C degrees, warning level <warn-degrees> C degrees, shutdown level <shutdown-degrees> C degrees	<p>Indicates an overtemperature condition on the active module.</p> <p>The <degrees> value indicates the temperature of the module.</p> <p>The <warn-degrees> value is the warning threshold temperature configured for the module.</p> <p>The <shutdown-degrees> value is the shutdown temperature configured for the module.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Alert	<num-modules> modules and 1 power supply, need more power supply!!	Indicates that the chassis needs more power supplies to run the modules in the chassis. The <num-modules> parameter indicates the number of modules in the chassis.
Alert	OSPF Memory Overflow	OSPF has run out of memory.
Alert	OSPF LSA Overflow, LSA Type = <lsa-type>	Indicates an LSA database overflow. The <lsa-type> parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following: <ul style="list-style-type: none"> • 1 – Router • 2 – Network • 3 – Summary • 4 – Summary • 5 – External
Alert	ISIS MEMORY USE EXCEEDED	IS-IS is requesting more memory than is available.
Alert	MAC Authentication failed for <mac-address> on <portnum> (Invalid User)	RADIUS authentication failed for the specified <mac-address> on the specified <portnum> because the MAC address sent to the RADIUS server was not found in the RADIUS server's users database.
Alert	MAC Authentication failed for <mac-address> on <portnum>	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, the VLAN returned in the RADIUS Access-Accept message did not refer to a valid VLAN or VLAN ID on the Foundry device. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (No VLAN Info received from RADIUS server)	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, dynamic VLAN assignment was enabled for the port, but the RADIUS Access-Accept message did not include VLAN information. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (RADIUS given VLAN does not match with TAGGED vlan)	Multi-device port authentication failed for the <mac-address> on a tagged port because the packet with this MAC address as the source was tagged with a VLAN ID different from the RADIUS-supplied VLAN ID.

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Alert	MAC Authentication failed for <mac-address> on <portnum> (RADIUS given vlan does not exist)	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, the RADIUS Access-Accept message specified a VLAN that does not exist in the Foundry device's configuration. This is treated as an authentication failure.
Alert	MAC Authentication failed for <mac-address> on <portnum> (Port is already in another radius given vlan)	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>; however, the RADIUS Access-Accept message specified a VLAN ID, although the port had previously been moved to a different RADIUS-assigned VLAN. This is treated as an authentication failure.
Critical	Authentication shut down <portnum> due to DOS attack	Denial of Service (DoS) attack protection was enabled for multi-device port authentication on the specified <portnum>, and the per-second rate of RADIUS authentication attempts for the port exceeded the configured limit. The Foundry device considers this to be a DoS attack and disables the port.
Error	No of prefixes received from BGP peer <ip-addr> exceeds maximum prefix-limit...shutdown	The Layer 3 Switch has received more than the specified maximum number of prefixes from the neighbor, and the Layer 3 Switch is therefore shutting down its BGP4 session with the neighbor.
Warning	Locked address violation at interface e<portnum>, address <mac-address>	Indicates that a port on which you have configured a lock-address filter received a packet that was dropped because the packet's source MAC address did not match an address learned by the port before the lock took effect. The e<portnum> is the port number. The <mac-address> is the MAC address that was denied by the address lock. Assuming that you configured the port to learn only the addresses that have valid access to the port, this message indicates a security violation.
Warning	NTP server <ip-addr> failed to respond	Indicates that a Simple Network Time Protocol (SNTP) server did not respond to the device's query for the current time. The <ip-addr> indicates the IP address of the SNTP server.

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Warning	Dup IP <ip-addr> detected, sent from MAC <mac-addr> interface <portnum>	<p>Indicates that the Foundry device received a packet from another device on the network with an IP address that is also configured on the Foundry device.</p> <p>The <ip-addr> is the duplicate IP address.</p> <p>The <mac-addr> is the MAC address of the device with the duplicate IP address.</p> <p>The <portnum> is the Foundry port that received the packet with the duplicate IP address. The address is the packet's source IP address.</p>
Warning	mac filter group denied packets on port <portnum> src macaddr <mac-addr>, <num> packets	<p>Indicates that a Layer 2 MAC filter group configured on a port has denied packets.</p> <p>The <portnum> is the port on which the packets were denied.</p> <p>The <mac-addr> is the source MAC address of the denied packets.</p> <p>The <num> indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.</p>
Warning	list <acl-num> denied <ip-proto> <src-ip-addr> (<src-tcp/udp-port>) (Ethernet <portnum> <mac-addr>) -> <dst-ip-addr> (<dst-tcp/udp-port>), 1 event(s)	<p>Indicates that an Access Control List (ACL) denied (dropped) packets.</p> <p>The <acl-num> indicates the ACL number. Numbers 1 – 99 indicate standard ACLs. Numbers 100 – 199 indicate extended ACLs.</p> <p>The <ip-proto> indicates the IP protocol of the denied packets.</p> <p>The <src-ip-addr> is the source IP address of the denied packets.</p> <p>The <src-tcp/udp-port> is the source TCP or UDP port, if applicable, of the denied packets.</p> <p>The <portnum> indicates the port number on which the packet was denied.</p> <p>The <mac-addr> indicates the source MAC address of the denied packets.</p> <p>The <dst-ip-addr> indicates the destination IP address of the denied packets.</p> <p>The <dst-tcp/udp-port> indicates the destination TCP or UDP port number, if applicable, of the denied packets.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Warning	rip filter list <list-num> <direction> V1 V2 denied <ip-addr>, <num> packets	<p>Indicates that a RIP route filter denied (dropped) packets.</p> <p>The <list-num> is the ID of the filter list.</p> <p>The <direction> indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following:</p> <ul style="list-style-type: none"> • in • out <p>The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2).</p> <p>The <ip-addr> indicates the network number in the denied updates.</p> <p>The <num> indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.</p>
Warning	No of prefixes received from BGP peer <ip-addr> exceeds warning limit <num>	<p>The Layer 3 Switch has received more than the allowed percentage of prefixes from the neighbor.</p> <p>The <ip-addr> is the IP address of the neighbor.</p> <p>The <num> is the number of prefixes that matches the percentage you specified. For example, if you specified a threshold of 100 prefixes and 75 percent as the warning threshold, this message is generated if the Layer 3 Switch receives a 76th prefix from the neighbor.</p>
Warning	DOT1X security violation at port <portnum>, malicious mac address detected: <mac-address>	<p>A security violation was encountered at the specified port number.</p>
Notification	Module was inserted to slot <slot-num>	<p>Indicates that a module was inserted into a chassis slot.</p> <p>The <slot-num> is the number of the chassis slot into which the module was inserted.</p>
Notification	Module was removed from slot <slot-num>	<p>Indicates that a module was removed from a chassis slot.</p> <p>The <slot-num> is the number of the chassis slot from which the module was removed.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	ACL insufficient L4 session resource, using flow based ACL instead	<p>The device does not have enough Layer 4 session entries.</p> <p>To correct this condition, allocate more memory for sessions. To allocate more memory, enter the following command at the global CONFIG level of the CLI interface:</p> <p>system-max session-limit <num></p>
Notification	ACL exceed max DMA L4 cam resource, using flow based ACL instead	<p>The port does not have enough Layer 4 CAM entries for the ACL.</p> <p>To correct this condition, allocate more Layer 4 CAM entries. To allocate more Layer 4 CAM entries, enter the following command at the CLI configuration level for the interface:</p> <p>ip access-group max-l4-cam <num></p>
Notification	ACL insufficient L4 cam resource, using flow based ACL instead	<p>The port does not have a large enough CAM partition for the ACLs. To re-partition the CAM, see the "Changing CAM Partitions" chapter in the <i>Foundry Diagnostic Guide</i>.</p>
Notification	ACL system fragment packet inspect rate <rate> exceeded	<p>The fragment rate allowed on the device has been exceeded.</p> <p>The <rate> indicates the maximum rate allowed.</p> <p>This message can occur if fragment throttling is enabled.</p>
Notification	ACL port fragment packet inspect rate <rate> exceeded on port <portnum>	<p>The fragment rate allowed on an individual interface has been exceeded.</p> <p>The <rate> indicates the maximum rate allowed.</p> <p>The <portnum> indicates the port.</p> <p>This message can occur if fragment throttling is enabled.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF interface state changed, rid <router-id>, intf addr <ip-addr>, state <ospf-state>	<p>Indicates that the state of an OSPF interface has changed.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the interface's IP address.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown
Notification	OSPF virtual intf state changed, rid <router-id>, area <area-id>, nbr <ip-addr>, state <ospf-state>	<p>Indicates that the state of an OSPF virtual routing interface has changed.</p> <p>The <router-id> is the router ID of the router the interface is on.</p> <p>The <area-id> is the area the interface is in.</p> <p>The <ip-addr> is the IP address of the OSPF neighbor.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • loopback • waiting • point-to-point • designated router • backup designated router • other designated router • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF nbr state changed, rid <router-id>, nbr addr <ip-addr>, nbr rid <nbr-router-Id>, state <ospf-state>	<p>Indicates that the state of an OSPF neighbor has changed.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the neighbor.</p> <p>The <nbr-router-id> is the router ID of the neighbor.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF virtual nbr state changed, rid <router-id>, nbr addr <ip-addr>, nbr rid <nbr-router-id>, state <ospf-state>	<p>Indicates that the state of an OSPF virtual neighbor has changed.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the neighbor.</p> <p>The <nbr-router-id> is the router ID of the neighbor.</p> <p>The <ospf-state> indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none"> • down • attempt • initializing • 2-way • exchange start • exchange • loading • full • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF intf config error, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF interface configuration error has occurred.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the Foundry device received the error packet.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF virtual intf config error, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF virtual routing interface configuration error has occurred.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the Foundry device received the error packet.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF intf authen failure, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF interface authentication failure has occurred.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the Foundry device received the authentication failure.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF virtual intf authen failure, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, error type <error-type>, pkt type <pkt-type>	<p>Indicates that an OSPF virtual routing interface authentication failure has occurred.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the Foundry device received the authentication failure.</p> <p>The <error-type> can be one of the following:</p> <ul style="list-style-type: none"> • bad version • area mismatch • unknown NBMA neighbor • unknown virtual neighbor • authentication type mismatch • authentication failure • network mask mismatch • hello interval mismatch • dead interval mismatch • option mismatch • unknown <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF intf rcvd bad pkt, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, pkt type <pkt-type>	<p>Indicates that an OSPF interface received a bad packet.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the Foundry device received the authentication failure.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown
Notification	OSPF virtual intf rcvd bad pkt, rid <router-id>, intf addr <ip-addr>, pkt src addr <src-ip-addr>, pkt type <pkt-type>	<p>Indicates that an OSPF interface received a bad packet.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <src-ip-addr> is the IP address of the interface from which the Foundry device received the authentication failure.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF intf retransmit, rid <router-id>, intf addr <ip-addr>, nbr rid <nbr-router-id>, pkt type is <pkt-type>, LSA type <lsa-type>, LSA id <lsa-id>, LSA rid <lsa-router-id>	<p>An OSPF interface on the Foundry device has retransmitted a Link State Advertisement (LSA).</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <nbr-router-id> is the router ID of the neighbor router.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF virtual intf retransmit, rid <router-id>, intf addr <ip-addr>, nbr rid <nbr-router-id>, pkt type is <pkt-type>, LSA type <lsa-type>, LSA id <lsa-id>, LSA rid <lsa-router-id>	<p>An OSPF interface on the Foundry device has retransmitted a Link State Advertisement (LSA).</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <ip-addr> is the IP address of the interface on the Foundry device.</p> <p>The <nbr-router-id> is the router ID of the neighbor router.</p> <p>The <packet-type> can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state ack • unknown <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>
Notification	OSPF originate LSA, rid <router-id>, area <area-id>, LSA type <lsa-type>, LSA id <lsa-id>, LSA router id <lsa-router-id>	<p>An OSPF interface has originated an LSA.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <area-id> is the OSPF area.</p> <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>
Notification	OSPF max age LSA, rid <router-id>, area <area-id>, LSA type <lsa-type>, LSA id <lsa-id>, LSA rid <lsa-router-id>	<p>An LSA has reached its maximum age.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <area-id> is the OSPF area.</p> <p>The <lsa-type> is the type of LSA.</p> <p>The <lsa-id> is the LSA ID.</p> <p>The <lsa-router-id> is the LSA router ID.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF LSDB overflow, rid <router-id>, limit <num>	<p>A Link State Database Overflow (LSDB) condition has occurred.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <num> is the number of LSAs.</p>
Notification	OSPF LSDB approaching overflow, rid <router-id>, limit <num>	<p>The software is close to an LSDB condition.</p> <p>The <router-id> is the router ID of the Foundry device.</p> <p>The <num> is the number of LSAs.</p>
Notification	OSPF intf rcvd bad pkt: Bad Checksum, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The device received an OSPF packet that had an invalid checksum.</p> <p>The rid <ip-addr> is Foundry device's router ID.</p> <p>The intf addr <ip-addr> is the IP address of the Foundry interface that received the packet.</p> <p>The pkt size <num> is the number of bytes in the packet.</p> <p>The checksum <num> is the checksum value for the packet.</p> <p>The pkt src addr <ip-addr> is the IP address of the neighbor that sent the packet.</p> <p>The pkt type <type> is the OSPF packet type and can be one of the following:</p> <ul style="list-style-type: none"> • hello • database description • link state request • link state update • link state acknowledgement • unknown (indicates an invalid packet type)
Notification	OSPF intf rcvd bad pkt: Bad Packet type, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The device received an OSPF packet with an invalid type.</p> <p>The parameters are the same as for the Bad Checksum message. The pkt type <type> value is "unknown", indicating that the packet type is invalid.</p>
Notification	OSPF intf rcvd bad pkt: Unable to find associated neighbor, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The neighbor IP address in the packet is not on the Foundry device's list of OSPF neighbors.</p> <p>The parameters are the same as for the Bad Checksum message.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	OSPF intf rcvd bad pkt: Invalid packet size, rid <ip-addr>, intf addr <ip-addr>, pkt size <num>, checksum <num>, pkt src addr <ip-addr>, pkt type <type>	<p>The device received an OSPF packet with an invalid packet size.</p> <p>The parameters are the same as for the Bad Checksum message.</p>
Notification	FSRP intf state changed, intf <portnum>, addr <ip-addr>, state <fsrp-state>	<p>A state change has occurred in a Foundry Standby Router Protocol (FSRP) interface.</p> <p>The <portnum> is the port.</p> <p>The <ip-addr> is the IP address of the FSRP interface.</p> <p>The <fsrp-state> can be one of the following:</p> <ul style="list-style-type: none"> • init • negotiating • standby • active • unknown
Notification	VRRP intf state changed, intf <portnum>, vrid <virtual-router-id>, state <vrrp-state>	<p>A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) interface.</p> <p>The <portnum> is the port.</p> <p>The <virtual-router-id> is the virtual router ID (VRID) configured on the interface.</p> <p>The <vrrp-state> can be one of the following:</p> <ul style="list-style-type: none"> • init • master • backup • unknown
Notification	BGP Peer <ip-addr> UP (ESTABLISHED)	<p>Indicates that a BGP4 neighbor has come up.</p> <p>The <ip-addr> is the IP address of the neighbor's BGP4 interface with the Foundry device.</p>
Notification	BGP Peer <ip-addr> DOWN (IDLE)	<p>Indicates that a BGP4 neighbor has gone down.</p> <p>The <ip-addr> is the IP address of the neighbor's BGP4 interface with the Foundry device.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	Local ICMP exceeds <burst-max> burst packets, stopping for <lockup> seconds!!	<p>The number of ICMP packets exceeds the <burst-max> threshold set by the ip icmp burst command. The Foundry device may be the victim of a Denial of Service (DoS) attack.</p> <p>All ICMP packets will be dropped for the number of seconds specified by the <lockup> value. When the lockup period expires, the packet counter is reset and measurement is restarted.</p>
Notification	Local TCP exceeds <burst-max> burst packets, stopping for <lockup> seconds!!	<p>The number of TCP SYN packets exceeds the <burst-max> threshold set by the ip tcp burst command. The Foundry device may be the victim of a TCP SYN DoS attack.</p> <p>All TCP SYN packets will be dropped for the number of seconds specified by the <lockup> value. When the lockup period expires, the packet counter is reset and measurement is restarted.</p>
Notification	Transit ICMP in interface <portnum> exceeds <num> burst packets, stopping for <num> seconds!!	<p>Threshold parameters for ICMP transit (through) traffic have been configured on an interface, and the maximum burst size for ICMP packets on the interface has been exceeded.</p> <p>The <portnum> is the port number.</p> <p>The first <num> is the maximum burst size (maximum number of packets allowed).</p> <p>The second <num> is the number of seconds during which additional ICMP packets will be blocked on the interface.</p> <p>Note: This message can occur in response to an attempted Smurf attack.</p>
Notification	Local TCP exceeds <num> burst packets, stopping for <num> seconds!!	<p>Threshold parameters for local TCP traffic on the device have been configured, and the maximum burst size for TCP packets has been exceeded.</p> <p>The first <num> is the maximum burst size (maximum number of packets allowed).</p> <p>The second <num> is the number of seconds during which additional TCP packets will be blocked on the device.</p> <p>Note: This message can occur in response to an attempted TCP SYN attack.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	Transit TCP in interface <portnum> exceeds <num> burst packets, stopping for <num> seconds!!	<p>Threshold parameters for TCP transit (through) traffic have been configured on an interface, and the maximum burst size for TCP packets on the interface has been exceeded.</p> <p>The <portnum> is the port number.</p> <p>The first <num> is the maximum burst size (maximum number of packets allowed).</p> <p>The second <num> is the number of seconds during which additional TCP packets will be blocked on the interface.</p> <p>Note: This message can occur in response to an attempted TCP SYN attack.</p>
Notification	ISIS L1 ADJACENCY DOWN <system-id> on circuit <circuit-id>	<p>The Layer 3 Switch's adjacency with this Level-1 IS has gone down.</p> <p>The <system-id> is the system ID of the IS.</p> <p>The <circuit-id> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS L1 ADJACENCY UP <system-id> on circuit <circuit-id>	<p>The Layer 3 Switch's adjacency with this Level-1 IS has come up.</p> <p>The <system-id> is the system ID of the IS.</p> <p>The <circuit-id> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS L2 ADJACENCY DOWN <system-id> on circuit <circuit-id>	<p>The Layer 3 Switch's adjacency with this Level-2 IS has gone down.</p> <p>The <system-id> is the system ID of the IS.</p> <p>The <circuit-id> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS L2 ADJACENCY UP <system-id> on circuit <circuit-id>	<p>The Layer 3 Switch's adjacency with this Level-2 IS has come up.</p> <p>The <system-id> is the system ID of the IS.</p> <p>The <circuit-id> is the ID of the circuit over which the adjacency was established.</p>
Notification	ISIS ENTERED INTO OVERLOAD STATE	<p>The Layer 3 Switch has set the overload bit to on (1), indicating that the Layer 3 Switch's IS-IS resources are overloaded.</p>
Notification	ISIS EXITING FROM OVERLOAD STATE	<p>The Layer 3 Switch has set the overload bit to off (0), indicating that the Layer 3 Switch's IS-IS resources are no longer overloaded.</p>
Notification	DOT1X issues software but not physical port up indication of Port <portnum> to other software applications	<p>The device has indicated that the specified port has been authenticated, but the actual port may not be active.</p>

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Notification	DOT1X issues software but not physical port down indication of Port <portnum> to other software applications	The device has indicated that the specified is no longer authorized, but the actual port may still be active.
Notification	Authentication Enabled on <portnum>	The multi-device port authentication feature was enabled on the on the specified <portnum>.
Notification	Authentication Disabled on <portnum>	The multi-device port authentication feature was disabled on the on the specified <portnum>.
Notification	MAC Authentication succeeded for <mac-address> on <portnum>	RADIUS authentication was successful for the specified <mac-address> on the specified <portnum>.
Informational	Cold start	The device has been powered on.
Informational	Warm start	The system software (flash code) has been reloaded.
Informational	<user-name> login to USER EXEC mode	A user has logged into the USER EXEC mode of the CLI. The <user-name> is the user name.
Informational	<user-name> logout from USER EXEC mode	A user has logged out of the USER EXEC mode of the CLI. The <user-name> is the user name.
Informational	<user-name> login to PRIVILEGED mode	A user has logged into the Privileged EXEC mode of the CLI. The <user-name> is the user name.
Informational	<user-name> logout from PRIVILEGED mode	A user has logged out of Privileged EXEC mode of the CLI. The <user-name> is the user name.
Informational	SNMP Auth. failure, intruder IP: <ip-addr>	A user has tried to open a management session with the device using an invalid SNMP community string. The <ip-addr> is the IP address of the host that sent the invalid community string.
Informational	Interface <portnum>, state up	A port has come up. The <portnum> is the port number.
Informational	Interface <portnum>, state down	A port has gone down. The <portnum> is the port number.
Informational	Interface <portnum>, line protocol up	The line protocol on a port has come up. The <portnum> is the port number.

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Informational	Interface <portnum>, line protocol down	The line protocol on a port has gone down. The <portnum> is the port number.
Informational	Trunk group (<ports>) created by 802.3ad link-aggregation module.	802.3ad link aggregation is configured on the device, and the feature has dynamically created a trunk group (aggregate link). The <ports> is a list of the ports that were aggregated to make the trunk group.
Informational	Bridge root changed, vlan <vlan-id>, new root ID <string>, root interface <portnum>	A Spanning Tree Protocol (STP) topology change has occurred. The <vlan-id> is the ID of the VLAN in which the STP topology change occurred. The <root-id> is the STP bridge root ID. The <portnum> is the number of the port connected to the new root bridge.
Informational	Bridge is new root, vlan <vlan-id>, root ID <root-id>	A Spanning Tree Protocol (STP) topology change has occurred, resulting in the Foundry device becoming the root bridge. The <vlan-id> is the ID of the VLAN in which the STP topology change occurred. The <root-id> is the STP bridge root ID.
Informational	Bridge topology change, vlan <vlan-id>, interface <portnum>, changed state to <stp-state>	A Spanning Tree Protocol (STP) topology change has occurred on a port. The <vlan-id> is the ID of the VLAN in which the STP topology change occurred. The <portnum> is the port number. The <stp-state> is the new STP state and can be one of the following: <ul style="list-style-type: none"> • disabled • blocking • listening • learning • forwarding • unknown
Informational	startup-config was changed or startup-config was changed by <user-name>	A configuration change was saved to the startup-config file. The <user-name> is the user's ID, if they entered a user ID to log in.
Informational	vlan <vlan-id> interface <portnum> Bridge TC Event (DOT1wTransition)	802.1W recognized a topology change event in the bridge. The topology change event is the forwarding action that started on a non-edge Designated port or Root port.

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Informational	vlan <vlan-id> interface <portnum> STP state -> <state> (DOT1wTransition)	802.1W changed the state of a port to a new state: forwarding, learning, blocking. If the port changes to blocking, the bridge port is in discarding state.
Informational	vlan <vlan-id> New RootPort <portnum> (RootSelection)	802.1W changed the port's role to Root port, using the root selection computation.
Informational	vlan <vlan-id> New RootBridge <mac-address> RootPort <portnum> (BpduRcvd)	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.
Informational	vlan <vlan-id> Bridge is RootBridge <mac-address> (MgmtPriChg)	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Informational	vlan <vlan-id> Bridge is RootBridge <mac-address> (MsgAgeExpiry)	The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.
Informational	DOT1X: Port <portnum>, AuthControlledPortStatus change: authorized	The status of the interface's controlled port has changed from unauthorized to authorized.
Informational	DOT1X: Port <portnum>, AuthControlledPortStatus change: unauthorized	The status of the interface's controlled port has changed from authorized to unauthorized.
Informational	DOT1X: Port <portnum> currently used vlan-id changes to <vlan-id> due to dot1x-RADIUS vlan assignment	A user has completed 802.1X authentication. The profile received from the RADIUS server specifies a VLAN ID for the user. The port to which the user is connected has been moved to the VLAN indicated by <vlan-id>.
Informational	DOT1X: Port <portnum> currently used vlan-id is set back to port default vlan-id <vlan-id>	The user connected to <portnum> has disconnected, causing the port to be moved back into its default VLAN, <vlan-id>.
Informational	DOT1X Port <portnum> is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC address filters	802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred: <ul style="list-style-type: none"> Insufficient system resources were available on the device to apply an IP ACL or MAC address filter to the port Invalid information was received from the RADIUS server (for example, the Filter-ID attribute did not refer to an existing IP ACL or MAC address filter)
Informational	Port <portnum>, srcip-security max-ipaddr-per-int reached.Last IP=<ipaddr>	The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Informational	telnet SSH web access [by <username>] from src IP <source ip address>, src MAC <source MAC address> rejected, <n> attempt(s)	<p>There were failed web, SSH, or Telnet login access attempts from the specified source IP and MAC address.</p> <ul style="list-style-type: none"> [by <user> <username>] does not appear if telnet or SSH clients are specified. <n> is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.
Informational	user <username> added deleted modified from console telnet ssh web snmp	A user created, modified, or deleted a local user account via the Web, SNMP, console, SSH, or Telnet session.
Informational	vlan <vlan id> added deleted modified from console telnet ssh web snmp session	A user created, modified, or deleted a VLAN via the Web, SNMP, console, SSH, or Telnet session.
Informational	ACL <acl id> added deleted modified from console telnet ssh web snmp session	A user created, modified, deleted, or applied an ACL via the Web, SNMP, console, SSH, or Telnet session.
Informational	MAC Filter added deleted modified from console telnet ssh web snmp session filter id = <MAC filter ID>, src mac = <Source MAC address> any, dst mac = <Destination MAC address> any	A user created, modified, deleted, or applied this MAC filter via the Web, SNMP, console, SSH, or Telnet session.
Informational	SNMP read-only community read-write community contact location user group view engineid trap [host] [<value -str>] deleted added modified from console telnet ssh web snmp session	<p>A user made SNMP configuration changes via the Web, SNMP, console, SSH, or Telnet session.</p> <p>[<value-str>] does not appear in the message if SNMP community or engineid is specified.</p>
Informational	Syslog server <IP-address> deleted added modified from console telnet ssh web snmp OR Syslog operation enabled disabled from console telnet ssh web snmp	A user made Syslog configuration changes to the specified Syslog server address, or enabled or disabled a Syslog operation via the Web, SNMP, console, SSH, or Telnet session.
Informational	SSH telnet server enabled disabled from console telnet ssh web snmp session [by user <username>]	A user enabled or disabled an SSH or Telnet session, or changed the SSH enable/disable configuration via the Web, SNMP, console, SSH, or Telnet session.
Informational	Enable super port-config read-only password deleted added modified from console telnet ssh web snmp OR Line password deleted added modified from console telnet ssh web snmp	A user created, re-configured, or deleted an Enable or Line password via the Web, SNMP, console, SSH, or Telnet session.

Table A.2: Foundry Syslog Messages (Continued)

Message Level	Message	Explanation
Informational	Port <portnum>, srcip-security max-ipaddr-per-int reached.Last IP=<ipaddr>	The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Debug	BGP4: Not enough memory available to run BGP4	The device could not start the BGP4 routing protocol because there is not enough memory available.
Debug	DOT1X: Not enough memory	There is not enough system memory for 802.1X authentication to take place. Contact Foundry Technical Support.
Debug	<ul style="list-style-type: none"> • Out of Frame cleared • Loss of Signal cleared • Loss of Frame cleared • Loss of pointer cleared • Path AIS cleared • Line AIS cleared • RDI cleared • Unequipped Signal cleared • Signal Label Mismatch cleared • Out of Frame detected • Loss of Signal detected • Loss of Frame detected • Loss of pointer detected • Path AIS detected • Line AIS detected • RDI detected • Unequipped Signal detected • Signal Label Mismatch detected 	Used by Foundry Technical Support for troubleshooting Packet Over SONET (POS) interfaces.

Syslog Messages for Web Authentication

The following Syslog messages are available for the Web Authentication feature. See the Web Authentication chapter in the *Foundry Switch and Router Installation and Basic Configuration Guide* for a list of devices that support Web Authentication.

Message Level	Message	Explanation
Informational	Web Auth: Enabled on vlan <vlan id>	Web Authentication has been enabled on the specified VLAN.

Message Level	Message	Explanation
Informational	Web Auth: Disabled on vlan <vlan id>	Web Authentication has been disabled on the specified VLAN.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically authorized for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically authenticated for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically authorized for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically authenticated for infinity.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically authorized on port <port num> for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically authenticated on the specified interface for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically authorized on port <port num> for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically authenticated on the specified interface for infinity
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically blocked for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically blocked for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically blocked for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically blocked for infinity.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically blocked on port <port num> for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically blocked on the specified interface for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is statically blocked on port <port num> for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address has been configured to be automatically blocked on the specified interface for infinity.
Informational	Web Auth in Vlan <vlan id>: <mac address> is removed from the block list	On the specified VLAN where Web Authentication is enabled, the specified MAC address has been removed from the blocked MAC address list.
Informational	Web Auth in Vlan <vlan id>: <mac address> is removed from the authorized list	On the specified VLAN where Web Authentication is enabled, the specified MAC address has been removed from the authorized MAC address list.

Message Level	Message	Explanation
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically authorized for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be authenticated for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically authorized for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be authenticated for infinity.
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically authorized on port <port num> for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be authenticated on the specified port for the specified duration
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically authorized on port <port num> for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be authenticated on the specified port for the specified infinity
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically blocked for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated and it will be blocked from any Web Authentication attempts for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically blocked for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be blocked from Web Authentication for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically blocked on port <port_num> for a duration of <value> seconds	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be blocked from Web Authentication on the specified port for the specified duration.
Informational	Web Auth in Vlan <vlan id>: <mac address> is updated to be statically blocked on port <port_num> for a duration of infinity	On the specified VLAN on which Web Authentication is enabled, the specified MAC address was updated to be blocked from Web Authentication on the specified port for infinity.
Informational	Web Auth in Vlan <vlan id>: Authentication succeeded for user: <user name> using mac: <mac address>	The specified user using the specified MAC address passed Web Authentication on the specified VLAN.
Informational	Web Auth in Vlan <vlan id>: Authentication failed (# of attempts tried: <value>) for user: <user name> using mac: <mac address>	The specified user using the specified MAC address failed the specified number of Web Authentication attempts on the specified VLAN.

Message Level	Message	Explanation
Informational	Web Auth in Vlan <vlan id>: User: <user name> using mac: <mac address> is blocked because of reaching the maximum number of authentication attempts (# of attempts tried: <value>)	On the specified VLAN on which Web Authentication is enabled, the specified user who is using the specified MAC address has been blocked from any additional Web Authentication attempts because the user reached the maximum number of attempts allowed.
Informational	Web Auth in VLAN (vlan id>: Authentication succeeded for user: <username> using mac: <mac address> but is blocked because the maximum number of authenticated hosts reached (max number of hosts: <value>)	On the specified VLAN on which Web Authentication is enabled, the specified user who is using the specified MAC address has been blocked from any additional Web Authentication attempts because the maximum number of hosts that can be authenticated at one time has been reached.

Appendix B

Hardware Specifications

NOTE: For information about JetCore chassis modules, see “JetCore Chassis Modules” on page D-1.

Chassis Devices

Foundry Networks offers three families of Chassis devices:

- BigIron
- NetIron Metro Router / NetIron Internet Backbone router
- FastIron

BigIron

Foundry Networks' BigIron Layer 2 and Layer 3 Switches provide next-generation, hardware-based Layer 2/3/4 switching and multi-protocol routing on a single Chassis device.

Enterprises and Internet service providers (ISPs) can use BigIron to build very high-performance, end-to-end packet networks that provide the Quality of Service (QoS) needed to support delay-sensitive traffic.

The BigIron family includes the following:

- BigIron 4000 (a 4-slot chassis) – see Figure B.1
- BigIron 8000 (an 8-slot chassis) – see Figure B.2
- BigIron 15000 (a 15-slot chassis) – see Figure B.3

Each slot of the Chassis device can be populated by either a forwarding module or a management module.

Each system requires one **management module** and can accept a second one for redundancy. Management modules are available with 10/100 Mbps, 100 Mbps fiber ports or Gigabit Ethernet ports and provide a serial port for console access. You can install a management module in any chassis slot.

All non-management modules (those without a serial management port), are referred to as **forwarding modules**.

Figure B.1 BigIron 4000




Figure B.2 BigIron 8000 Chassis





Figure B.3 BigIron 15000 Chassis



NetIron Router

The NetIron 400, 800, and 1500 are Chassis-based routers for Metro and ISP networks. The NetIron is based on the BigIron architecture and provides the same high-throughput, non-blocking performance as the BigIron Layer 3 Switches. The following models are available. Figure B.4 shows an example of the NetIron 400.

Figure B.4 NetIron 400 (model NI400-4)





Figure B.5 shows an example of the NetIron 800.

Figure B.5 NetIron 800 (model NI800-4)



The NetIron supports many of the forwarding modules supported by the BigIron. The NetIron requires the Management 4 module.

FastIron II Family

The FastIron II, FastIron II Plus, and FastIron III are next-generation closet solutions that provide enterprises with the resiliency required in mission critical networks.

Foundry offers many configurations of the FastIron II products with various combinations of 10/100, SX, LX, and GC ports to meet your networking needs. You also of course can order individual modules as needed for upgrades, replacements, or spares.

Redundant Management Modules

The Chassis devices support **redundant management modules**. Redundant management modules provide redundancy to prevent downtime in the event that a management module stops operating.

You can use one or two redundant management modules in a Chassis device. Using two redundant management modules adds fault protection against system outage. The two modules work together as active and standby management modules. If the active module becomes unavailable, the standby module automatically takes over system operation.

For more information and complete configuration and management information, see “Using Redundant Management Modules” on page 3-1.

Stackable Devices

Foundry Networks provides the following Layer 2 and Layer 3 Stackable devices:

- FastIron 4802
- FastIron Edge Switch
- NetIron stackable Layer 3 Switch
- NetIron 4802

NOTE: For information about the FastIron Edge Switch, see the release notes that come with that product.

FastIron 4802

The FastIron 4802 is a 1-1/2 rack-unit (RU) high switch that provides 48 10/100 Ethernet ports and can support two 1000 Mbps Ethernet uplink ports. The uplink ports require mini-GBIC connectors and can support 1000BaseSX and 1000BaseLX.

FastIron Edge Switch

The FastIron Edge Switches provide high 10/100 port density and Gigabit Ethernet uplinks in a compact form factor.

- The FastIron Edge Switch 2402 has 24 10/100 ports and two Gigabit uplink ports.
- The FastIron Edge Switch 4802 has 48 10/100 ports and two Gigabit uplink ports.
- The FastIron Edge Switch 9604 has 96 10/100 ports and four Gigabit uplink ports.

For more information about the FastIron Edge Switch, see the release notes that come with your product.

NetIron Layer 3 Switch

The NetIron Layer 3 Switch is used to build low-cost, wire-speed collapsed router backbones.

The NetIron Layer 3 Switch comes in either a 16-port or 24-port 10BaseT/100BaseTX configuration and either an 8-port or 16-port 100BaseFX Fiber configuration. Optional Fast Ethernet and Gigabit expansion modules are also available for building large, fast Gigabit networks.

NetIron 4802

The NetIron 4802 is a JetCore device. It has 48 10BaseT/100BaseT (10/100) ports and two mini-GBIC slots for 1000BaseSX, LX, LHA, or LHB fiber connections, in a compact, stackable form factor.

Control Features

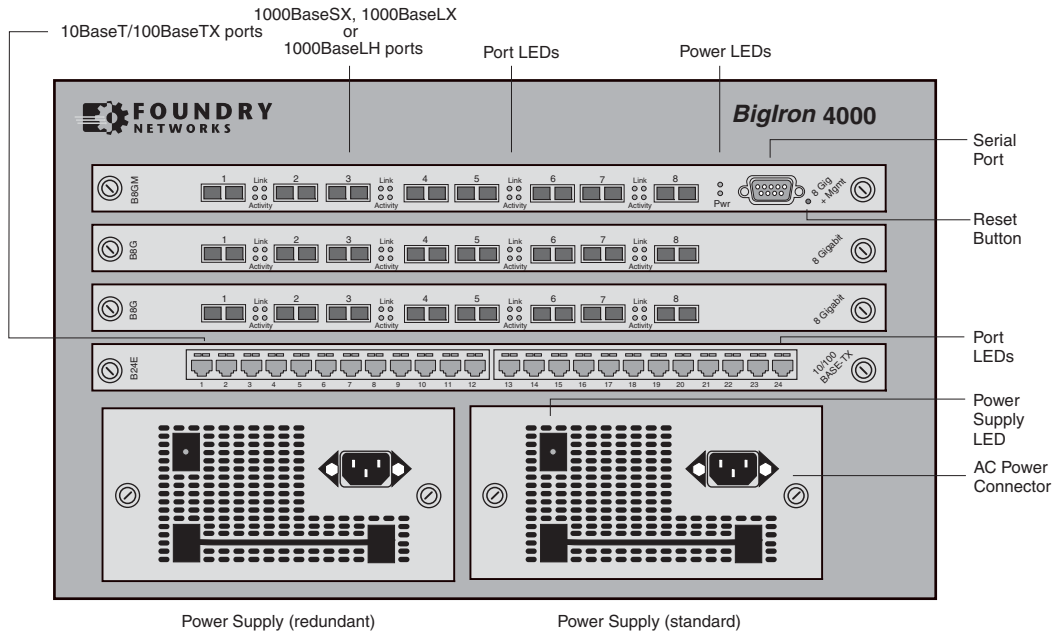
This section describes the external control features of the Foundry Chassis devices and Stackable devices.

Control Panels

The following sections show the control features of Foundry devices.

Chassis Device

Figure B.6 Front panel of a BigIron 4000 chassis



FastIron 4802

Figure B.7 Front panel of a FastIron 4802


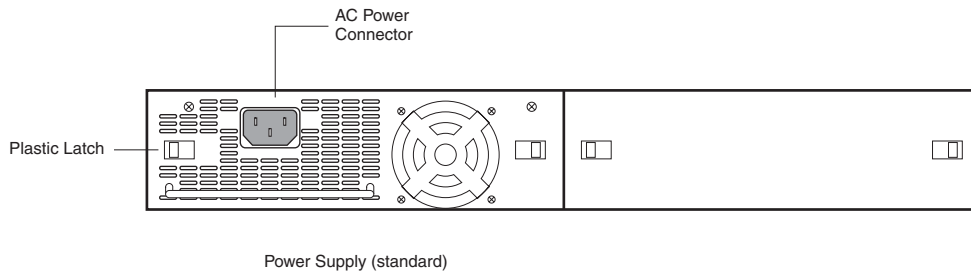


Figure B.8 Rear panel of a FastIron 4802



Netron Stackable Device

Figure B.9 Front panel of a Netron stackable Layer 3 Switch

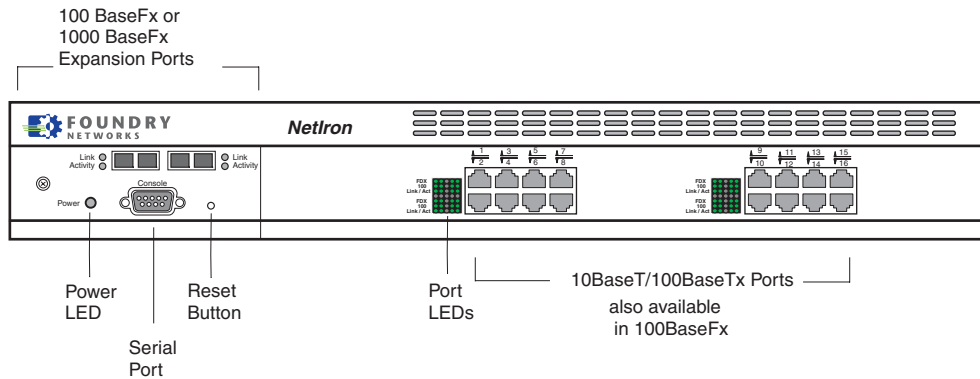




Figure B.10 Rear panel of a Netron stackable Layer 3 Switch



NOTE: The rear panel of a Chassis device does not provide network or power connections and therefore is not shown.

Netron 4802

Figure B.11 Front Panel of a Netron 4802



Ports

The following port types are supported on Foundry devices.

10 Gigabit Ethernet

The 10 Gigabit ports are compliant with the 10 Gigabit Ethernet standard, IEEE 802.3ae. Foundry 10 Gigabit ports support 1310nm and 1510nm serial connections to single-mode fiber. The port connectors are SC connectors.

1000BaseT Gigabit Copper (GC) Ports

The 1000BaseT Gigabit Copper (GC) ports are compliant with the IEEE 802.3ab standard and can provide Gigabit throughput over standard category-5 (“Cat-5”) copper wiring. The port connectors are RJ-45s, the same as the connectors on Foundry’s 10/100 modules. Thus, you can immediately deploy the GC ports without recabling.

Starting with Enterprise IronWare release 07.7.00 and Service Provider IronWare release 09.1.00, copper mini-GBICs are supported on Gigabit Ethernet modules. You can use fiber mini-GBICs on some ports on a Gigabit Ethernet module, and copper mini-GBICs on others.

Copper mini-GBICs are supported on both IronCore or IronCore and JetCore Gigabit Ethernet modules, with the exception of the J-F2404GMR4 module. Copper mini-GBICs are also supported on the FastIron 4802. Note that the copper mini-GBICs can operate in 1000 Mbps autonegotiation mode only. You cannot configure them to operate in other modes.

10BaseT/100BaseTX Ports

The 10BaseT/100BaseTX ports are auto-sensing, auto-negotiating ports. Most have RJ-45 UTP connectors. These ports accept category-5 Unshielded Twisted Pair (UTP) cables. The JetCore 48-port Telco module uses RJ-21 connectors.

100BaseFX Ports

The 100BaseFX ports are equipped with MT-RJ connectors and operate at 100 Mbps in full-duplex mode.

1000BaseSX Ports

The 1000BaseSX ports operate in full-duplex mode and are equipped with SC connectors on fixed-configuration modules and can be mini-GBICs with LC connectors for mini-GBIC modules that support this port type. Multi-mode fiber cabling is supported.

1000BaseLX

The 1000BaseLX ports operate in full-duplex mode and are equipped with SC connectors on fixed-configuration modules and come as mini-GBICs for mini-GBIC modules that support this port type. Both single-mode fiber (SMF) and multi-mode fiber (MMF) cabling is supported. The 1000BaseLX ports must be connected to another 1000BaseLX port. Connection to a 1000BaseSX port is not supported.

1000BaseLH

The 1000BaseLH ports operate in full-duplex mode and are equipped with SC connectors on fixed-configuration modules. Single-mode fiber cabling is supported.

NOTE: 1000BaseSX, 1000BaseLX, and 1000BaseLH ports also support auto-negotiation when the auto-gig option is enabled on the system.

NOTE: 1000BaseSX, 1000BaseLX, and 1000BaseLH ports operate only at full-duplex.

CWDM Mini-GBICs

Starting with Enterprise IronWare release 07.7.00 and Service Provider IronWare release 09.1.00, support is included for Coarse Wavelength Division Multiplexing (CWDM) mini-GBICs on Gigabit Ethernet modules. CWDM mini-GBICs allow you to connect up to 8 Gigabit Ethernet ports over a single fiber link using mux/demux

equipment at both ends of the link. CWDM mini-GBICs are supported on both IronCore and JetCore Gigabit Ethernet modules.

The Foundry device automatically detects the CWDM mini-GBICs. Valid part numbers are FWDM-1519-7D-xx for transceivers with a reach of 80 km, where xx and yy represent the wavelength of the transceivers.

Each of the ports must operate on different wavelengths. The following wavelengths are supported:

- 1470 nm
- 1490 nm
- 1510 nm
- 1530 nm
- 1550 nm
- 1570 nm
- 1590 nm
- 1610 nm

Packet Over SONET (POS) Ports

The POS ports use SC duplex connectors. For more information, see “Using Packet Over SONET Modules” on page 7-1.

Asynchronous Transfer Mode (ATM) Ports

The ATM ports use SC duplex connectors. For more information, see “Using Asynchronous Transfer Mode Modules” on page 8-1.

Slot and Port Numbers

The port numbers on all Stackable devices and Chassis devices are labeled on the hardware. However, the method you use to enter or select a port number differs depending on whether you are managing a Stackable device or a Chassis device.

Stackable Devices

To specify a port number in the software, enter or select the number associated with the port on the device's front panel. For example, to assign a name to port 8 on a Stackable device, enter the following CLI commands:

```
FastIron(config)# interface ethernet 8
FastIron(config-if-8)# port-name pdtmarketing
```

Syntax: interface ethernet <portnum>

Syntax: port-name <string>

Chassis Devices

The port numbers on the modules in s are labeled, but the slot numbers are not labeled.

- Slots on a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots on an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots on a 15-slot chassis are numbered 1 – 15, from left to right.

You can place a management module in any slot. The slot numbers are absolute and do not change based on the position of the management module.

To specify a port on a Layer 3 Switch, enter the slot number, a forward slash (/), and the number associated with the port on the device's front panel. For example, to assign a name to Ethernet port 8 on the module installed in Chassis slot 2, enter the following commands:

```
BigIron(config)# interface e 2/8
BigIron(config-if-2/8)# port-name pdtmarketing
```

Syntax: interface ethernet <portnum>

Syntax: port-name <string>

NOTE: The Stackable devices do not contain separate slots and thus do not use slot numbers.

LEDs

Each Foundry device is equipped with LEDs that denote port and power supply status. The tables below reflect the different port and expansion module port states.

Table B.1: Chassis 100BaseFX, 1000BaseSX/LX, and 1000BaseT LEDs

LED	Position	State	Meaning
Link	Top	On	Port is connected.
		Off	No port connection exists.
Activity	Bottom	On	Traffic is being transmitted and received on that port.
		Off	No traffic is being transmitted.
		Blinking	Traffic is being transmitted and received on that port.

Table B.2: Chassis 10BaseT/100BaseTX LEDs

LED	Position	State	Meaning
Link/Activity	Left	On	Port is connected.
		Off	No port connection exists.
		Blinking	Traffic is being transmitted and received on that port.
FDX	Right	On	The port is operating at full-duplex.
		Off	The port is operating at half-duplex.

Table B.3: FastIron 4802 LEDs for 10/100 Mbps Ports

LED	Position	State	Meaning
FDX	Right side of port connector	On	Full-duplex connection found or configured. Note: This LED also is lit if you configure the port to 10 Mbps full-duplex or 100 Mbps full-duplex. This is true even when no link is present.
		Off	Half-duplex connection or no port connection exists.
Link/Activity	Left side of port connector	On	Connection established, no activity.
		Off	No connection established.
		Blinking	Connection established with activity on the link.

Table B.4: Stackable NetIron Layer 3 Switch LEDs

LED	Position	State	Meaning
FDX/HDX	Top	On	The port is operating at full-duplex.
		Off	The port is operating at half-duplex.
100	Middle	On	The port is operating at 100 Mbps.
		Off	The port is not operating at 100 Mbps.
Link/Act	Bottom	On	Port is connected.
		Off	No port connection exists.
		Blinking	Traffic is being transmitted or received on that port.

Table B.5: NetIron 4802 LEDs for 10/100 Mbps Ports

LED	Position	State	Meaning
Lnk/Act	Left LED above port	On	Link is up.
		Off	Link is down.
		Blinking	Port is transmitting or receiving.

Table B.5: NetIron 4802 LEDs for 10/100 Mbps Ports (Continued)

LED	Position	State	Meaning
FDX	Right LED above port	On	Full-duplex connection found or configured. Note: This LED also is lit if you configure the port to 10 Mbps full-duplex or 100 Mbps full-duplex. This is true even when no link is present.
		Off	Half-duplex connection or no port connection exists.
		Blinking	Collisions are being detected.

Table B.6: NetIron 4802 LEDs for 1000 Mbps Ports

LED	Position	State	Meaning
Lnk	Top	On	Fiber port is connected.
		Off	No fiber port connection exists.
Act	Bottom	On	Traffic is being transmitted and received on the fiber port.
		Off	No traffic is being transmitted on the fiber port.
		Blinking	Traffic is being transmitted and received on the fiber port.

Table B.7: POS Port LEDs

LED	Position	State	Meaning
Link	Upper left	On	Port is connected.
		Off	No port connection exists.
Alarm	Upper right	On	At least one of the following SONET alarm conditions has been detected: <ul style="list-style-type: none"> • LOS – Loss of Signal • LOF – Loss of Frame • LOP – Loss of Pointer • AIS – Alarm Indication Signal
		Off	None of the alarm conditions listed above have been detected.
TxAct	Lower left	Blinking	The port is transmitting traffic.
RxAct	Lower right	Blinking	The port is receiving traffic.

Table B.8: ATM Port LEDs

LED	Position	State	Meaning
Link	Upper left	On	The port is connected.
		Off	The port is not connected.

Table B.8: ATM Port LEDs (Continued)

LED	Position	State	Meaning
Alarm	Upper right	On	At least one of the following SONET alarm conditions has been detected: <ul style="list-style-type: none"> • LOS – Loss of Signal • LOF – Loss of Frame • LOP – Loss of Pointer • AIS – Alarm Indication Signal
		Off	None of the alarm conditions listed above have been detected.
TxAct	Lower left	Blinking	The port is transmitting traffic.
RxAct	Lower right	Blinking	The port is receiving traffic.

Reset Button

The reset button allows you to restart the system. The reset button is recessed to prevent it from being pushed accidentally.

- For Stackable devices, the reset button is located to the right of the serial port as labeled in Figure B.9.
- For Chassis devices, the reset button is located to the right of the serial port on the management module as labeled in Figure B.6.

Power Specifications

The following table lists the power consumption for Foundry devices. This information is subject to updates. For the latest information, see the following web page:

<http://www.foundrynet.com/services/faqs/power.html#modules>

Power Specifications for Chassis Devices

Table B.9: Foundry Chassis device Power Supply Ratings

	4-slot	8-slot	8-slot	15-slot
Minimum 100 – 120v AC Wall Outlet Circuit per Power Supply	15A	15A	15A	20A
Maximum 100 – 120v AC Circuit Rating ^a	7.5A	7.5A	15A	30A
Minimum 200 – 240v AC Wall Outlet Circuit per Power Supply	4A	4A	4A	7.5A
Maximum 200 – 240v AC Circuit Rating ^a	4A	4A	8A	15A
Maximum -70 to -40v DC Rating per Power Supply	22A	22A	22A	44A
Minimum number of Power Supplies Required for Operation	1	1 for 1 – 3 modules	2 for 4 or more modules	2

Table B.9: Foundry Chassis device Power Supply Ratings

	4-slot	8-slot	8-slot	15-slot
Minimum number of Power Supplies for N+1 Redundancy ^b	2	3	3	3
Maximum number of Power Supplies for 100% Redundancy ^b	2	4	4	4

a. Assumes that the **minimum** number of Power Supplies required for operation are connected to the same wall outlet/circuit.

b. Number of power supplies installed for N+1 or 100% redundancy in the chassis does not increase Maximum Ratings.

The following equation is used to calculate the values listed in Table B.9:

$$\text{CurrentDraw} = \frac{\text{PowerOutput}}{(\text{VoltsIn})(\text{Efficiency})(\text{PowerFactor})}$$

Table B.10 lists the current draw equations for each Chassis device.

Table B.10: Current Draw Calculations for Chassis devices

CurrentDraw	4-slot 8-slot	15-slot
110VAC	$\frac{550}{(110)(0.7)(0.95)} = 7.5A$	$\frac{1100}{(110)(0.7)(0.95)} = 15.1A$
220VAC	$\frac{550}{(220)(0.7)(0.95)} = 3.8A$	$\frac{1100}{(220)(0.7)(0.95)} = 7.5A$
36VDC	$\frac{550}{(36)(0.7)(1)} = 22A$	$\frac{1100}{(36)(0.7)(1)} = 44A$

Table B.11: Maximum Power Calculations for Chassis devices

Product	Watts	BTUs
FastIron II	~550	~1877
FastIron II Plus	~1100	3753
FastIron III	~2200	7508
BigIron 4000	~550	~1877
BigIron 8000	~1100	~3753
BigIron 15000	~2200	7508
NetIron 400	~550	~1877
NetIron 800	~1100	~3753
NetIron 1500	~2200	7508

Table B.12: Wattage Consumed by the Chassis Themselves

Chassis Size	Watts
4-slot or 8-slot chassis	125
15-slot chassis	TBD

Table B.13: Wattage Consumed by Individual Modules

Module	Maximum Power Consumption, in Watts
JetCore Management Modules	

Table B.13: Wattage Consumed by Individual Modules

J-FxGMR4-BASE	65
J-FxGMR4	65
J-BxGMR4	70
J-F2404-GMR4	85
JetCore Ethernet Interface Modules	
J-F48E	95
J-B48E	100
J-F48T	95
J-B48T	100
J-FxG	60
J-BxG	65
J-F16Gx	95
J-B16Gx	100
J-F16GC	120
J-B16GC	125
10 Gig Module	115
IronCore Management Modules	
2-port SX/LX	105
4-port SX/LX	110
8-port SX/LX	120
16-port 10/100	112
IronCore Ethernet Interface Modules	
B2G	105
B4G	105
B8G	105
B24E	105
B24FX	105
WAN Interface Modules	
2-port OC-3c	TBD
4-port OC-3c	TBD
2-port OC-12c	TBD
2-port OC-48c	TBD
2-port OC-48c NPA	TBD

Power Specifications for Stackable Devices

Table B.14: Foundry Stackable device Power Supply Ratings

Product	Minimum 100 – 120v AC Wall Outlet Circuit per Power Supply	Minimum 200 – 240v AC Wall Outlet Circuit per Power Supply	Maximum -70 to -40v DC Rating per Power Supply	Minimum # of Power Supplies Required for Operation	Minimum # of Power Supplies for N+1 Redundancy^a
FastIron Workgroup Switch (FWS24)	2.5A	1.4A	5.7A	1	2
FastIron Wiring Closet Switch (FWS4802)	4A	2A	10A	1	2
NetIron Stackable Router (NSR16/24)	2.5A	1.4A	5.7A	1	2
NetIron 4802	2.5A	1.4A	5.7A	1	2
TurboIron/8 Switch & Router	7.5A	3.75A	N/A	1	2
ServerIronXL	2.5A	1.4A	5.7A	1	2
ServerIronXL/G	7.5A	3.75A	N/A	1	2

a.Number of power supplies installed for N+1 or 100% redundancy in the chassis does not increase Maximum Ratings.

Table B.15: Foundry Stackable device Power Supply Ratings

Product	Minimum 100 – 120v AC Wall Outlet Circuit per Power Supply	Minimum 200 – 240v AC Wall Outlet Circuit per Power Supply	Minimum # of Power Supplies Required for Operation	Minimum # of Power Supplies for N+1 Redundancy^a
TurboIron/8	7.5A	3.75A	1	2
TurboIron/8				

a.Number of power supplies installed for N+1 or 100% redundancy in the chassis does not increase Maximum Ratings.

Physical Dimensions

Table B.16: Physical dimensions for Foundry devices

Platform	Depth	Width	Length (Height)	Weight
15-slot Layer 3 Switch	15"	17.5"	29.75"	256 lbs. fully populated
8-slot Layer 3 Switch	15"	17.5"	23"	69.1 lbs. fully populated
4-slot Layer 3 Switch	15"	17.5"	9"	47.7 lbs. fully populated
Turbolron/8 FastIron Workgroup FastIron Backbone NetIron stackable ServerIron ServerIronXL ServerIronXL/G Turbolron (4- to 6- port Gigabit Layer 2 Switch)	16.75"	17.5"	2.75"	18 – 22 lbs.
NetIron 4802	18.64"	17.5"	2.75"	23.5 lbs. fully loaded

Operating Environment

- Operating Temperature: 32° – 104° F, 0° – 40° C (0° – 50° for DC power supply)
- Relative Humidity: 5% – 90%, non-condensing (20% – 90% non-condensing for DC power supply)
- Operating Altitude: 0 – 6,562 feet (2,000 meters)

Storage Environment

- Storage Temperature: 9° – 158° F, -25° – 70° C
- Storage Humidity: 95% maximum, non-condensing
- Storage Altitude: 10,000 feet (3,000 meter) maximum

Electromagnetic Emissions

- FCC Class A, Part 15, Subpart B
- EN 55022A Class A
- VCCI Class A
- EN50082-1

Safety Agency Approvals

- UL 1950
- CSA-C22.2 No. 950 93
- TUV EN 60950, EN 60825

Laser

- Class 1 Laser Product
- Laser Klasse 1
- Complies with IEC 825-2:1993

Appendix C

Software Specifications

This appendix lists the following information:

- IEEE compliance
- RFC support
- ISO/IEC specification support
- Internet draft support

NOTE: For a list of features supported on a specific product, see the data sheet for that product.

IEEE Compliance

Foundry devices support the following standards.

- 802.1D Bridging
- 802.1p/q VLAN Tagging
- 802.1w Rapid Spanning Tree (RSTP)
- 802.1X Port-Based Network Access Control
- 802.3, 10BaseT
- 802.3ad Link Aggregation
- 802.3ae 1000BaseX
- 802.3u, 100BaseTX, 100BaseFX
- 802.3z 1000BaseSX, 1000BaseLX
- 802.3x Flow Control

RFC Support

The following table lists the RFCs supported by Foundry devices.

NOTE: Some devices support only a subset of the RFCs. For example, Layer 2 Switches and the ServerIron do not support router-specific RFCs.

Table C.1: Foundry RFC Support

RFC Number	Protocol or Standard
768	User Datagram Protocol (UDP)
783	Trivial File Transfer Protocol (TFTP)
791	Internet Protocol (IP)
792	Internet Control Message Protocol (ICMP)
793	Transmission Control Protocol (TCP)
826	Ethernet Address Resolution Protocol (ARP)
854, 855, and 857	Telnet
894	IP over Ethernet frames
903	Reverse ARP (RARP)
906	Bootstrap loading using TFTP
919	Broadcast Internet datagrams
920	Domain requirements
922	Broadcast Internet datagrams in the presence of subnets
950	Internet standard subnetting procedure
951	Bootstrap Protocol (BootP)
1027	Proxy ARP
1042	IP datagrams over IEEE 802 networks (for Ethernet)
1058	Route Information Protocol (RIP) version 1
1075	Distance Vector Multicast Routing Protocol
1112	Internet Gateway Management Protocol (IGMP)
1122 and 1123	Requirements for Internet hosts (routers)
1141	Incremental updating of the Internet checksum
1155	Structure and Identification of Management Information (SMI)
1157	Simple Network Management Protocol (SNMP) version 1
1195	Use of OSI IS-IS for Routing in TCP/IP and Dual Environments

Table C.1: Foundry RFC Support (Continued)

RFC Number	Protocol or Standard
1212	Concise MIB Definitions
1213	MIB II Definitions
1215	SNMP generic traps
1256	ICMP Router Discovery Protocol (IRDP)
1267	Border Gateway Protocol version 3
1321	The MD5 Message-Digest Algorithm
1340	Assigned numbers (where applicable)
1354	IP Forwarding Table MIB
1377	The PPP OSI Network Layer Control Protocol (OSINLCP)
1398	Ethernet-Like MIB
1492	An Access Control Protocol, Sometimes Called TACACS
1493	Bridge MIB (excluding filtering of objects)
1541 and 1542	Dynamic Host Configuration Protocol (DHCP)
1583	Open Shortest Path First (OSPF)
1587	OSPF Not-So-Stubby Areas (NSSAs)
1626	Default IP MTU for use over ATM AAL5
1657	Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol
1661	The Point-to-Point Protocol (PPP)
1662	PPP in HDLC-like Framing
1723	RIP version 2
1742	AppleTalk Management Information Base II
1745	OSPF Interactions
1757	Remote Monitoring (RMON) groups 1, 2, 3, 9
1765	OSPF Database Overflow
1771	Border Gateway Protocol (BGP) version 4
1812	Requirements for IP version 4 routers
1850	Open Shortest Path First (OSPF) version 2 MIB
1905	Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)
1906	Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)
1966	BGP Route Reflection
1977	BGP Communities

Table C.1: Foundry RFC Support (Continued)

RFC Number	Protocol or Standard
1997	BGP Communities Attributes
2003	IP Tunneling
2030	Simple Network Time Protocol (SNTP) version 4
2068	HTTP
2138	Remote Authentication Dial In User Server (RADIUS)
2139	RADIUS Accounting
2178	Open Shortest Path First (OSPF)
2205	Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification
2225	Classical IP and ARP over ATM
2233	The Interfaces Group MIB using SMIv2
2328	OSPF Version 2 Note: AS External LSA reduction is supported.
2336	IGMP Version 2
2338	Virtual Router Redundancy Protocol (VRRP)
2362	IP Multicast PIM Sparse
2370	The OSPF Opaque LSA Option
2385	TCP MD5 Signature Option (for BGP4)
2439	BGP Route Flap Dampening
2453	BGP Route Information Protocol (RIP) version 2
2515	Definitions of managed objects for ATM management
2570	Introduction to Version 3 of the Internet-standard Network Management Framework
2571	An Architecture of Describing SNMP Management Frameworks
2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
2574	User-based Security (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
2615	PPP over SONET/SDH
2665	Ethernet Like MIB (incorporates RFC 1398)
2674	Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
2684	Multiprotocol Encapsulation over ATM Adaptation Layer 5
2702	Requirements for Traffic Engineering Over MPLS
2763	Dynamic Host Name Exchange Mechanism for IS-IS

Table C.1: Foundry RFC Support (Continued)

RFC Number	Protocol or Standard
2796	BGP Route Reflection
2842	BGP Capability Advertisement
2858	BGP Multi-protocol Extension
2869	RADIUS Extensions
2918	Route Refresh Capability for BGP-4
2966	Domain-wide Prefix Distribution with Two-Level IS-IS
3031	Multiprotocol Label Switching Architecture
3032	MPLS Label Stack Encoding
3036	LDP Specification
3065	BGP Confederations
3176	InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks

ISO/IEC Specifications

- ISO/IEC 10589 – Information Technology – Telecommunication and information exchange between systems – Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)
- ISO/IEC 8473 – Information processing systems – Data Communications – Protocols for providing the connectionless-mode network service
- ISO/IEC 9542 – Information Technology – Telecommunication and information exchange between systems – End system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)

Internet Drafts

In addition to the RFCs listed in “RFC Support” on page C-2, the Layer 3 Switches support the following Internet drafts:

- ietf-idmr-dvmp version 3.05, obsoletes RFC 1075
- draft-ietf-pim-dm-05 (V1)
- draft-ietf-pim-v2-dm-03 (V2)
- draft-ietf-mpls-icmp-02.txt
- draft-ietf-mpls-rsvp-lsp-tunnel-08.txt
- draft-katz-yeung-ospf-traffic-03.txt
- draft-martini-l2circuit-trans-mpls-07.txt
- draft-martini-l2circuit-encap-mpls-03.txt
- The TACACS+ Protocol version 1.78
- IS-IS extensions for Traffic Engineering

NOTE: Foundry supports the portions of this draft that describe the Extended IP reachability TLV (TLV type 135) and the extended Intermediate System (IS) reachability TLV (TLV type 22) to provide support for wide metrics.

Appendix D

JetCore Chassis Modules

NOTE: This appendix describes JetCore chassis modules. For general hardware information, including power specifications, see the “Hardware Specifications” on page B-1.

Foundry JetCore modules provide enhanced system performance through custom-designed ASICs.

Determining Your Device Type

Chassis devices are either JetCore or IronCore devices, depending on whether the management module is a JetCore or IronCore module. To determine whether a management module is JetCore or IronCore, check the part number. If the part number begins with “J-”, then the management module is a JetCore module. Otherwise, the management module is an IronCore module. JetCore modules are listed in Table D.1.

The FastIron 4802 is a JetCore device.

JetCore Chassis Modules

Table D.1 lists the JetCore modules for BigIron and FastIron s.

NOTE: Equivalents of most of these JetCore modules are available for NetIron Metro routers. Contact Foundry Networks. If you have a NetIron Metro router that uses JetCore modules, the module descriptions in this appendix also apply to your equivalent NetIron Metro JetCore modules.

NOTE: *You cannot use JetCore modules and non-JetCore modules in the same chassis.*

Table D.1: BigIron and FastIron JetCore Modules

Model	Chassis Type	Description	Ports	Interface Type
J-24FX	BigIron FastIron	Forwarding module	24 100Base-FX Fiber ports	MT-RJ connectors

Table D.1: BigIron and FastIron JetCore Modules (Continued)

Model	Chassis Type	Description	Ports	Interface Type
J-B2GMR4	BigIron	Management module	Two Gigabit Ethernet Fiber ports	Fiber mini-GBIC transceivers, Copper mini-GBICs, or CWDM mini-GBICs
J-BxGMR4	BigIron	Management module	Eight Gigabit Ethernet ports	Mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber
J-BxG	BigIron	Forwarding module	Eight Gigabit Ethernet ports	Mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber
J-B16GC	BigIron	Forwarding module	16 Gigabit Ethernet Copper ports	RJ-45s for Cat-5 copper (100/1000 Mbps)
J-B16Gx	BigIron	Forwarding module	16 Gigabit Ethernet Fiber ports	Mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber
J-B48E ^a	BigIron	Forwarding module	48 10/100 Ethernet ports	RJ-45s for Cat-5 copper
J-B48E-A ^a	BigIron	Forwarding module	48 10/100 Ethernet ports	RJ-45s for Cat-5 copper
J-B48T	BigIron	Forwarding module	48 10/100 Ethernet ports	50-pin Telco connectors for Cat-5 copper (12 ports per connector)
J-B48T-A	BigIron	Forwarding module	48 10/100 Ethernet ports	50-pin Telco connectors for Cat-5 copper (12 ports per connector)
J-F2GMR4	FastIron	Management module	2 Gigabit Ethernet Fiber ports	Fiber mini-GBIC transceivers, Copper mini-GBICs, or CWDM mini-GBICs
J-FixGMR4	FastIron	Management module	Eight Gigabit Ethernet ports	Mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber
J-F2404GMR4	FastIron	Management module	24 10/100 ports and four Gigabit Ethernet ports	RJ-45s for Cat-5 copper (10/100 and 1000 Mbps); mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber
J-FixG	FastIron	Forwarding module	Eight Gigabit Ethernet ports	Mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber
J-FI48E ^b	FastIron	Forwarding module	48 10/100 Ethernet ports	RJ-45s for Cat-5 copper
J-F48E-A ^b	FastIron	Forwarding module	48 10/100 Ethernet ports	RJ-45s for Cat-5 copper

Table D.1: BigIron and FastIron JetCore Modules (Continued)

Model	Chassis Type	Description	Ports	Interface Type
J-FI48T	FastIron	Forwarding module	48 10/100 Ethernet ports	50-pin Telco connectors for Cat-5 copper (12 ports per connector)
J-F48T-A	BigIron	Forwarding module	48 10/100 Ethernet ports	50-pin Telco connectors for Cat-5 copper (12 ports per connector)
J-F16GC	FastIron	Forwarding module	16 Gigabit Ethernet Copper ports	RJ-45s for Cat-5 copper (100/1000 Mbps)
J-F16Gx	FastIron	Forwarding module	16 Gigabit Ethernet Fiber ports	Mini-GBIC slots for 1000BaseSX or 1000BaseLX fiber

- a. The J-B48E and J-B48E-A modules are double-wide modules. They occupy two chassis slots.
b. The J-FI48E and J-F48E-A modules are double-wide modules. They occupy two chassis slots.

The modules listed in Table D.1 are described in the following sections.

The JetCore Management Module

Hardware Overview

JetCore ASICs

JetCore module ports are managed by the following custom ASICs:

- Integrated Gigabit Controllers (IGCs) – Ethernet packet controllers for Gigabit ports. Each Gigabit Ethernet module contains two IGCs.
- Integrated Packet Controllers (IPCs) – Ethernet packet controllers for 10/100 ports. Each 10/100 Ethernet module contains two IPCs.

These custom ASICs perform address lookup, data formatting and data movement for Ethernet packets. The Gigabit Ethernet modules use IGCs. The 10/100 modules use IPCs.

Each Gigabit Ethernet management or forwarding module has two IGCs.

- IGC 1 manages ports 1 – 4 on the module.
- IGC 2 manages ports 5 – 8 on the module.

Each 10/100 forwarding module has two IPCs:

- IPC 1 manages ports 1 – 24 on the module.
- IPC 2 manages ports 25 – 48 on the module.

Generally, you do not need to know which IGC or IPC is managing a port. However, the information is useful for a few features such as jumbo packet support and port monitoring. The documentation repeats the IGC and IPC port mapping information where needed.

Serial Management Interface

On management modules, the serial management interface enables you to configure and manage the device using a third-party terminal emulation application on a directly connected PC. A straight-through EIA/TIA DB-9 serial cable (M/F) is shipped with the device.

Reset Button

On management modules, the reset button allows you to restart the system. The button is recessed to prevent it from being pushed accidentally.

Temperature Sensor

Every JetCore module contains a temperature sensor. Depending on the temperature reported by the sensor, the software can send a warning if the temperature exceeds the normal threshold and can even shut the device down if the temperature exceeds the safe threshold. The software reads the temperature sensor according to the system poll time, which is 60 seconds by default.

You can display the temperature of the device. You also can change the warning and shutdown temperatures and the chassis poll time. See “Using the Temperature Sensor” on page 9-52.

J-B2GMR4 and J-F2GMR4 JetCore Management Modules

The J-B2GMR4 and J-F2GMR4 JetCore Management modules are supported starting in software releases 07.6.05 and 07.8.00.

NOTE: These Management Modules require boot code 07.6.05 or later.

Figure D.1 shows the front panel of a JetCore J-B2GMR4 Gigabit management module. The J-B2GMR4 is based on M4 technology for enhanced performance.

Figure D.1 J-B2GMR4 2-port JetCore management module



System Status LEDs

The LEDs listed in Table D.2 provide status information for the management processor and system power.

Table 1: System Status LEDs

LED	Position	State	Meaning
Link	Left side of serial interface, top	On	The management processor is active.
		Off	The management processor is not active.
Power	Left side of serial interface, bottom	On	The power status is good.
		Off	The power status is not good.

Gigabit Ethernet Network Interfaces

The JetCore Gigabit Ethernet management modules provide two miniature Gigabit Interface Converter (mini-GBIC) slots. The LEDs listed in Table D.3 provide status information for the ports.


Table 2: LEDs for 1000 Mbps Ports

LED	Position	State	Meaning
Link	Top	On	Port is connected.
		Off	No port connection exists.
Activity	Bottom	On	Traffic is being transmitted and received on that port.
		Off	No traffic is being transmitted.
		Blinking	Traffic is being transmitted and received on that port.

J-BxGMR4 and J-FixGMR4 JetCore Management Modules

The J-BxGMR4 and J-FixGMR4 management modules are based on M4 technology for enhanced performance. Figure D.2 shows the front panel of a JetCore Gigabit management module.

Figure D.2 J-FixGMR4 JetCore Gigabit management module



System Status LEDs

The LEDs listed in Table D.2 provide status information for the Management Processor and system power.

Table D.2: System Status LEDs

LED	Position	State	Meaning
Active	Left side of serial interface, top	On	The Management Processor is active.
		Off	The Management Processor is not active.
Pwr	Left side of serial interface, bottom	On	The power status is good.
		Off	The power status is not good.

Gigabit Ethernet Network Interfaces

The JetCore Gigabit Ethernet management modules provide eight miniature Gigabit Interface Converter (mini-GBIC) slots. You can insert a 1000BaseSX or 1000BaseLX or 1000Base-T mini-GBIC fiber connector into each slot, in any combination.

The LEDs listed in Table D.3 provide status information for the ports.


Table D.3: LEDs for 1000 Mbps Ports

LED	Position	State	Meaning
Link	Top	On	Port is connected.
		Off	No port connection exists.
Activity	Bottom	On	Traffic is being transmitted and received on that port.
		Off	No traffic is being transmitted.
		Blinking	Traffic is being transmitted and received on that port.

J-F2404GMR4 JetCore Management Module

Figure D.3 shows the front panel of a J-F2404GMR4.

Figure D.3 J-F2404GMR4 JetCore 10/100 and Gigabit management module



NOTE: The 2404 module is supported only in FastIron JetCore Chassis devices.

Management Redundancy

The J-F2404GMR4 supports management redundancy.

Slot Numbering

The J-F2404GMR4 is a double-wide module: the module occupies two chassis slots but uses only one of the slot's backplane connections.

The slot number used by the module is the slot that contains the lower half of the module. For example, in a 15-slot or 8-slot chassis, if you insert the module into slots 1 and 2, the module uses slot number 1. In a 4-slot chassis, if you insert the module into slots 1 and 2, the module uses slot 2.

Network Interfaces

The J-F2404GMR4 modules provide 24 10/100 ports and four 1000 Mbps ports. The 10/100 ports have RJ-45 connectors for Cat5 cabling. For flexibility, the module provides two types of physical interfaces for the 1000 Mbps ports: mini-GBIC slots for fiber cabling and RJ-45 connectors for Cat5 cabling. The Gigabit Copper interfaces support 1000 Mbps connections only. They do not support 10 Mbps or 100 Mbps connections.

The module supports a total of four active 1000 Mbps ports. One port out of each pair of copper and fiber connectors can be active at a time. For example, you can use either copper port connector 25 or fiber port connector 25, but not both at the same time. You can use a combination of fiber and copper connectors or all copper or all fiber connectors, as needed.

The module uses the following rules when selecting the copper or fiber connector to be the active connector for a port:

- If the module is active, the first connector you plug into the network becomes the active connector for the port.

- If both the copper and fiber connectors for the same port are attached to the network when the device boots or when you enable the port, the fiber connector takes precedence over the copper connector and will be the active connector for the port.
- If only one connector is attached to the network, and you then attach the other connector to the network, the connector that was first attached remains the active connector unless the link goes down, in which case the port fails over to the other connector, if the link on the other connector is up.

Although the fiber connector takes precedence over the copper connector, the device can boot over the network using either connector. For example, if you attach both the fiber and copper connectors for port 25 to the network and then boot the device using a BootP server, the first connector that receives a valid BootP reply from the server becomes the active connector for the duration of the boot process. After the device is booted, the fiber connector becomes the active connector. During this process, only the fiber link LED is lit, even if the copper port is used.

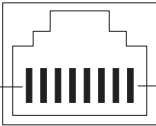
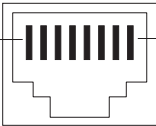
NOTE: If you plug both the copper and fiber connectors for the same port into the network, the remote ends of both links can be active even though the local end of the fiber link is inactive. This is normal; however, this condition can cause traffic flooding.

Gigabit Copper Interfaces

The J-F2404GMR4 provides four RJ-45 connectors for its Gigabit Copper interfaces. The connectors use Category5 (Cat5) cabling. The 16-port Gigabit Copper forwarding modules (J-B16GC, J-F16GC, and ML-16GC) provide 16 RJ-45 connectors for the 100/1000 interfaces. These connectors also use Category5 (Cat5) cabling. Figure D.4 shows the pin assignments and signalling for crossover connections on the Gigabit Copper (100/1000BaseT) ports.

NOTE: For comparison, the figure also shows 10BaseT pin and signal information.

Figure D.4 Crossover pin assignment and signalling for 100BaseTX and 1000BaseTX ports

Pin Assignment	10BaseT		100BaseTX and 1000BaseT	
	Pin Number	MDI-X ports	Pin Number	MDI-X ports
	1	RD+	1	RD+
	2	RD-	2	RD-
	3	TD+	3	TD+
	4	Not used	4	CMT
	5	Not used	5	CMT
	6	TD-	6	TD-
	7	Not used	7	CMT
	8	Not used	8	CMT

NOTE: The copper interfaces support autonegotiation but do not support half-duplex mode.

NOTE: The Gigabit Copper interfaces on the J-F2404GMR4 support 1000 Mbps connections only. They do not support 10 Mbps or 100 Mbps connections.

Forcing the Port Speed

The port's autonegotiation selects the correct port speed for the link. However, you can force the port speed to be either 100 Mbps or 1000 Mbps.

- To force the port to run at 1000 Mbps, set one of the link's ports to be the master for the link. To set a port as a Gigabit master port, enter the following command at the interface configuration level for the port:
speed-duplex 1000-master

- To force a Gigabit port to run at 100 Mbps on a link where both ends support 1000 Mbps, use a crossover cable to connect to the remote end of the link, and set each end of the link to 100 Mbps. You need to use the crossover cable in addition to setting the port speed on each side of the link to 100 Mbps.

Gigabit Fiber Interfaces

The J-F2404GMR4 provides four mini-GBIC connectors for 1000BaseSX or 1000BaseLX transceivers. The 16-port Fiber forwarding modules (J-B16Gx, J-F16Gx, and ML-16Gx) provide 16 mini-GBIC connectors for 1000BaseSX or 1000BaseLX transceivers.

The LEDs listed in Table D.3 provide status information for the fiber and copper Gigabit ports.

Table D.4: LEDs for 1000 Mbps Ports

LED	Position	State	Meaning
Link	Top	On	Port is connected.
		Off	No port connection exists.
Activity	Bottom	On	Traffic is being transmitted and received on that port.
		Off	No traffic is being transmitted.
		Blinking	Traffic is being transmitted and received on that port.

10/100 Ethernet Network Interfaces

Each 10/100 port on the J-F2404GMR4 modules has its own RJ-45 connector. Each connector supports a 10/100 Ethernet network segment on Cat5 cabling.

- Use a crossover cable to connect to another Layer 2 Switch or Layer 3 Switch. A crossover cable swaps the wires so that the send signal on one port connects to the receive signal on the other port, and so on.
- Use a straight-through cable to connect to an end station or server. A straight-through cable does not swap the wires.

For the pin assignments and signalling for crossover connections on the 10/100 ports, see Figure D.4.

The LEDs listed in Table D.5 provide status information for the 10/100 ports.

Table D.5: LEDs for 10/100 Mbps Ports


LED	Position	State	Meaning
Link/Activity	Left	On	Port is connected.
		Off	No port connection exists.
		Blinking	Traffic is being transmitted and received on that port.
FDX	Right	On	The port is operating at full-duplex.
		Off	The port is operating at half-duplex.

JetCore Gigabit Ethernet Forwarding Modules

J-BxG and J-FixG 8-Port Forwarding Modules

Figure D.5 shows the front panel of a JetCore 8-port Gigabit forwarding module.

Figure D.5 J-FixG JetCore Gigabit forwarding module




The JetCore Gigabit Ethernet forwarding modules provide eight mini-GBIC slots. You can insert any combination of 1000BaseSX or 1000BaseLX fiber connectors in to the slots. See “Gigabit Ethernet Network Interfaces” on page D-5.

Like the JetCore Gigabit Ethernet management modules, the Gigabit forwarding modules contain two IGCs that manage the ports. See “JetCore ASICs” on page D-3.

J-B16GC and J-F16GC 16-Port Forwarding Modules

Figure D.6 shows the front panel of a JetCore 16-port Gigabit forwarding module.

Figure D.6 J-F16GC JetCore Gigabit Copper forwarding module




The J-B16GC and J-F16GC modules provide 16 RJ-45 connectors for Cat5 cabling. You can connect each port to a 100 Mbps or 1000 Mbps segment. The ports automatically detect the speed of the network and configure themselves accordingly. You also can manually configure a port for 100 Mbps or 1000 Mbps. The ports also support automatic MDI/MDIX crossover.

The pin assignments and the status LEDs are the same as the ones for the 100 and 1000 Mbps ports on other Foundry modules. See Table D.3 on page D-6.

Figure D.7 shows the front panel of a J-F16Gx.

Figure D.7 J-F16Gx JetCore Gigabit Fiber forwarding module



The J-B16Gx and J-F16Gx modules provide 16 slots for mini-GBIC connectors. You can insert mini-GBICs for 1000BaseSX or 1000BaseLX fiber cables.

For information about the status LEDs, see Table D.3 on page D-6. All these modules use the same status LEDs for Gigabit Ethernet ports.

JetCore 10/100 Ethernet Forwarding Modules

24-Port 100BaseFX Forwarding Module

Software release 07.6.02 adds support for a new forwarding module, the JetCore 24-port Fast Ethernet 100Base-FX Fiber Module – part number J-24FX. Figure D.8 shows the front panel of a J-24FX module.

Figure D.8 J-24FX JetCore 24-Port 100BaseFX forwarding module



The J-24FX module provides 24 100BaseFX ports for connection to fiber-optic cable. The 100BaseFX ports are equipped with MT-RJ connectors and operate at 100 Mbps in full-duplex mode. A single Integrated Packet Controller (IPC) manages the ports.

Table D.6 describes the J-24FX module’s LEDs.


Table D.6: LEDs for 100BaseFX Ports

LED	Position	State	Meaning
Link	Top	On	Port is connected.
		Off	No port connection exists.
Activity	Bottom	On	Traffic is being transmitted and received on that port.
		Off	No traffic is being transmitted.
		Blinking	Traffic is being transmitted and received on that port.

J-B48E and J-FI48E 48-Port Enterprise Forwarding Modules

Figure D.9 shows the front panel of a JetCore 10/100 RJ-45 forwarding module. This module occupies two chassis slots.

Figure D.9 J-F48E JetCore 10/100 forwarding module



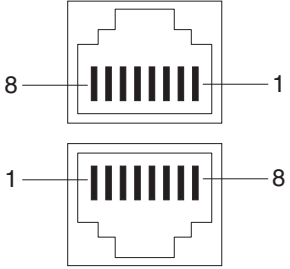
RJ-45 Interfaces

Each port on the J-FI48E module has its own RJ-45 connector. Each connector supports a 10/100 Ethernet network segment on Category 5 (Cat5) wire.

- Use a crossover cable to connect to another Layer 2 Switch or Layer 3 Switch. A crossover cable swaps the wires so that the send signal on one port connects to the receive signal on the other port, and so on.
- Use a straight-through cable to connect to an end station or server. A straight-through cable does not swap the wires.

Figure D.10 shows the pin assignments and signalling for crossover connections on the 10/100 ports.

Figure D.10 Crossover pin assignment and signalling for 10/100BaseTX ports

Pin Assignment	10BaseT		100BaseTX and 1000BaseT	
	Pin Number	MDI-X ports	Pin Number	MDI-X ports
	1	RD+	1	RD+
	2	RD-	2	RD-
	3	TD+	3	TD+
	4	Not used	4	CMT
	5	Not used	5	CMT
	6	TD-	6	TD-
	7	Not used	7	CMT
	8	Not used	8	CMT


J-B48E-A and J-F48E-A 48-Port Forwarding Modules

NOTE: Supported in Enterprise software releases 07.6.05 and later and in Service Provider software release 09.1.03 and later.

The J-B48E-A and J-F48E-A forwarding modules, available in software releases 07.6.05 and later, are identical in functionality to the J-B48E and J-FI48E forwarding modules. See “J-B48E and J-FI48E 48-Port Enterprise Forwarding Modules” on page D-10.

J-B48T and J-FI48T 48-Port Telco Forwarding Modules


Figure D.11 shows the front panel of a JetCore 10/100 RJ-21 forwarding module. This module occupies one chassis slot.

Figure D.11 J-F48T JetCore 10/100 forwarding module

RJ-21 Interfaces

The JetCore Telco modules provide four 50-pin connectors for attaching to 48 10/100 Ethernet segments. The connectors use the RJ-21 wiring standard, which uses four wires for each network segment. Each connector supports 12 segments. Figure D.12 shows an example of a Telco serial cable.

Figure D.12 RJ-21 Telco serial cable



To connect the JetCore module to the network, you can use a cable that terminates in another 50-pin connector or one that terminates in 12 RJ-45 connectors, depending on the patch panel you are using.

NOTE: Foundry does not provide the cables or patch panels. However, you can order cables and patch panels from Superior Module Products, www.superiormod.com.

Figure D.13 shows an example of a patch panel that accepts a 50-pin connector, and converts the signals to 12 RJ-45 sockets. Each of the RJ-45 sockets uses four signals per the RJ-21 wiring standard. You can use Cat5 cables with RJ-45 connectors to plug your network devices into the patch panel.

Figure D.13 Telco patch panel




Table D.7 shows the output signals on each JetCore Telco 50-pin connector. Notice that each 10/100 port uses four signals. Two of the signals are for transmit and the other two are for receive. Signals 25 and 50 are not used.

Table D.7: Output Signals for RJ-21

10/100 Port	Pin Number	Signal	Pin Number	Signal
1	1	RxD (-)	26	RxD (+)
	2	TxD (-)	27	TxD (+)
2	3	RxD (-)	28	RxD (+)
	4	TxD (-)	29	TxD (+)

Table D.7: Output Signals for RJ-21 (Continued)

10/100 Port	Pin Number	Signal	Pin Number	Signal
3	5	RxD (-)	30	RxD (+)
	6	TxD (-)	31	TxD (+)
4	7	RxD (-)	32	RxD (+)
	8	TxD (-)	33	TxD (+)
5	9	RxD (-)	34	RxD (+)
	10	TxD (-)	35	TxD (+)
6	11	RxD (-)	36	RxD (+)
	12	TxD (-)	37	TxD (+)
7	13	RxD (-)	38	RxD (+)
	14	TxD (-)	39	TxD (+)
8	15	RxD (-)	40	RxD (+)
	16	TxD (-)	41	TxD (+)
9	17	RxD (-)	42	RxD (+)
	18	TxD (-)	43	TxD (+)
10	19	RxD (-)	44	RxD (+)
	20	TxD (-)	45	TxD (+)
11	21	RxD (-)	46	RxD (+)
	22	TxD (-)	47	TxD (+)
12	23	RxD (-)	48	RxD (+)
	24	TxD (-)	49	TxD (+)
N/A	25	Not used	50	Not used

J-B48T-A and J-F48T-A 48-Port Forwarding Modules

NOTE: Supported in software releases 07.6.05 and later and in Service Provider software release 09.1.03 and later.

The J-B48T-A and J-F48T-A forwarding modules, available in software releases 07.6.05 and later, are identical in functionality to the J-B48T and J-FI48T forwarding modules. See “J-B48T and J-FI48T 48-Port Telco Forwarding Modules” on page D-11.

Configuration Considerations

- BigIron JetCore modules do not require a new chassis. You can use the modules in your installed chassis.
- FastIron JetCore modules require a FastIron 400, FastIron 800, or FastIron 1500.
- NetIron JetCore modules require a NetIron Metro router model N400, N800, or N1500.

- You cannot use JetCore modules and IronCore (non-JetCore) modules in the same chassis.
- You cannot use a JetCore module from one chassis type (BigIron, FastIron, or NetIron) in another chassis type. For example, you cannot use a BigIron JetCore module in a FastIron JetCore chassis.

Appendix E

Cautions and Warnings

The cautions and warnings that appear in this manual are listed below in English, German, French, and Spanish.

Cautions

A caution calls your attention to a possible hazard that can damage equipment.

"Vorsicht" weist auf eine mögliche Beschädigung des Geräts hin. Sie finden die folgenden Vorsichtshinweise in diesem Handbuch.

Une mise en garde attire votre attention sur un risque possible d'endommagement de l'équipement. Ci-dessous, vous trouverez les mises en garde utilisées dans ce manuel.

Un mensaje de precaución le advierte sobre un posible peligro que pueda dañar el equipo. Las siguientes son precauciones utilizadas en este manual.

CAUTION:	All devices with DC power supplies are intended for installation in restricted access areas only. A restricted access area is where access can be gained only by service personnel through the use of a special tool, lock and key, or other means of security, and is controlled by the authority responsible for the location.
VORSICHT:	Alle Geräte mit DC-Netzteil sind nur für die Installation in Bereichen mit beschränktem Zugang gedacht. Ein Bereich mit beschränktem Zugang ist ein Bereich, zu dem nur Wartungspersonal mit Spezialwerkzeug, Schlüssel oder anderen Sicherheitsvorrichtungen Zugang hat. Dieser Zugang wird von für den Bereich zuständigen Personen überwacht.
MISE EN GARDE:	Tous les dispositifs avec bloc d'alimentation C.C. sont conçus pour l'installation dans des zones à accès réglementé uniquement. Une zone à accès réglementé est une zone dont l'accès n'est possible qu'au personnel de service utilisant un verrou, une clé ou un outil spécial, ou d'autres moyens de sécurité, et qui est contrôlée par les autorités responsables du site.
PRECAUCIÓN:	Todos los instrumentos con suministros de corriente continua han sido diseñados únicamente para instalación en áreas restringidas. Se entiende como área de acceso restringido un lugar al que solo puede acceder personal de servicio mediante el uso de una herramienta especial, llave y cerrojo u otro medio de seguridad similar, y que esté controlado por la autoridad responsable de esa ubicación.

CAUTION: By default, the delete option deletes all files on the flash card. Make sure you specify the files you want to delete.

VORSICHT: Gemäß Vorgabe löscht die Option "Delete" (Löschen) alle Dateien auf der Flash-Karte. Stellen Sie sicher, dass Sie die zu löschenden Dateien angeben.

MISE EN GARDE: Par défaut, l'option de suppression supprime tous les fichiers de la carte mémoire. Assurez-vous de spécifier les fichiers que vous voulez supprimer.

PRECAUCIÓN: Por defecto, la opción de anular anula todos los archivos de la tarjeta flash. Verifique que especifica los archivos que quiere anular.

CAUTION: Carefully follow the mechanical guides on each side of the power supply slot and make sure the power supply is properly inserted in the guides. Never insert the power supply upside down.

VORSICHT: Beachten Sie mechanischen Führungen an jeder Seite des Netzteils, das ordnungsgemäß in die Führungen gesteckt werden muss. Das Netzteil darf niemals umgedreht eingesteckt werden.

MISE EN GARDE: Suivez attentivement les repères mécaniques de chaque côté du slot du bloc d'alimentation et assurez-vous que le bloc d'alimentation est bien inséré dans les repères. N'insérez jamais le bloc d'alimentation à l'envers.

PRECAUCIÓN: Siga cuidadosamente las guías mecánicas de cada lado de la ranura del suministro de energía y verifique que el suministro de energía está insertado correctamente en las guías. No inserte nunca el suministro de energía de manera invertida.

CAUTION: Do not add or remove a flash card while a file operation involving the flash card's slot is in progress. Doing so can result in corruption of the flash card. If this occurs, you may need to reformat the flash card to make it usable again. Reformatting the card erases all data stored on the card.

VORSICHT: Eine Flash-Karte darf nur dann eingesteckt oder herausgenommen werden, wenn keine Dateifunktion läuft, die der Flash-Karte bedarf. Wenn dies nicht beachtet wird, kann dies zur Korruption der Flash-Karte führen. Die Karte kann dann erst nach Neuformatierung wieder benutzt werden. Bei Neuformatierung gehen alle auf der Karte gespeicherten Daten verloren.

MISE EN GARDE: N'ajoutez pas ou ne supprimez pas une carte mémoire au cours d'une opération de fichier dans laquelle le slot de carte mémoire est impliqué. Vous risquez sinon de corrompre la carte mémoire. Si cela se produit, vous devrez peut-être reformater la carte mémoire pour qu'elle soit à nouveau utilisable. Le reformatage de la carte efface toutes les données qui y sont stockées.

PRECAUCIÓN: No añada ni quite una tarjeta flash mientras una operación de archivo que conlleve el uso de una ranura de tarjeta flash se encuentre en uso. De hacerlo así se podría dar lugar a la corrupción de la tarjeta flash. Si esto ocurriera, podría ser necesario que vuelva a formatear la tarjeta flash para hacer que vuelva sea utilizable. Cuando se formatea la tarjeta se borran todos los datos almacenados en la tarjeta.

CAUTION:	Do not install the device in an environment where the operating ambient temperature might exceed 40° C (104° F).
VORSICHT:	Das Gerät darf nicht in einer Umgebung mit einer Umgebungsbetriebstemperatur von über 40° C (104° F) installiert werden.
MISE EN GARDE:	N'installez pas le dispositif dans un environnement où la température d'exploitation ambiante risque de dépasser 40° C (104° F).
PRECAUCIÓN:	No instale el instrumento en un entorno en el que la temperatura ambiente de operación pueda exceder los 40°C (104° F).

CAUTION:	Ensure that the device does not overload the power circuits, wiring, and over-current protection. To determine the possibility of overloading the supply circuits, add the ampere (amp) ratings of all devices installed on the same circuit as the device. Compare this total with the rating limit for the circuit. The maximum ampere ratings are usually printed on the devices near the input power connectors.
VORSICHT:	Stromkreise, Verdrahtung und Überlastschutz dürfen nicht durch das Gerät überbelastet werden. Addieren Sie die Nennstromleistung (in Ampere) aller Geräte, die am selben Stromkreis wie das Gerät installiert sind. Somit können Sie feststellen, ob die Gefahr einer Überbelastung der Versorgungsstromkreise vorliegt. Vergleichen Sie diese Summe mit der Nennstromgrenze des Stromkreises. Die Höchstnennströme (in Ampere) stehen normalerweise auf der Geräterückseite neben den Eingangsstromanschlüssen.
MISE EN GARDE:	Assurez-vous que le dispositif ne risque pas de surcharger les circuits d'alimentation, le câblage et la protection de surintensité. Pour déterminer le risque de surcharge des circuits d'alimentation, additionnez l'intensité nominale (ampères) de tous les dispositifs installés sur le même circuit que le dispositif en question. Comparez alors ce total avec la limite de charge du circuit. L'intensité nominale maximum en ampères est généralement imprimée sur chaque dispositif près des connecteurs d'entrée d'alimentation.
PRECAUCIÓN:	Verifique que el instrumento no sobrecargue los circuitos de corriente, el cableado y la protección para sobrecargas. Para determinar la posibilidad de sobrecarga en los circuitos de suministros, añada las capacidades nominales de corriente (amp) de todos los instrumentos instalados en el mismo circuito que el instrumento. Compare esta suma con el límite nominal para el circuito. Las capacidades nominales de corriente máximas están generalmente impresas en los instrumentos, cerca de los conectores de corriente de entrada.

CAUTION:	For DC power supplies in 15-slot Chassis devices (part number RPS4DC), use a grounding wire of at least 10 American Wire Gauge (AWG). For DC power supplies in 4-slot and 8-slot Chassis devices (part number RPS3DC), use at least 14 AWG. For DC power supplies in the FastIron 4802 (part number RPS5DC), use at least 14 AWG.
VORSICHT:	Benutzen Sie für DC-Netzteile in 15-Steckplatz-Chassisgeräte (Teilenummer RPS4DC) einen Erdungsdraht mit wenigstens 10 AWG (American Wire Gauge). Benutzen Sie für DC-Netzteile in 4- bis 8-Steckplatz-Chassisgeräte (Teilenummer RPS3DC) einen Erdungsdraht mit wenigstens 14 AWG. < DC-Netzteile im FastIron 4802 (Teilenummer RPS5DC) erfordern wenigstens 14 AWG.
MISE EN GARDE:	Pour les blocs d'alimentation C.C. des dispositifs en châssis à 15 slots (numéro de pièce RPS4DC), utilisez un fil de mise à la terre de calibre 10 AWG (American Wire Gauge) minimum. Pour les blocs d'alimentation C.C. des dispositifs en châssis à 4 et 8 slots (numéro de pièce RPS3DC), utilisez un fil de mise à la terre de calibre 14 AWG minimum. Pour les

blocs d'alimentation du FastIron 4802 (numéro de pièce RPS5DC), utilisez un fil de mise à la terre de calibre 14 AWG minimum.

PRECAUCIÓN: Para suministros de corriente continua en los instrumentos del chasis de 15 ranuras (pieza número RPS4DC), use un hilo de tierra de al menos 10 AWG (American Wire Gauge [calibración de hilos americana]). Para suministros de corriente continua en instrumentos de chasis de 4 y 8 ranuras (pieza número RPS3DC), utilice al menos 14 AWG. Para suministros de corriente continua en el FastIron 4802 (pieza número RPS5DC), use al menos 14 AWG.

CAUTION: For the DC input circuit to a 15-slot Chassis device (DC power supply part number RPS4DC), make sure there is a 30-amp circuit breaker on the input to the power supply.

VORSICHT: Für den DC-Eingangstromkreis zu einem 15-Steckplatz-Chassisgerät (DC-Netzteil-Teilenummer RPS4DC) müssen Sie sicherstellen, dass sich am Eingang zum Netzteil ein 13 A-Überlastschalter befindet.

MISE EN GARDE: Pour le circuit d'alimentation C.C. d'un dispositif en châssis à 15 slots (numéro de pièce du bloc d'alimentation C.C. RPS4DC), assurez-vous de la présence d'un disjoncteur de 30 ampères sur l'entrée vers le bloc d'alimentation.

PRECAUCIÓN: Para el circuito de entrada de corriente continua a un instrumento de chasis de 15 ranuras (suministro de corriente continua, pieza número RPS4DC), asegúrese que haya un cortacircuitos de 30 amperios a la entrada del suministro de energía.

CAUTION: For the DC input circuit to a FastIron 4802 (DC power supply part number RPS5DC), make sure there is a 10-amp circuit breaker when installed in the end system.

VORSICHT: Für den DC-Eingangstromkreis zu einem FastIron 4802 (DC-Netzteil-Teilenummer RPS5DC) müssen Sie sicherstellen, dass ein 10 A-Überlastschalter im Endsystem installiert wird.

MISE EN GARDE: Pour le circuit d'alimentation C.C. d'un FastIron 4802 (numéro de pièce du bloc d'alimentation C.C. RPS5DC), assurez-vous de la présence d'un disjoncteur de 10 ampères lors de l'installation sur le système d'extrémité.

PRECAUCIÓN: Para el circuito de entrada de corriente continua a un FastIron 4802 (suministro de corriente continua, pieza número RPS5DC), asegúrese de haya un cortacircuitos de 10 amperios cuando esté instalado en el sistema final.

CAUTION: Make sure the air flow around the front, sides, and back of the device is not restricted.

VORSICHT: Stellen Sie sicher, dass an der Vorderseite, den Seiten und an der Rückseite der Luftstrom nicht behindert wird.

MISE EN GARDE: Vérifiez que rien ne restreint la circulation d'air devant, derrière et sur les côtés du dispositif et qu'elle peut se faire librement.

PRECAUCIÓN: Asegúrese de que el flujo de aire en las inmediaciones de las partes anterior, laterales y posterior del instrumento no esté restringido.

- CAUTION:** Make sure the flash card is empty or does not contain files you want to keep. Formatting a flash card completely erases all files on the card.
- VORSICHT:** Stellen Sie sicher, dass die Flash-Karte leer ist oder keine Dateien auf ihr gespeichert sind, die Sie behalten möchten. Die Formatierung einer Flash-Karte löscht alle Dateien auf der Karte.
- MISE EN GARDE:** Vérifiez que la carte mémoire est vide ou ne contient pas de fichiers que vous voulez conserver. Le reformatage de la carte mémoire efface tous les fichiers qui s'y trouvent.
- PRECAUCIÓN:** Verifique que la tarjeta flash esté vacía o que no contenga archivos que desee conservar. Al formatear una tarjeta flash todos los archivos de ésta se borran.
-

- CAUTION:** Make sure you insert the power supply right-side up. It is possible to insert the supply upside down, although the supply will not engage with the power backplane when upside down. The power supply is right-side up when the power connector is on the left and the fan vent is on the right.
- VORSICHT:** Sicher Sie sicher, dass Sie das Netzteil mit der richtigen Seite nach oben weisend einstecken. Man kann die Karte auch umgekehrt einstecken. Allerdings rastet das umgekehrte Netzteil nicht in die Netzstrom-Rückwandplatine ein. Die rechte Seite des Netzteils weist nach oben, wenn sich der Stromanschlussstecker links und der Ventilatorschlitz rechts befindet.
- MISE EN GARDE:** Assurez-vous d'insérer le bloc d'alimentation dans le bon sens. Il est possible de l'insérer " la tête en bas ", mais le bloc d'alimentation ne s'enclenchera pas dans la face arrière d'alimentation s'il est inséré à l'envers. Le bloc d'alimentation est dans le bon sens lorsque le connecteur se trouve sur le côté gauche et le ventilateur sur la droite.
- PRECAUCIÓN:** Verifique que inserta el suministro de corriente con la cara correcta hacia arriba. Es posible insertar el suministro hacia abajo, pese a que este no se conectará con el enchufe posterior de esta forma. El suministro de potencia estará con la cara correcta hacia arriba cuando el conector de corriente quede a la izquierda y la abertura del ventilador queda a la derecha.
-

- CAUTION:** Never leave tools inside the chassis.
- VORSICHT:** Lassen Sie keine Werkzeuge im Chassis zurück.
- MISE EN GARDE:** Ne laissez jamais d'outils à l'intérieur du châssis.
- PRECAUCIÓN:** No deje nunca herramientas en el interior del chasis.
-

- CAUTION:** Once you start the formatting process, you cannot stop it. Even if you enter CTRL-C to stop the CLI output and a new prompt appears, the formatting continues. Make sure you want to format the card before you enter the command.
- VORSICHT:** Wenn Sie mit dem Formattieren beginnen, können Sie diesen Prozess nicht anhalten. Selbst wenn zum Anhalten der CLI-Ausgabe Strg-C drücken und eine neue Aufforderung gezeigt wird, wird mit dem Formattieren fortgefahren. Stellen Sie sicher, dass Sie die Karte formattieren wollen, bevor Sie den Befehl eingeben.
- MISE EN GARDE:** Une fois le processus de formatage commencé, vous ne pouvez pas l'interrompre. Même si vous appuyez sur CTRL-C pour arrêter la sortie CLI et si une nouvelle invite apparaît, le formatage continue. Soyez bien sûr de vouloir formater la carte avant d'entrer la commande.
- PRECAUCIÓN:** Una vez que empieza con el proceso de formateado, no se puede detener. Incluso si pulsa CTRL-C para detener la salida de CLI y aparece un nuevo indicador, el formateado continuará. Esté seguro que desea formatear la tarjeta antes de introducir el comando.
-

- CAUTION:** Remove the power cord from a power supply before you install it in or remove it from the device. Otherwise, the power supply or the device could be damaged as a result. (The device can be running while a power supply is being installed or removed, but the power supply itself should not be connected to a power source.)
- VORSICHT:** Nehmen Sie vor dem Anschließen oder Abtrennen des Geräts das Stromkabel vom Netzteil ab. Ansonsten könnten das Netzteil oder das Gerät beschädigt werden. (Das Gerät kann während des Anschließens oder Annehmens des Netzteils laufen. Nur das Netzteil sollte nicht an eine Stromquelle angeschlossen sein.)
- MISE EN GARDE:** Enlevez le cordon d'alimentation d'un bloc d'alimentation avant de l'installer ou de l'enlever du dispositif. Sinon, le bloc d'alimentation ou le dispositif risque d'être endommagé. (Le dispositif peut être en train de fonctionner lorsque vous installez ou enlevez un bloc d'alimentation, mais le bloc d'alimentation lui-même ne doit pas être connecté à une source d'alimentation.)
- PRECAUCIÓN:** Retire el cordón de corriente del suministro de corriente antes de instalarlo o retirarlo del instrumento. De no hacerse así, el suministro de corriente o el instrumento podrían resultar dañados. (El instrumento puede estar encendido mientras se instala o retira un suministro de corriente, pero el suministro de corriente en sí no deberá conectado a la corriente).
-

- CAUTION:** The software does not have an undelete option. Make sure you really want to delete the file.
- VORSICHT:** Die Software verfügt über keine Option "Undelete" (Löschung rückgängig machen). Stellen Sie sicher, dass Sie die Datei wirklich löschen wollen.
- MISE EN GARDE:** Le logiciel n'a pas d'option permettant d'annuler la suppression. Soyez donc bien sûr de vouloir supprimer le fichier.
- PRECAUCIÓN:** El software no dispone de una opción de recuperar lo anulado. Esté plenamente seguro de que quiere anular el archivo.
-

- CAUTION:** To provide additional safety and proper airflow to the device, make sure that slot cover plates are installed on all chassis slots that do not have either a module or power supply installed.
- VORSICHT:** Für mehr Sicherheit und eine bessere Luftzufuhr zum Gerät müssen Sie sicherstellen, dass die Abdeckplatten für die Steckplätze an allen Chassissteckplätzen montiert sind und dass in diesen keine Module oder Netzteile installiert sind.
- MISE EN GARDE:** Pour fournir une sécurité supplémentaire et une circulation d'air adéquate pour le dispositif, vérifiez que des caches de slots sont installés sur tous les slots du châssis dans lesquels un module ou un bloc d'alimentation n'est pas installé.
- PRECAUCIÓN:** Para proporcionar seguridad adicional y un flujo de aire apropiado al instrumento, verifique que las placas de cierre de las ranuras estén instaladas en todas las ranuras del chasis que no tengan un módulo o un suministro de corriente instalado.
-

- CAUTION:** Use at least two separate branch circuits for the power. This provides redundancy in case one of the circuits fails.
- VORSICHT:** Verwenden Sie wenigstens zwei getrennte Stromkreise für die Stromversorgung. Somit steht Ihnen im Fall des Ausfalls eines Stromkreises ein Ersatzstromkreis zur Verfügung.
- MISE EN GARDE:** Utilisez au moins deux circuits de dérivation différents pour l'alimentation. Ainsi, il y aura un circuit redondant en cas de panne d'un des circuits.
- PRECAUCIÓN:** Use al menos dos circuitos derivados separados para la corriente. Esto proporciona redundancia en el caso que uno de los circuitos falle.
-

CAUTION: Use the erase startup-config command only for new systems. If you enter this command on a system you have already configured, the command erases the configuration. If you accidentally do erase the configuration on a configured system, enter the write memory command to save the running configuration to the startup-config file.

VORSICHT: Verwenden Sie den Befehl "Erase startup-config" (Löschen Startup-Konfig) nur für neue Systeme. Wenn Sie diesen Befehl in ein bereits konfiguriertes System eingeben, löscht der Befehl die Konfiguration. Falls Sie aus Versehen die Konfiguration eines bereits konfigurierten Systems löschen, geben Sie den Befehl "Write Memory" (Speicher schreiben) ein, um die laufende Konfiguration in der Startup-Konfig-Datei zu speichern.

MISE EN GARDE: N'utilisez la commande erase startup-config que pour les nouveaux systèmes. Si vous entrez cette commande sur un système que vous avez déjà configuré, elle efface la configuration. Si vous effacez la configuration par accident sur un système configuré, entrez la commande write memory pour enregistrer la configuration actuelle dans le fichier startup-config.

PRECAUCIÓN: Use el comando erase startup-config (borrar configuración de inicio) para sistemas nuevos solamente. Si usted introduce este comando en un sistema que ya ha configurado, el comando borrará la configuración. Si usted borra accidentalmente la configuración en un sistema ya configurado, introduzca el comando write memory (escribir memoria) para guardar la configuración en ejecución en el archivo startup-config.

CAUTION: When you connect a fan cable to a fan connector on the backplane, make sure the red wire in the connector is on the right side (for horizontally oriented connectors) or facing down (for vertically oriented connectors). If you accidentally reverse the wires, the fan will not operate.

Also, make sure the fan cable connector is seated over all three pins on the backplane connector.

VORSICHT: Wenn Sie einen Ventilator an den Ventilatoranschlussstecker auf der Rückplatine anschließen, müssen Sie sicherstellen, dass sich der rote Draht im Anschlussstecker rechts befindet (bei waagrecht angeordneten Anschlusssteckern) oder nach unten weist (bei senkrecht angeordneten Anschlusssteckern). Wenn Sie die Drähte aus Versehen vertauschen, läuft der Ventilator nicht.

Stellen Sie auch sicher, dass der Anschlussstecker des Ventilatorkabels auf allen drei Stiften am Rückplatinen-Anschlussstecker sitzt.

MISE EN GARDE: Lorsque vous connectez le câble d'un ventilateur à un connecteur de ventilateur sur la face arrière, vérifiez que le fil rouge dans le connecteur est bien sur le côté droit (pour les connecteurs orientés horizontalement) ou vers le bas (pour les connecteurs orientés verticalement). Si vous inversez les fils, le ventilateur ne fonctionnera pas.

De plus, vérifiez que le connecteur du câble du ventilateur est bien en place sur les trois broches du connecteur de la face arrière.

PRECAUCIÓN: Cuando conecte un cable de ventilador a un conector de ventilador en el enchufe posterior, verifique que el cable rojo del conector está en el lado derecho (para conectores orientados horizontalmente) u orientado hacia abajo (para conectores orientados verticalmente). Si usted invierte los cables accidentalmente, el ventilador no funcionará.

Asimismo, verifique que el conector del cable del ventilador queda asentado sobre las tres clavijas en el conector del enchufe posterior.

Warnings

A warning calls your attention to a possible hazard that can cause injury or death. The following are the warnings used in this manual.

"Achtung" weist auf eine mögliche Gefährdung hin, die zu Verletzungen oder Tod führen können. Sie finden die folgenden Warnhinweise in diesem Handbuch:

Un avertissement attire votre attention sur un risque possible de blessure ou de décès. Ci-dessous, vous trouverez les avertissements utilisés dans ce manuel.

Una advertencia le llama la atención sobre cualquier posible peligro que pueda ocasionar daños personales o la muerte. A continuación se dan las advertencias utilizadas en este manual.

WARNING:	The procedures in this manual are for qualified service personnel.
ACHTUNG:	Die Verfahren in diesem Handbuch sind nur für qualifiziertes Wartungspersonal gedacht.
AVERTISSEMENT:	Les procédures décrites dans ce manuel doivent être effectuées par le personnel de service qualifié uniquement.
ADVERTENCIA:	Los procedimientos de este manual se han hecho para personal de servicio cualificado.

WARNING:	All fiber-optic interfaces except LHB interfaces use Class 1 Lasers.
ACHTUNG:	Alle Glasfaser-Schnittstellen verwenden Laser der Klasse 1.
AVERTISSEMENT:	Toutes les interfaces en fibres optiques utilisent des lasers de classe 1.
ADVERTENCIA:	Todas las interfaces de fibra óptica usan Láser de Clase 1.

WARNING:	Before beginning the installation, see the precautions in "Power Precautions" on page 2-4.
ACHTUNG:	Vor der Installation siehe Vorsichtsmaßnahmen unter " Power Precautions " (Vorsichtsmaßnahmen in Bezug auf elektrische Ablagen) auf den Seiten 2 - 4.
AVERTISSEMENT:	Avant de commencer l'installation, consultez les précautions décrites dans " Power Precautions " (Précautions quant à l'alimentation), pages 2-4.
ADVERTENCIA:	Antes de comenzar la instalación, consulte las precauciones en la sección " Power Precautions " (Precauciones sobre corriente) que se encuentra en las páginas 2-4.

WARNING:	Disconnect the power cord from all power sources to completely remove power from the device.
ACHTUNG:	Ziehen Sie das Stromkabel aus allen Stromquellen, um sicherzustellen, dass dem Gerät kein Strom zugeführt wird.
AVERTISSEMENT:	Débranchez le cordon d'alimentation de toutes les sources d'alimentation pour couper complètement l'alimentation du dispositif.
ADVERTENCIA:	Para desconectar completamente la corriente del instrumento, desconecte el cordón de corriente de todas las fuentes de corriente.

WARNING:	Do not lift the 15-slot chassis using the lifting handles unless the chassis is empty. Remove the power supplies and interface modules before lifting the chassis.
ACHTUNG:	Sie dürfen das 15-Steckplatz-Chassis nur dann an den Hebegriffen anheben, wenn das Chassis leer ist. Trennen Sie die Netzteile und Schnittstellenmodule vor dem Anheben des Chassis ab.
AVERTISSEMENT:	Ne soulevez le châssis à 15 slots à l'aide des poignées de levage que si le châssis est vide. Enlevez les blocs d'alimentation et les modules d'interface avant de soulever le châssis.

ADVERTENCIA: No alce el chasis de 15 ranuras usando las asas de alzado a menos que el chasis esté vacío. Retire los suministros de corriente y los módulos de interfaz antes de alzar el chasis.

WARNING: Do not use the handles on the power supply units to lift or carry a Layer 3 Switch.

ACHTUNG: Die Griffe an den Netzteilen dürfen nicht zum Anheben oder Tragen eines Chassisgeräts verwendet werden.

AVERTISSEMENT: N'utilisez pas les poignées des unités de bloc d'alimentation pour soulever ou porter un dispositif en châssis.

ADVERTENCIA: No use las asas de las unidades de suministro de corriente para alzar o transportar un instrumento de chasis.

WARNING: If the installation requires a different power cord than the one supplied with the device, make sure you use a power cord displaying the mark of the safety agency that defines the regulations for power cords in your country. The mark is your assurance that the power cord can be used safely with the device.

ACHTUNG: Falls für die Installation ein anderes Stromkabel erforderlich ist (wenn das mit dem Gerät gelieferte Kabel nicht passt), müssen Sie sicherstellen, dass Sie ein Stromkabel mit dem Siegel einer Sicherheitsbehörde verwenden, die für die Zertifizierung von Stromkabeln in Ihrem Land zuständig ist. Das Siegel ist Ihre Garantie, dass das Stromkabel sicher mit Ihrem Gerät verwendet werden kann.

AVERTISSEMENT: Si l'installation nécessite un cordon d'alimentation autre que celui fourni avec le dispositif, assurez-vous d'utiliser un cordon d'alimentation portant la marque de l'organisation responsable de la sécurité qui définit les normes et réglementations pour les cordons d'alimentation dans votre pays. Cette marque vous assure que vous pouvez utiliser le cordon d'alimentation avec le dispositif en toute sécurité.

ADVERTENCIA: Si la instalación requiere un cordón de corriente distinto al que se ha suministrado con el instrumento, verifique que usa un cordón de corriente que venga con la marca de la agencia de seguridad que defina las regulaciones para cordones de corriente en su país. Esta marca será su garantía de que el cordón de corriente puede ser utilizado con seguridad con el instrumento.

WARNING: Make sure that the power source circuits are properly grounded, then use the power cord supplied with the device to connect it to the power source.

ACHTUNG: Stellen Sie sicher, dass die Stromkreise ordnungsgemäß geerdet sind. Benutzen Sie dann das mit dem Gerät gelieferte Stromkabel, um es an die Stromquelle anzuschließen.

AVERTISSEMENT: Vérifiez que les circuits de sources d'alimentation sont bien mis à la terre, puis utilisez le cordon d'alimentation fourni avec le dispositif pour le connecter à la source d'alimentation.

ADVERTENCIA: Verifique que circuitos de la fuente de corriente están conectados a tierra correctamente; luego use el cordón de potencia suministrado con el instrumento para conectarlo a la fuente de corriente.

WARNING: Make sure the rack or cabinet housing the device is adequately secured to prevent it from becoming unstable or falling over.

ACHTUNG: Stellen Sie sicher, dass das Gestell oder der Schrank für die Unterbringung des Geräts auf angemessene Weise gesichert ist, so dass das Gestell oder der Schrank nicht wackeln oder umfallen kann.

AVERTISSEMENT: Vérifiez que le bâti ou le support abritant le dispositif est bien fixé afin qu'il ne devienne pas instable ou qu'il ne risque pas de tomber.

ADVERTENCIA: Verifique que el bastidor o armario que alberga el instrumento está asegurado correctamente para evitar que pueda hacerse inestable o que caiga.

WARNING: Mount the devices you install in a rack or cabinet as low as possible. Place the heaviest device at the bottom and progressively place lighter devices above.

ACHTUNG: Montieren Sie die Geräte im Gestell oder Schrank so tief wie möglich. Platzieren Sie das schwerste Gerät ganz unten, während leichtere Geräte je nach Gewicht (je schwerer desto tiefer) darüber untergebracht werden.

AVERTISSEMENT: Montez les dispositifs que vous installez dans un bâti ou support aussi bas que possible. Placez le dispositif le plus lourd en bas et le plus léger en haut, en plaçant tous les dispositifs progressivement de bas en haut du plus lourd au plus léger.

ADVERTENCIA: Monte los instrumentos que instale en un bastidor o armario lo más bajos posible. Ponga el instrumento más pesado en la parte inferior y los instrumentos progresivamente más livianos más arriba.

WARNING: Power supplies are hot swappable. However, Foundry Networks recommends that you disconnect the power supply from AC power before installing or removing the supply. The device can be running while a power supply is being installed or removed, but the power supply itself should not be connected to a power source. Otherwise, you could be injured or the power supply or other parts of the device could be damaged.

ACHTUNG: Netzteile können unter Strom stehend ausgetauscht werden. Allerdings empfiehlt Foundry Networks, dass Sie das Netzteil vom Netzstrom abtrennen, bevor Sie das Netzteil anschließen oder abtrennen. Das Gerät kann während des Anschließens oder Abnehmens des Netzteils laufen. Nur das Netzteil sollte nicht an eine Stromquelle angeschlossen sein. Ansonsten können Sie verletzt oder das Netzteil bzw. andere Geräteteile beschädigt werden.

AVERTISSEMENT: Les blocs d'alimentation peuvent être changés à chaud. Cependant, Foundry Networks vous conseille de débrancher le bloc d'alimentation de l'alimentation C.A. avant d'installer ou d'enlever le bloc d'alimentation. Le dispositif peut être en cours de fonctionnement pendant que vous installez ou enlevez un bloc d'alimentation, mais le bloc d'alimentation lui-même ne doit pas être connecté à une source d'alimentation. Sinon, vous risquez d'être blessé ou le bloc d'alimentation ou d'autres pièces du dispositif risquent d'être endommagés.

ADVERTENCIA: Los suministros de corriente pueden intercambiarse sin necesidad de ajustes. No obstante, Foundry Networks recomienda que desconecte el suministro de corriente de la toma de corriente alterna antes de instalar o retirar el suministro. El instrumento puede estar activado cuando se esté instalando o retirando un suministro de corriente, pero el suministro de corriente en sí no deberá estar conectado a la fuente de corriente. De no hacerlo así, podría sufrir daños personales o el suministro de corriente u otras piezas podrían resultar dañadas.

WARNING: The s are very heavy, especially when fully populated with modules and power supplies. TWO OR MORE PEOPLE ARE REQUIRED WHEN LIFTING, HANDLING, OR MOUNTING THESE DEVICES.

ACHTUNG: Die Chassisgeräte sind sehr schwer. Dies gilt insbesondere, wenn sie mit vielen Modulen und Netzteilen bestückt sind. FÜR DAS ANHEBEN, TRANSPORTIEREN ODER MONTIEREN DIESER GERÄTE SIND WENIGSTENS ZWEI PERSONEN ERFORDERLICH.

AVERTISSEMENT: Les dispositifs en châssis sont très lourds, surtout s'ils sont entièrement remplis de modules et de blocs d'alimentation. POUR SOULEVER, MANIPULER OU MONTER CES DISPOSITIFS, DEUX PERSONNES MINIMUM SONT NÉCESSAIRES.

ADVERTENCIA: Los instrumentos del chasis son más pesados, especialmente cuando están muy cargados con módulos y suministros de corriente. SE REQUERIRÁN DOS O MÁS PERSONAS CUANDO ESTOS INSTRUMENTOS SE VAYAN A ALZAR, MANEJAR O MONTAR.

- WARNING:** To avoid risk of shock, do not attach the clip end to the air flow panel of the power supply.
- ACHTUNG:** Das Klemmemende darf nicht an die Belüftungsplatte des Netzteils angeschlossen werden. Andernfalls setzen Sie sich dem Risiko eines elektrischen Schlags aus.
- AVERTISSEMENT:** Pour éviter le risque de choc électrique, n'attachez pas l'extrémité du clip au panneau de circulation d'air du bloc d'alimentation.
- ADVERTENCIA:** Para evitar riesgos de electrocución, no acople el extremo del clip al panel de flujo de aire del suministro de corriente.
-

- WARNING:** You can lift the 4-slot and 8-slot s when they contain modules and power supplies. However, fully populated chassis are heavy. TWO OR MORE PEOPLE ARE REQUIRED WHEN LIFTING, HANDLING, OR MOUNTING THESE DEVICES.
- ACHTUNG:** Sie können ein 4-Steckplatz- und 8-Steckplatz-Chassis anheben, das mit Modulen und Netzteilen bestückt ist. Allerdings sind voll bestückte Chassis schwer. FÜR DAS ANHEBEN, TRANSPORTIEREN ODER MONTIEREN DIESER GERÄTE SIND WENIGSTENS ZWEI PERSONEN ERFORDERLICH.
- AVERTISSEMENT:** Vous pouvez soulever les dispositifs en châssis à 4 ou 8 slots lorsqu'ils contiennent des modules et des blocs d'alimentation. Cependant, les châssis sont lourds quand ils sont entièrement remplis. POUR SOULEVER, MANIPULER OU MONTER CES DISPOSITIFS, DEUX PERSONNES MINIMUM SONT NÉCESSAIRES.
- ADVERTENCIA:** Puede alzar los instrumentos de chasis de 4 y 8 ranuras cuando contengan módulos y suministros de corriente. Sin embargo, los chasis muy concurridos son pesados. SE REQUERIRÁN DOS O MÁS PERSONAS CUANDO ESTOS INSTRUMENTOS SE VAYAN A ALZAR, MANEJAR O MONTAR.
-

